

Project Operating Guide

Human accountability and AI assistance across multidisciplinary projects

Creation, research, revision, execution, verification, and release

Copyright and License

© 2026 Marius Tanase.

Project Operating Guide is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0).

<https://creativecommons.org/licenses/by/4.0/>

If you find this guide useful, consider supporting its creator's work with a voluntary contribution at www.mariustanase.org/support/. To explore more projects, access <https://www.mariustanase.org/resources/>. No contribution is required to use the guide under CC BY 4.0.

Master rule. Define the purpose, authority and scope of the next activity. Use qualified inputs, accountable owners, clear handoffs and evidence-based acceptance. Select controls proportionate to the project's risks and dependencies.

Purpose

This guide explains how to plan, create, revise, verify and release work across disciplines. It covers documents, datasets, code, models, media, research outputs, reusable templates, controlled derivatives, execution records and delivery packages.

The project owner determines the objective, approves the working boundary, assigns decision rights, and authorizes release. A human operator coordinates execution. Generative AI may assist with analysis, research, drafting, comparison, transformation, and testing; it does not supply missing authority, evidence, independent review, or signoff.

How to use the guide

When adopted for a project, use this guide as a governing input to the human–AI workflow. AI may help explain its terms, identify applicable controls and adapt their application to the current work, within the decision rights in Section 1.2. Apply the input and context controls in Sections 4.2 and 4.7.

Begin with Start Here and the Operator Control Sheet. Section 1 states the governing rules. Sections 2–4 cover project structure, change, execution and evidence. Activate the relevant Section 5 modules; use M6–M6.1 during content planning and drafting. Section 6 governs verification and release. The appendices provide working records, task prompts, definitions, conditional overlays and source notes.

Use existing project records wherever they can carry the required information. A control function does not require a separate document. A blueprint, artifact register, change matrix, and run record may be views or sections of one controlled system; they must not become competing sources of the same fact.

Boundary

This guide does not replace discipline-specific methods, applicable law, ethics and safety controls, information-security requirements, or required expert review. Define project-specific rights, approvals, numerical limits and contractual terms in an approved project overlay: the additional rules and choices for that project. Record implementation claims only when supported by evidence.

The workflow must remain understandable and executable by competent humans without AI. Approved artifacts and recorded decisions—not chat history, model memory, or unapproved summaries—control downstream work.

Contents and Task Routes

Use the linked entries or Word's Navigation Pane. Each rule has one authoritative location, called its canonical home. Records, prompts and overlays apply or point to that rule; they do not create additional authority.

Section	Use	Page
Start Here	Select controls for the task.	3
Operator Control Sheet	Record project decisions once.	4
1. Canonical Rules and Authority	Apply governing rules, decision rights and conflict handling.	5
2. Project and Artifact Architecture	Define identities, authority, status, and package activation.	8
3. Plan, Audit, and Control Change	Qualify baselines, bound scope, and control change.	12
4. Human-AI Execution	Run work units, research, claims, and handoffs.	19
5. Conditional Control Modules	Select only relevant preservation, evidence, action, and outcome controls.	24
6. Verification, Release, and Maintenance	Verify, authorize, deliver and maintain accepted work.	39
Appendix A. Working Record Templates	Create artifact, change, run and acceptance records.	43
Appendix B. Task Prompt Library	Launch scoped work under approved controls.	45
Appendix C. Glossary	Resolve controlled terminology.	47
Appendix D. Conditional Blueprint Overlays	Configure knowledge assets or controlled transactions.	49
Appendix E. Research Sources and Use Limits	Locate evidence, source versions and limits for cited controls.	51

Quick task routes

Need	Go to
Begin purposeful work; choose scope and commitment	Sections 3.1–3.1.1; M10 only when active.
Resolve authority, custody, or activation	Sections 1.2, 2.5, and 4.2.
Select baselines or assess dependency change	Sections 3.2–3.3.
Apply exact correction, migration, or protected-copy work	Sections 3.4–3.6; M1–M2; M8.
Manage context, parallel work and accepted handoffs	Sections 4.1–4.4; 4.7.
Test sources, derived claims, or assumptions	Sections 4.5–4.6; M7; M9.
Prepare restricted review or a live external action	M3–M4; Appendix D2 where applicable.
Verify and deliver the exact authorized output	M8; Sections 6.1–6.5.
Plan, draft and review substantive content	Sections 3.1 and 4.1; M6–M6.1; Appendix D1 when relevant; Section 6.1.

Start Here

Choose the mode, control level and intended use for the next activity, while retaining the project requirements that apply to it. Purposeful creation needs no universal novelty or demand gate; use Section 3.1.1. Keep permissions and outputs distinguishable where modes overlap.

Level	Use when	Minimum control
0 — Direct	A single low-risk task has no sensitive material, material evidence claim, or consequential dependency.	Objective and source boundary; direct work; local check.
1 — Structured	A small package or routine external output needs clear scope and coordination.	Compact control record; artifact roles; authorized outputs; relevant checks.
2 — Governed	Work is multidisciplinary, dependency-bearing, research-intensive, technically complex, or spans internal and external layers.	Blueprint or requirements; artifact register; scoped work units; evidence and integration checks.
3 — High-risk	Legal, safety, regulated, sensitive-data, rights-transfer, protected-copy, high-impact public, deployment, or irreversible-action risk is present.	Level 2 plus the relevant preservation, security, expert-review, action, and release gates.

Large multidisciplinary projects normally begin at Level 2. Escalate for the identified risk, not for document length alone. A low-risk public summary does not require the same control package as a safety-critical release. Conversely, a one-sentence high-impact claim can require Level 3.

Select the work path

Mode	Operating path
New build	Set purpose, resource boundary and sufficient provisional architecture. Use M6–M6.1 to construct content. Reconcile dependencies before broad reliance; verify the resulting content or behavior.
Research or analysis	Frame the question or exploratory purpose and safeguards. Refine through inquiry; declare method, evidence standard and validation route before qualifying conclusions.
Revision or report execution	Qualify the baseline; distinguish defects from suggestions; authorize changes and unchanged boundaries. Use M6–M6.1 for substantive revision within that scope.
Pairwise or reusable-master work	Protect history; compare by role and authority; import only approved, portable improvements.
Derivative, pilot, or live action	Define recipient, parent authority, allowed transformation, exposure, action scope, and required gates.
Hygiene or final release	Use an explicit change envelope; verify technical behavior, package identity, intended use, and acceptance.

Operating sequence

Purpose and boundary → Bounded work ↔ Architecture and qualified inputs → Verification → Authorized use → Stabilization. Exploration, reading, prototyping and drafting may develop together; resolve prerequisites before reliance. Revisions require diagnosis and authorized changes. Live actions retain applicable route, pilot and scale gates. Use M10 only when uptake or post-delivery benefits are necessary to the objective or commitment.

Stop when a required authority, input, permission, or acceptance condition is missing. A smaller authorized task may proceed only when it is independent of the blocker. Record the limitation rather than concealing it with a broad completion label.

Operator Control Sheet

At Level 2 or 3, record or reference the decisions needed for the next authorized scope; identify provisional choices and when they become due. Use a compact subset for Level 1. Artifact and work-unit details stay in their own records. Section 3.1.1 governs the basis to begin.

Control	Decision or authoritative record reference
Identity and objective	[Project/workstream; controlled identity; purpose; completion criteria suited to that purpose; outcome evidence only where applicable.]
Accountability	[Owner; operator; artifact owners; required reviewers; release authority.]
Mode and risk	[Active modes; control level; escalation reason; relevant Section 5 modules.]
Scope and non-scope	[Included work; forbidden outputs/actions; explicit exclusions; change authority.]
Active artifact universe	[Register location; current identities; missing, historical, protected, and excluded items.]
Governing controls	[Approved requirements/blueprint, policies, work orders, source domains, and package authority view where needed; compatibility decision.]
Baseline and dependencies	[Eligible baseline records; lineage limitations; required upstream acceptance; dependent work.]
Evidence and currentness	[Source/method plan; claim and assumption records; dates and review triggers; unresolved evidence or validation limits.]
Tools, data, and rights	[Approved environments and repositories; data classes; permissions; prohibited transfers/actions.]
Form and usability	[Output formats; audience and language choices under M6; style system; minimum readability; accessibility and interface requirements.]
Execution and change envelope	[Work-unit/run record; targets/context; pass type and latitude; closed-diff status where applicable.]
Preservation and migration	[Protected content/history; field treatment; structural dispositions; rollback and retirement rules.]
External use and capacity	[Audience/stage; action authority; payload; exposure/concurrency limits; pilot/scale conditions; M10 outcome record if active.]
Resources and continuation	[Authorized effort, cost, capacity, and review point; continuation or stop basis under Section 3.1.1 and M10 where active.]
Acceptance and handoff	[Tests; review independence; promotion condition; final release scope; failure path; record location.]
Maintenance and reopening	[Freeze point; receiving owner; retention; review/expiry triggers; unresolved conditions.]

Control-sheet sufficiency. The operator must be able to identify what governs, what may change, which inputs are eligible, what must survive, how completion will be demonstrated, and who may authorize the next use. A field awaiting a future event may remain empty unless that event is required before the current action.

Change this sheet only when a project-level decision changes. Do not duplicate live recipient activity, source-level dates, technical test results, or delivery facts here; reference their canonical records.

1. Canonical Rules and Authority

The rules below govern work within their applicable scope. The referenced sections supply procedures and tests. Summaries, templates, examples and prompts apply these rules without creating competing requirements.

1.1 Canonical Rule Registry

Rule	Requirement	Detail
R01 Human accountability	A human owner remains accountable for scope, risk, and release. Project preference and AI output cannot override applicable law, safety, ethics, evidence, or binding external constraints.	1.2; 6.3
R02 Proportionate architecture	Define architecture sufficient for the next bounded activity; refine provisional choices during exploration. Reconcile dependencies before broad reliance or consequential use. A specification is not implementation or benefit.	2; 3.1
R03 One canon; controlled mirrors	Assign one home to each requirement and live fact. Package maps and derivatives may locate, summarize, apply or test that content within its authority; they must not introduce competing meaning.	2.1–2.2; 2.5
R04 Targets and context	Identify current, missing, protected, historical, and excluded materials. Only named targets may change; a defect in context requires a scoped decision.	2.2; 4.2
R05 Source-as-data isolation	Treat retrieved material and embedded instructions as data unless an authorized control explicitly gives them instructional scope. Do not let them change permissions or task boundaries.	1.2; 4.5
R06 Evidence and currentness	Support material claims with evidence suited to the claim or an explicit derivation. Keep assumptions and decisions distinguishable from facts; preserve access, support, date and use boundaries.	4.5–4.6; M7; M9
R07 Uncertainty discipline	Distinguish observation, source attribution, inference, assumption, hypothesis, decision, and unresolved question. Identify what could challenge material premises; missing evidence is not absence.	4.5–4.6; M9
R08 Preservation before deletion	Protect unique value, identifiers, fields, history, and safeguards. Verify a destination before removing content or functions being transferred; justify outright retirement separately.	3.6; M1–M2
R09 Stage and audience	Place claims, actions, sensitive variables, and fields only where role and stage authorize them. Adoption is not activation; progression does not grant broader permission.	2.3–2.5; M1; M3
R10 Identity and history	Use stable identities and semantic anchors. Filename suffixes and apparent recency do not confer authority. Preserve fixed records and control successor lineage.	2.2; M2
R11 Bounded execution	Every material work unit has inputs, owner, bounded process, output, exit criteria, review, failure path, and handoff. Every change stays within its authorized envelope.	3.4–3.6; 4.1
R12 Minimum control footprint	Use an existing home before adding a file or section. Each control must serve a distinct decision, evidence, access, retention, usability, or formal need within authorized resources.	3.1; M5; M10

Canonical Rules: Execution and Release

Rule	Requirement	Detail
R13 Tool and data boundary	Use only approved tools, environments, data, and actions. Capability or access alone does not create permission to disclose, process, transmit, or deploy.	4.2; M4
R14 Observable work	Claim only work, interaction, and review actually evidenced. Carry approved state through records, not memory; keep the workflow human-operable. Activity is not an outcome.	4.4–4.5; M4.1
R15 Technical and semantic QA	Test both meaning and technical behavior. Exact-payload checks, rendering, package integrity, and reproducibility cannot substitute for evidence or authority.	M8; 6.1
R16 Gates and review	Accept work through the applicable verification regime. Distinguish a fresh self-reconciliation from genuinely independent review and from specialist validation.	6.1–6.3
R17 Reusable-master genericity	Admit a lesson only after abstraction, coverage and contradiction checks, canonical placement, and testing across relevant project types. Local defaults do not become universal rules.	M5
R18 Clean artifact boundary	Keep work orders, prompts, analysis, and QA records outside final outputs unless their role requires them. Preserve required source attribution and substantive limitations.	2.1; M8
R19 Compatible input and run set	Confirm governing controls, active inputs, exact targets, run order, outputs, and ownership are compatible. Lock one current instance per identity; competing copies block promotion.	4.1–4.3
R20 Layer and narrator	Use the voice and authority of the artifact's layer. A wrapper, internal note, certificate, or model response must not silently create substantive or institutional authority.	2.1; M6
R21 Recipient actions and field completion	Request only actions within the recipient's remit. Populate a field only from the authorized source or event that establishes its value.	2.4; M1; M3
R22 Qualified promotion	Only accepted, non-superseded outputs control downstream work. Acceptance names its purpose. Assess changed dependencies and affected verification before any reuse.	3.2–3.3; 4.4
R23 Deliverable integrity	Return the authorized set in the specified form, with accurate per-item disposition. Preserve unchanged dependencies and distinguish support records from deliverables.	4.4; M8; 6.4
R24 Guardrail non-regression	Compression, cleanup, or simplification must not weaken safety, evidence, rights, access, privacy, AI-use, or claim limits. Preserve or escalate uncertain cases.	3.6; M6

1.2 Decision Rights and Authority Domains

Role	Accountability and limit
Owner / sponsor	Sets objectives, resources, authorized risk, and final release scope within binding constraints. Does not fabricate facts or waive mandatory safeguards.
Operator / controller	Maintains the working boundary, selects eligible inputs, directs work, resolves routing, and assembles acceptance evidence. Does not silently broaden scope.
Artifact or workstream owner	Maintains the assigned canonical content, approves local changes within authority, and accepts or returns handoffs.
Reviewer / specialist	Makes the named technical, scientific, legal, safety, security, or accessibility judgment. Certifies only the scope actually reviewed.
AI system / execution tool	Performs authorized assistance or observable operations. Neither output polish nor a successful call establishes factual validity, approval, or review independence beyond the recorded check.

Resolve by domain, then by authorized precedence

Question	Controlling source and boundary
What is legally or otherwise mandatory?	Applicable law, safety/ethics requirements, binding agreements, and competent authority within their domains.
What does the evidence support?	Qualified observations and permitted sources. An owner may choose policy but may not override contrary facts.
What is the project intended to do?	Owner-approved objectives, blueprint, requirements, and decisions; none silently changes external obligations.
What is the active set and permissible change?	Artifact register, approved findings, and change records, within their assigned scope.
How is this run executed?	The compatible work-unit contract and launcher; procedure may select and sequence authorized work, not invent new work.
What does this record prove?	The accepted source, signed instrument, or observed event specified by the record's role—not its title, age, or formatting.

Where an artifact combines external obligations with project-adopted requirements, identify the source and scope of each. A project may adopt a stricter internal requirement where permitted, without presenting it as an external mandate or overriding binding constraints. Clarifying provenance must preserve the force of requirements, recommendations, permissions and prohibitions.

An apparent conflict may be a difference in domain, stage, representation, or recorded date. Check those distinctions first. If the same material question still has incompatible controlling answers, stop the affected work and obtain the appropriate decision. Do not resolve it by recency, prose volume, model confidence, or a more detailed launcher.

1.3 Stop and contain

Stop for missing or conflicting authority; ambiguous targets or baselines; unsafe source/currentness gaps; threatened historical or protected content; unknown final facts; an unauthorized tool/action; a failed technical or release condition; or an unapproved change in scope. Record the affected item, unmet condition, dependency impact, responsible resolver, and permitted next step. Independent work may continue only within an explicit unaffected boundary.

2. Project and Artifact Architecture

2.1 Artifact universe and canonical homes

Inventory every controlled work product that matters to the objective, including missing and excluded items. Assign each artifact one primary job, owner, audience, representation, and intended action. A single physical file may contain multiple clearly separated functions; one function may have controlled derivatives, but only one authoritative home.

Layer / home	Owns	Must not imply
Core / canonical	Requirements, doctrine, methods, schemas, terminology, source rules, or technical behavior.	That a stated requirement is implemented or validated.
Evidence / data	Source identities, observations, measurements, provenance, claim support, and dated source states.	That a citation, accessible link, or completed source record proves a stronger claim.
Code / models / computation	Implementations, configurations, environments, tests, scripts, notebooks, and reproducibility records.	That an executable build is safe, authorized, or ready for every intended use.
Media / design	Source assets, compositions, captions, accessibility variants, and format-specific design decisions.	That inclusion establishes rights, attribution, or permission for reuse.
Wrapper / derivative	Identity, classification, local use boundary, summaries, extracts, exports, or adapted representations.	That a label or conversion changes the parent's evidence, rights, or release status.
Support / formal / access	Party-defined terms, use restrictions, permissions, scope, reviewers, remedies, and formal actions.	That informal explanatory copy amends binding terms or opens access.
Internal operations	Plans, routes, exact-copy masters, work orders, current activity, issues, and approval references.	That an internal strategy, example, or draft is an observed event or external commitment.
Delivery / acceptance	Frozen perimeter, integrity values, receipt, verification, exceptions, and acceptance facts.	That a manifest, hash, or certificate adds rights, source validity, or product clearance.
History / exclusions	Issued or signed records, superseded identities, retention evidence, and unavailable or excluded material.	That historical material is the current baseline or that absence from the working set proves nonexistence.

Derived views identify their source and version: requirement summaries derive from governing provisions; factual views derive from identified evidence, observations or builds. A requirements crosswalk does not prove implementation. A summary of obligations either includes the full applicable set or is explicitly abbreviated with a direct route to omitted requirements.

Use descriptive, role-appropriate voice. An artifact may define, summarize, specify, or record; it cannot itself perform a human judgment, grant authority it does not hold, or establish a real-world action that has not occurred. Internal instructions belong in internal controls, not in the recipient-facing body.

2.2 Artifact Identity and Register

Maintain one artifact record per controlled identity. The record doubles as the role card; do not create a separate role map containing the same live values. Template A1 in Appendix A supplies a compact entry form.

Record field	Required meaning
Identity and representation	Stable title/ID, version or edition where used, physical location, format, and whether governing, derived, template, example, historical snapshot, or event record.
Role, layer, and audience	One-sentence job, owner, intended recipient/action, stage, distribution class, and content that belongs elsewhere.
Working-set state	Target, read-only context, control input, evidence source, candidate, accepted, superseded/rejected, historical, absent, or excluded.
Baseline and lineage	Source identity and fixed version binding where required; comparison basis, acceptance or qualification reference, known lineage limits, represented units, and permitted downstream use.
Authority and dependencies	Governing controls; canonical content, term and state homes; upstream imports; dependent outputs; derivative transformation/update rules.
Status and field treatment	Relevant status dimensions from Section 2.3; how each field must be treated, including whether it is descriptive, fillable, created by an event, protected, or generated.
Preservation and environment	Exact-copy or history obligations; code/data behavior; visual and accessibility requirements; approved tools, environments, and formats.
Evidence and currentness	Required sources, claim ceiling, verification scope, access limitations, date semantics, and refresh triggers.
Output and acceptance	Authorized delivery form, QA/approver references, release scope, failure route, retention, and reopening condition.

Identity is not the filename alone

Distinguish the stable logical identity, the actual accessible path or repository reference, and the required output name. Normalize interface-added suffixes only when the project's naming rule authorizes it. Never invent a path from a display title or assume a file exists because another artifact names it.

Bibliographic edition, working revision, release instance, and temporary review label identify different things. Coordinate titles, headers, properties, manifests, and links according to their roles; they need not repeat an identical full string everywhere. A technical save or upload suffix does not establish a new edition or approved version.

When a revision controls review, reuse or delivery, link its identifier to an exact source state recoverable from existing version history, retained snapshots or equivalent records. A mutable location or date alone is insufficient. Distinguish the source revision, the identity of a derivative or copy set, and the issue and check dates. If an output combines several source revisions, identify each one.

Representation and version control

A governing item states requirements. A derivative traces to a parent and has a defined transformation/update rule. A template contains fields for later completion. An example illustrates without setting defaults. A snapshot records an as-of view. An event record states only what the authorized event established. Label these distinctions where misunderstanding is plausible.

When accepting an object for downstream use or issuing it, version the object that changes. A change to governed content, behavior or structure used by references requires a successor version of that object. A change only to format, watermark, packaging or access may instead update the derivative or release/access record while retaining the unchanged source revision. Check that meaning, locators and function remain intact. Do not silently overwrite an issued representation. Changed permissions still require the applicable authority. Use M8.1 for affected checks and Section 6.4 to reconcile the reviewed and delivered outputs.

2.3 Status Dimensions and Stage Meaning

Record separate kinds of state only where material. A project may have an adopted architecture, incomplete implementation, limited validation, restricted release, and high planning confidence at the same time. Do not collapse these into “ready” or “approved.”

Dimension	Question and evidence
Governance / adoption	What rule, requirement, artifact, or baseline controls? Record proposal, adoption, freeze, supersession, and the approving authority.
Implementation / execution	What has actually been built, populated, configured, run, or completed? Require an observable result or event record.
Evidence / validation	What was tested, extracted, reviewed, or independently confirmed, by whom, and to what standard? Preserve scope and unresolved limits.
Release / use	Who may access, publish, rely on, transfer, deploy, or reuse the item, in which environment and for which purpose?
Confidence / dependency	How uncertain is the assessment and which conditions could change or block it? Confidence is not a passed gate.

Each material transition identifies the object, actor, prerequisite, evidence, date, and resulting scope. Scheduled events, signatures, adoption, effectiveness, activation, and verified operation may be distinct. Apply Section 2.5 to custody/activation and M7 to dates; do not invent a linear sequence when activities can legitimately overlap.

2.4 Audience and coordinated components

Recipient situation	Appropriate function and boundary
Public or first qualified read	Define the object, purpose, audience, present value, concise limits, and next decision. Disclose only what the approved scope and evidence support.
Sample or restricted evaluation	Provide bounded proof and provenance, explain omissions, and state allowed actions and the next gate. Apply M3.
Specialist or technical review	Expose methods, assumptions, sources, tests, defects, and uncertainty needed for the specified judgment.
Legal or formal action	Present operative terms and completion fields for authorized parties. Keep internal strategy and unsupported event facts outside.
Internal operation	Maintain policy, exact-copy blocks, activity, issues, and evidence pointers in their assigned homes; distinguish copy to send from notes.
Delivery or acceptance	Provide the identifiers and evidence needed to verify the package, then record the actual result under the governing acceptance procedure.

Each component has one principal job and works both alone and in the intended sequence. Repeat essential identity, boundary, and handoff anchors where they protect an isolated or forwarded component. Do not repeat the entire argument, evidence set, or operating procedure. Test expected and out-of-order reading; label optional branches as optional, and add a gate only for a real prerequisite.

2.5 Package Governance and Activation

When a multi-artifact package divides substantive, evidence, rights, execution, and release authority, maintain a package-level view of the artifact records in Section 2.2. Use the existing blueprint or control system. A separate constitution is justified only by a distinct governance or audience need; it is not a second set of artifact records.

Package decision	Required treatment
Perimeter and authority	Identify the governed package, supporting layers, exclusions, and the domain each authority controls. Complementary editorial, technical, research, and formal authorities need explicit interfaces, not an invented universal hierarchy.
Canonical ownership	Expose each artifact's primary job, canonical home, imports, exports, prohibited uses, and update rule from its existing record. Keep a single writer for each active identity.
Package-wide rules	Record only the few project-specific rules that all derivatives must inherit, with their authoritative locations. A local summary may not weaken them or silently turn a proposal into an adopted requirement.
Current-state containment	Keep volatile evidence, assumptions, relationship facts, and execution events in their designated live records. The architecture points to these records without certifying or duplicating their current values.
Change propagation	When authority, an interface, a package-wide rule or a permitted use changes, review the artifacts and systems that rely on it under Section 3.3. Release linked changes together where required. Leave unaffected artifacts unchanged.

Custody, activation, and project work

Distinguish the authority to hold, version, govern, or evaluate a package from permission to activate a specific use. The same person may hold several roles, but custody, project sponsorship, operation, assurance, and release decisions retain their separate scope.

Record the intended use: hold and govern; bounded development or evaluation; an authorized operational use; or reuse within a demonstrated scope. These are project choices, not mandatory stages. Activation becomes project work when a named environment, counterparty, resource commitment or live action must satisfy specific permissions, contracts, evidence or operating conditions. Use Section 4.1 and M4 for that work.

Package control supplies neither operating proof nor implied support services. A reuse or replication case must identify what remains equivalent and what needs fresh proof; architectural similarity alone is insufficient. Required obligations and independent evidence retain their own authority under Section 1.2.

3. Plan, Audit, and Control Change

3.1 Blueprint and audit decisions

For a new build, establish a proportionate blueprint under Section 3.1.1; provisional architecture may develop with exploration. For existing work, distinguish audit, bounded correction, authorized redesign and release check. An observation or attractive alternative is not automatically a defect.

Blueprint component	Decision to establish
Objective and perimeter	Purpose, reader or user need, intended contribution, success test, included workstreams, non-scope, forbidden outcomes, and change authority.
Architecture and ownership	Artifact roles, canonical homes, recipient stages, decision rights, dependencies, and milestone handoffs.
Evidence and methods	Questions, hypotheses, available substance and sources, material gaps, research/computation method, evidence standard, validation, limitations, and currentness rules.
Tools and format	Approved environments, data/rights/security boundary, output formats, terminology and language choices under M6, visual identity, and accessibility.
Execution and release	Work-unit completion criteria, ordering, review independence, output rules, package scope, retention, and reopening triggers.
Assumptions and intended value	Reference M9 for material planning premises and M10 where outcomes depend on uptake. Define resource limits, learning/continuation decisions, and separately due evidence.

Choose whole-project, artifact-by-artifact or workstream-batch production according to dependencies, complexity, context capacity and verification load. Use the common work-unit contract in Section 4.1. For content work, connect the reader's need, intended contribution, available substance and unresolved gaps in the existing blueprint or task record; M6.1 governs construction. A small task may need only a few sentences. Provisional choices may develop through inquiry under Section 3.1.1. Appendix D supplies conditional overlays for knowledge/publication assets and controlled evaluation or transactions.

Audit coverage

Examine the applicable domains: project structure and authority; identity and history; artifact role, audience and permitted claims; evidence and currentness; legal, safety, privacy and rights requirements; redundancy and usability; technical behavior and reproducibility; capacity for external actions; and release and delivery. State which materials, pages, records, technical components, tests and external facts were actually examined.

Distinguish an internally consistent source set from externally verified facts, and both from independent specialist validation. A search hit or automated warning is a candidate finding; inspect context before classifying it as a defect. A no-change outcome is valid when the artifact already fulfills its role.

Convert only accepted findings

For each candidate, assign severity, evidence, affected and unaffected scope, value of change, and action posture. Then authorize its target, instruction, latitude, preservation, dependencies, and acceptance test. Template A2 in Appendix A supplies the change record. A report and its launcher may share a record, but their substantive and procedural roles remain distinct.

Value-of-change test. A change must improve a decision, evidence support, accuracy, usability, technical behavior, preservation, handoff, or release control. Additional wording, a new table, or a different design is not enough. Preserve accurate, role-fit material and consolidate before expanding.

3.1.1 Purpose, Scope and Development Commitment

A coherent purpose and an authorized, proportionate commitment can justify exploration or creation. Learning, artistic expression, research, education, preservation, public benefit and internal use are legitimate purposes; they need not be converted into revenue or externally measurable utility. Do not require exhaustive landscape research, an unoccupied niche, unprecedented novelty, a buyer or prior demand before work may begin.

Architecture and research may be provisional and sufficient for the next bounded activity. Reading, prototyping, exploration and drafting may develop together. Reconcile dependencies before broad reliance or consequential action, not before every act of discovery. An explicit purpose and tolerable resource basis may authorize substantial noncommercial production; do not confine it to perpetual pilots.

Contribution and commercial viability

Keep contribution or originality, relevance to a purpose or audience, anticipated or demonstrated usefulness, and commercial viability distinct. A contribution may be synthesis, replication, interpretation, accessibility, pedagogy, localization, integration, timeliness, execution quality or a different perspective. Existing alternatives do not by themselves invalidate the work.

A proportionate landscape check can improve accuracy and positioning. Investigate claims such as “first,” “unique,” “no equivalent” or comparative superiority at their actual scope under Sections 4.5–4.6 and M7. Narrow or remove an unsupported claim; hold only work that materially depends on it. Preserve attribution and applicable disciplinary or formal originality requirements. This policy supplies no ownership, rights clearance or protectability.

Match evidence to the decision

Decision	Requirement for that decision
Begin or explore	A coherent purpose, tolerable authorized commitment and safeguards due for the activity; no universal market test.
Commit more resources	Reassess material assumptions, alternatives, dependencies and review capacity. The rationale may be noncommercial.
Make consequential claims or release	Meet claim-specific evidence, permissions and quality requirements, even for a cheap or small output.
Pursue uptake or commercial return	Test the uptake or economic assumptions on which the declared objective or commitment actually depends; activate M10.
Expand exposure or scale	Obtain evidence and authority proportionate to added consequences. Require a pilot only where justified.

These are decision triggers, not a compulsory linear funnel. Research effort follows risk, uncertainty, reversibility and intended reliance—not cost alone. Explicit uncertainty is permissible where the activity can safely tolerate it; mandatory prerequisites remain due.

Scope Sufficiency and Continuation

Define what “enough” means before breadth, polish, or control overhead becomes the default objective. Each added unit should serve a required function, advance the chosen creative or learning purpose, materially improve a decision, supply relevant counterevidence, resolve a consequential gap, or meet a named reader or user need.

Scope question	Decision rule
Required coverage	Map required topics, functions, interfaces, and acceptance cases to their homes. Track material omissions in the existing gap record with their consequence, owner, due stage, and release effect.
Saturation and stopping	For open-ended research or content, define a scope-appropriate saturation test. Stop discretionary expansion when credible additions mainly repeat covered functions without a stated creative, learning or coverage purpose. Record the search or comparison basis; file count or an elegant outline does not prove completeness.
Coverage debt	Keep important but unresolved coverage visible. An acknowledged omission is not fulfillment of a mandatory requirement. Narrow the consuming scope only under Section 3.3, or retain the applicable hold.
Core and conditional scope	Separate the current deliverable from conditional extensions and exploratory possibilities when their evidence or maturity differs. Define what would admit a conditional item. Core membership itself does not establish deployment readiness.
Whole-system capacity	Include research, coordination, specialist review, correction, documentation, delivery, and any agreed maintenance—not production alone. Reserve capacity for failures and rework rather than treating full nominal utilization as resilience.

Bound the next commitment

Record the next work scope, expected information or outcome, authorized time/cost/effort, required review capacity, and decision point in the existing control record. Separate substantive creation from support/collateral work so either can be challenged when it no longer advances the objective. A control or improvement should justify its recurring cost as well as its initial usefulness.

For an outcome-dependent investment, M10 governs the value and continuation case. Exploratory work may proceed under an explicitly bounded uncertainty budget; prior revenue or external commitment is not a universal prerequisite. Mandatory compliance, safety, and preservation work is governed by its applicable obligations, not by a commercial return test.

Reassess without weakening the standard

When effort, dependencies, or benefits depart materially from the approved basis, reassess the remaining work and alternatives. Reduce discretionary scope, slow, hold, redirect, or stop through the proper authority. Sunk effort and output polish are not evidence of future value. A resource constraint never authorizes fabricating evidence, omitting a mandatory safeguard, or changing pass criteria after the result.

3.2 Baseline Qualification

Baseline / state	Permitted use	Qualification limit
Original source	Compare the exact eligible pre-change or authoritative current artifact with the output against the approved change set.	May establish source-preserving qualification for the surfaces actually checked.
Authorized corrective baseline	Use an already-revised artifact when the original is unavailable or not selected and an accountable authority approves bounded correction.	Qualifies the authorized delta and downstream use; does not certify unavailable source history.
Accepted carry-forward	Use an unchanged accepted output read-only for its named dependency purpose.	No implied fresh review, verification, resave, or redelivery.
Candidate output	Inspect, test, and correct during the active work unit against the eligible baseline and approved controls.	Not canonical, released, or eligible to control downstream work until accepted.
Rejected / superseded	Retain for history or diagnosis where required; restart authoritative comparison from an eligible baseline.	Not a retry baseline or current authority unless explicitly requalified.
Historical reference	Use for comparison, provenance, or a bounded historical claim.	Does not become the current operating baseline merely by being accessible.

Acceptance records the exact identity, accepting authority, date, purpose, evidence, and provenance limits. Acceptance for the next work unit is not publication approval, legal clearance, execution, deployment, or permission for unrestricted use.

3.3 Dependency-impact decision

Identify what each dependent artifact or system uses from the changed input: meaning, data, status, identifiers, interfaces, references, behavior, or exact bytes. The affected-dependency list defines what to review; it does not instruct the executor to rewrite every listed item.

Observed impact	Disposition
No relevant dependency affected	No downstream action beyond recording the determination. Do not reopen work for convenience.
Possible impact; output remains adequate	Revalidate the relevant section, reference, behavior, or check. Close as “Reviewed—no change” with basis, reviewer, and scope.
Consumed meaning or behavior must change	Authorize and rerun only the affected derivation, calculation, passage, test, or work unit; preserve unaffected work.
Baseline or impact is unresolved	Block dependent promotion/use until the eligible baseline and required correction or check are established.

A completed dependency review records the item examined, result, responsible owner, and any separate adaptation required. It does not refresh unrelated source dates, certify the entire parent, or create a new release permission. Apply M8 to checks or approvals tied to a changed representation.

A dependency is satisfied only when its required evidence, action and review are complete. Where governing requirements permit, the authorized decision maker may narrow, defer or exclude the work that relies on it. Record the removed scope. Confirm that the retained output no longer relies on the unmet dependency and still meets its own requirements. The underlying requirement remains unsatisfied; reducing scope neither completes the evidence nor waives mandatory conditions.

3.4 Change Posture, Pass Type, and Latitude

Record three distinct decisions in one change record: whether to act now (action posture), what kind of work is authorized (pass type), and how much discretion the executor has (latitude). These decisions describe the same work; they do not require separate workflows.

Action posture	Required handling
Apply	Perform the authorized change now and test the stated postcondition.
Conditional control	Record the trigger and restriction now. Activate the control only when evidence establishes that its condition has been met.
Watch	Record the signal and refresh trigger without treating it as a current fact, permission, or mandatory edit.
No action	Record why no target change is authorized; preserve the unchanged boundary.
Block	Hold affected work for the named missing authority, evidence, input, or safety condition.
Completed interlock	Use the accepted results of a prior bounded verification/correction; do not reapply superseded inputs.

Pass type	Scope of work
Exact correction	A specified text, field, node, code, or format change. Use the matching contract in Section 3.5; incidental cleanup is not implied.
Targeted revision / synchronization	Repair a named defect or align a confirmed term, reference, identity, or state in affected locations only.
Harmonization / structural repair	Reassign, combine, or coordinate content/functions because the approved defect is structural, not merely stylistic.
Compression	Remove or consolidate functionless repetition after preservation and guardrail classification.
Identity / metadata / technical pass	Change defined identity or technical surfaces while preserving governed meaning and accepted variation. Use an allowed-diff envelope.
Research / currentness refresh	Perform the declared evidence or status review; source maintenance does not itself authorize a broader substantive cascade.
Generative build / reusable abstraction	Create within an approved blueprint, or abstract qualified lessons into an existing master under M5.

Latitude labels: L0 inspect/classify/report only; L1 exact specified operations only; L2 minimal authorized local coherence; L3 bounded passage/table rewrite; L4 authorized structural change; L5 generative build; L6 advisory options without editing. These are permission labels, not an automatic escalation ladder. The narrower active instruction prevails.

Exact-edit boundary. An L1 instruction does not permit grammar, punctuation, whitespace, style or field refresh outside its enumerated operations. L2 must be expressly granted when consequential wording changes are needed. A closed change list further limits the change surfaces as set out in Section 3.5.

3.5 Exact Correction and Closed Change Lists

Use this contract when an edit must be reproducible and tightly bounded. The change record must establish the following before mutation.

Control	Required specification
Eligible input	Exact artifact identity, working state, baseline qualification, and actual accessible path/reference.
Stable target	Semantic heading, clause, row/key, bookmark, schema/code path, or package part. Page numbers and table indices are secondary locators.
Match precondition	Old-text/value anchor, expected match count, allowed locations, and any expressly permitted already-satisfied state.
Operation and payload	Exact replacement, insertion, deletion, row addition, field repair, or bounded operation; sequencing where operations interact.
Change envelope	Permitted semantic objects and technical surfaces; protected content, relationships, identifiers, styling, fields, and behavior.
Locator normalization	Only declared equivalences needed to locate split runs, soft-wrap artifacts, or character variants. Location normalization must not rewrite unrelated content.
Postcondition and failure	Expected final value, count, preserved references, residual tests, and the exact mismatch that requires a stop.

Where a work order mixes instructions and proposed wording, distinguish authorized payload, decision-dependent candidate text and non-transferable commentary. Identify each payload's destination and boundaries; copy neither its label nor surrounding instructions. Resolve and record required method, schema, rights or acceptance decisions with the competent authority before using dependent text. A polished paragraph is not an adopted decision or a complete field/test contract. Unaffected edits may proceed; unresolved candidate wording must not enter an authoritative issue.

Dry run, then apply

Before editing, locate and count the required matches. For interacting operations, resolve all targets and required templates against the eligible baseline before the first change. Map original objects to their working-copy counterparts so insertions, deletions and reflows do not change target identity. Check overlaps and operation order. Unexpected zero, multiple or materially different matches block the affected operation. Report the mismatch; do not improvise a substitute. Independent authorized actions may proceed only when they do not depend on it.

A missing old string is not proof that the correction was made. Verify postconditions at the mapped destination; an old string may legitimately remain elsewhere. Where the contract permits reapplication without another change, verify the complete postcondition, required location/count and absence of mixed states, then record "Already satisfied—no mutation." Otherwise obtain a disposition. Never restore obsolete text just to make an old instruction match.

Closed-list acceptance

A closed change list enumerates every permitted operation and surface. For an Office Open XML (OOXML) document this may identify parts and text nodes; for code it may identify files/functions; for data it may identify fields/records. Compare before and after at the declared level. Untargeted surfaces must remain equal at the required semantic or byte level.

Do not confuse archive-container bytes with content equality. Compression, serialization, or application bookkeeping can alter bytes without changing meaning; such changes are allowed only when the contract explicitly permits and tests them. Conversely, equal rendered pages do not prove that hidden data, relationships, or fields were preserved.

Bounded example. A footer-label correction can permit the named footer text nodes while preserving the document body, style definitions, page fields, and unrelated metadata. A tool that broadly round-trips the package is unsuitable when it cannot demonstrate that envelope. Escalate rather than silently relaxing the test.

3.6 Structural Change, Migration, and Retirement

Assign a primary disposition to each affected unit: KEEP, CLARIFY, EXPAND, MERGE, MOVE, COMPRESS, REWRITE, DELETE, AUTHORITATIVE ONLY, or NO CHANGE. State the resulting location and function, not just the heading that will change.

Situation	Required treatment
Move or merge	Create and verify the destination before removing the source occurrence. Preserve unique evidence, controls, references, fields, and downstream function.
Delete without replacement	Establish that the material is obsolete, unauthorized, or functionless, and that no required dependency or safeguard relies on it. Do not invent a replacement merely to fill space.
Meaning unchanged	Preserve stable identifiers and legitimate aliases; update physical locators only as necessary.
Meaning changed but mappable	Approve the mapping, migration rules, affected consumers, and validation before promoting the successor.
Meaning incompatible	Create a new canonical identity or explicit schema version; retain the former mapping as history where required. Do not silently reuse an ID for a different concept.
Empty or legacy field	Classify as unknown, not yet applicable, deprecated, or requiring review. Absence alone does not create permission, completion, or a new value.
Retired control or artifact	Remove from active use and prevent later sessions from recreating it under a new label. Reintroduction requires a new approved need and canonical placement.

Preserve the useful function, not every repetition

A shortened or merged output must still perform the assigned job. Renumbering, restyling, changing filenames, or nesting substantially unchanged material does not demonstrate a requested functional transformation. Acceptance must show the actual move, merge, narrowing, removal, or improved handoff.

Separate distinct domains even when they concern the same subject: policy specifies what must happen; a configuration states how this instance is set up; an event record states what happened; a summary points to the relevant record. Co-locate them when useful, but do not maintain their live values independently in several places.

Cross-artifact change sets

Where a correction changes linked source IDs, bibliography entries, claim wording, schema fields, or contractual references, identify the coupled set. The required items may be edited in sequence, but dependent use and release wait until the set is mutually consistent. Do not expose a partially synchronized state as accepted. Unrelated accepted items need not be rerun.

Keep migration and acceptance evidence in the existing control system. Reader-facing outputs contain durable current content, not the rationale, source-file comparison, work-order codes, or discussion that produced the change. Required historical records and substantive provenance remain where their roles require them.

4. Human-AI Execution

4.1 One contract for each material work unit

A phase, session, research task, batch, migration, or release operation is a bounded work unit. Use INPUT → OWNER → PROCESS → OUTPUT → EXIT CRITERION → REVIEW → HANDOFF, with a failure path. The same record can serve as the session contract and run-matrix row; do not maintain parallel copies.

Contract field	Required entry
Objective and controls	Decision or outcome; reader or user purpose where relevant; governing control identity; approved creation scope or action rows; explicit non-scope.
Inputs and baseline	Complete required target/context set; accessible locations; qualified baseline and upstream acceptance where applicable; known lineage limits. For content work, reference the substance and gaps identified under Section 3.1.
Ownership and access	Executor, reviewer, accepting authority, integration owner, and canonical write owner; data/tool permissions.
Process and latitude	Method, pass type, action posture, permitted transformations, mandatory steps, read order, and prohibited actions. Apply M6–M6.1 during content construction and substantive revision.
Dependencies and sequence	Entry conditions, run mode, same-session promotion if authorized, coupled change set, and downstream consumers.
Preservation and outputs	Protected surfaces and history; exact output identities/formats; support-record treatment; no-change and blocked-item rules.
Exit and review	Definition of done; observable tests; required specialist or independent review; applicable Section 6 acceptance scope.
Failure and handoff	Hold, return, adjudication, rollback, or retry route; promoted state; receiving owner; dependency-impact and reopening conditions.

4.2 Input and control compatibility

Before changes begin, the operator selects and approves the working set. The executor confirms that the required inputs are accessible, readable, correctly identified, and used in their assigned roles. These responsibilities are distinct. Do not repeat completed manual preparation without a reason, and do not ignore an actual access failure because preparation was reported complete.

Compare the blueprint, work order, action record, run plan and launcher wherever they govern the same work: target scope, active versions, dependencies, preservation, output set, and acceptance. A launcher cannot resolve a conflicting governing instruction by implication. If a material mismatch remains, obtain a decision from the responsible authority before running the affected work.

Load the context needed for the authorized purpose; use Section 4.7 when context is compressed, split or transferred. Repeated references reuse the same controlled instance. Historical comparators may coexist with the current target only when explicitly labelled for comparison. Do not infer missing content, final file paths, source records, or private facts from filenames, summaries, or other sessions.

Permission to begin. Uploads, a complete template, or a technically available tool do not themselves authorize editing or external action. Use the approved work-unit instruction. Where a task begins in advisory mode, do not silently convert recommendations into file changes.

4.3 Run Order, Concurrency, and Recovery

Plan work in the order required by its actual dependencies. Do not impose a fixed number of conversations or a universal “foundation first” sequence. A foundational input must be accepted before dependent work relies on its meaning. Independent work may share the same accepted, frozen context.

Run control	Operating rule
Sequential dependency	Run the prerequisite first. Define the exact output and acceptance state required by the next unit.
Parallel independent work	Permit independent target sets that share qualified read-only inputs and do not require one another's unaccepted results.
Single canonical writer	Assign one writer to each active artifact identity at a time. Separate code or document branches are candidates until a named integration owner reconciles and accepts them.
Session sizing	Bound work by target complexity, context volume, tool constraints, review load, and risk—not file count alone. Do not start a batch too large to verify.
Same-session promotion	An earlier target becomes context for a later target only when the contract specifies the order, acceptance criteria, and authorized promotion.
Interlock	Insert a bounded verification/correction when needed between dependent stages. Consolidate the accepted result into the baseline and retire superseded inputs.
Retry and recovery	Correct an active candidate under the same contract, or restart from the last eligible baseline after rejection. Preserve the authoritative comparison basis and rollback evidence.
Coupled promotion	Accept linked changes together as a consistent set. A file that passes its own checks does not authorize use of an inconsistent dependency. A failed branch blocks work that depends on it; unrelated branches may proceed under their own acceptance conditions.

4.4 Completion, reporting, and downstream handoff

The producing owner is responsible for local content, preservation and technical QA in each work unit. Resolve avoidable omissions, unclear explanations and basic defects within the authorized scope before handoff. Route gaps requiring new evidence, specialist judgment or additional authority to the responsible owner. Additional independent or integration review may be required by risk; it supplements local acceptance and does not automatically authorize further editing.

Account for every named target as passed within scope, reviewed with no change, candidate, or blocked. Do not present a successful subset as a completed batch. Independent accepted items may be returned only when the output rule permits partial delivery and the incomplete remainder is explicit. A coupled set remains withheld until its group acceptance passes.

The handoff identifies the exact output, accepting authority and purpose, tests performed, unresolved limits, dependency disposition and receiving owner. Distinguish completed work, decisions for the recipient, substantive uncertainty and expressly deferred work, with responsibility and due stage where relevant. Keep unchanged carry-forward inputs read-only; do not resave or redeliver them unless the required package explicitly includes them.

Report only what occurred: inspected text is not a visual review; a successful save is not an open-file test; generated copy is not a sent message; a sample is not a whole-corpus audit. Later work may not reopen a passed artifact without an authorized target scope or defined reopening trigger.

4.5 Research-Grade Evidence and Analysis

Use AI to assist a declared research method, not to create an appearance of verification. Research records must let another competent person identify the question, evidence, operations, limitations, and basis for the conclusion.

Stage	Required method and evidence
Frame	State the exact question, decision, population or system boundary, non-scope, evidence standard, and what result could challenge the working hypothesis.
Plan sources	Fit source types to the claim and discipline; define access, search coverage, freshness, citation granularity, and exclusions. Use primary/official material where required, without imposing one source ladder on all domains.
Verify identity and access	Match title, author/issuer, stable identifier, source class, and version. Record what was actually retrieved; a resolving link or search snippet does not establish a complete source review.
Extract and interpret	Capture the outcome, comparison, population/context, intervention or configuration, timepoint, uncertainty, limitations, and adverse or contrary findings. Preserve qualifiers, exceptions and primary/secondary outcome distinctions that bound the claim.
Analyze and reproduce	Record search boundaries, sampling, transformations, calculations, configuration, and tests. Distinguish measurements, estimates, and assumptions; use M9 for material planning premises or scenario outputs.
Synthesize	Separate source statements from inferences and recommendations. Preserve mixed, null, and conflicting results; do not substitute a favorable secondary result for an unsupported primary claim. Apply Section 4.6 to material derived claims.
Review	Use the required technical or specialist review. When consistent human classification is necessary for acceptance, predefine a test set, reviewer independence, adjudication and acceptance standard. Apply Section 6.1 to the definition of metrics, diagnostics and decision boundaries.
Conclude and hand off	State the decision implications, coverage limits, unresolved questions, and conditions for reuse. Route source/date changes through M7 and dependent changes through Section 3.3.

Fact-state labels

Verified fact: supported by permitted evidence or observation within the stated scope. **Source-bounded fact:** attributed to an identified source/date, with access and currency limits. **Inference:** reasoned but not directly observed. **Hypothesis:** a proposition to test. **Assumption:** an explicit planning premise. **Owner decision:** an authorized choice, not empirical evidence. **Unresolved question:** a material unknown to qualify, investigate, or block.

Coverage and confidence

Record whether the inspection covered full content, selected excerpts, metadata, structured fields, rendered pages, external sources, or a tested sample. Do not upgrade a report's statement about earlier work into an independent review performed in the current task. Repeated source-derived recommendations are not repeated empirical confirmation.

A sample supports the inference allowed by its design and evidence standard. Do not certify uninspected units or generalize beyond the population, environment, outcome, and observation window actually supported. Lack of a detected defect is not proof of completeness; failure to retrieve evidence is not proof that it does not exist.

Treat a research agent's report and bibliography as candidate evidence, not proof of coverage. Check whether material questions, contrary results and source families were omitted. Compare the scope of the final summary with the actual sources; preserving citations while broadening population, certainty, causation or recommendation is not faithful synthesis. [E01, E02]

4.6 Claim, Inference, and Synthesis Control

Apply this section when a material claim depends on interpretation, comparison, aggregation, recommendation, or generalization beyond a direct source statement. Extend the claim record in M7 rather than creating a parallel evidence store. Use the granularity needed to inspect the reasoning; routine prose does not require a database object for every sentence.

Claim component	Required meaning
Formulation and type	State the actual output claim and location. Distinguish attribution, deduction, estimate, comparison, synthesis, recommendation, limitation, and owner decision where material.
Premises and derivation	Identify recoverable evidence, observations, arguments, or M9 assumptions and explain why they support this conclusion. A citation supporting a premise is not automatically support for the inference. Record value choices underlying a recommendation separately.
Scope and comparison	Name the relevant population, object, context, period, analytical scale, and comparison basis. Preserve material differences in meaning, measurement, conditions, and source terminology. Similar wording or function does not establish equivalence.
Challenge and uncertainty	Record relevant contrary evidence, negative cases, disanalogies, missing evidence, alternative explanations, and their disposition. Do not invent counterevidence; state limits on the search or test.
Interpretive ceiling	Specify a plausible stronger reading that is not supported where overgeneralization is a material risk. Keep necessary qualifiers adjacent to the released claim, not only in an inaccessible internal record.
Review and downstream use	Record owner, required review, accepted wording/version, permitted audience/use, and affected outputs through M7 and Section 6. Approval to use a formulation does not independently validate its premises.

Distinguish necessary from sufficient conditions. Evidence that is insufficient by itself may still contribute to a conclusion; a required condition may still be insufficient for acceptance. Showing that one proposal fails does not establish that an alternative succeeds. Support the alternative conclusion separately, including any claimed direction or magnitude of effect.

Support synthesis with a stated method

For a pattern or cumulative conclusion, state selection and coverage, the inference method, shared source lineage or dependence, counterexample treatment, and sensitivity to material inclusion, translation, classification, or modeling choices. Use discipline-appropriate review. A synthesis level or maturity label is needed only when the project defines it; no universal numbered ladder applies.

Distinguish a selected-case pattern from a representative or generalizable result. Repeated derivatives of one source are not independent confirmations. Demonstrating process capability does not validate every unit processed; release the actual supported units and claims under their applicable conditions.

Legitimate non-answers and model challenge

A question may remain unanswered for different reasons: insufficient sources, ambiguous wording, meaningful variation, a subject that cannot determine the answer, or categories that do not fit the subject. Record the actual reason. Do not invent a value to fill every cell or treat a non-answer as proof of absence. If categories or the schema repeatedly fail to represent the material, review the method instead of forcing the material into them. Revise through Section 3 before accepting dependent conclusions for use.

4.7 Context Continuity and Model-Assisted Work

Use these methods when model configuration, long context, staged generation or delegated work materially affects the result. For a simple low-risk task, retain only the information needed to perform and check it. These methods support the work-unit contract; they do not require a separate AI register.

Qualify the actual workflow

Distinguish generation or reasoning, retrieval assistance, research orchestration and actions through tools. Record material model/version or dated service identity, mode, relevant settings, tools, source states and review method in the existing run record. Mark unavailable configuration details unknown. A subscription label, larger context window, benchmark score or vendor description does not demonstrate fitness for the actual task.

Trigger	Action and observable check
Long or split context	The operator identifies required controls, sources and interfaces for each unit. Load relevant exact passages, not only summaries. Check that all required input portions were accessed; retrieval ranking and nominal context capacity do not establish coverage.
Compression, reset or handoff	Retain approved decisions, exact artifact/source bindings, changed requirements, unresolved questions and the next authorized action. Check the carry-forward record against those sources before dependent work; a summary is a route, not replacement authority.
Drift or repeated failed attempts	Pause the affected generation. Reconstruct the current contract from controlling inputs, remove superseded instructions from the active set and retry within the approved envelope. Preserve the failed result where required; do not keep elaborating a faulty premise.
Parallel or agent-assisted work	Apply Sections 4.1–4.4 to each material dependency. The integration owner checks actual outputs, missing components, shared assumptions and conflicting changes. Agent messages, consensus or a termination signal do not establish completion.
Material configuration change	Before reusing affected assurance, compare the changed model, mode, prompt, retrieval corpus or tools with the tested configuration. Repeat relevant task checks under M8.1; do not assume a newer or more capable system preserves behavior.

Context-loss and coordination studies support testing these failure paths, not a ban on conversation or a guarantee that compaction, more agents or a reset will fix them. Keep exploration flexible; reconcile state before reliance. [E03, E04, E17]

Reconcile the whole result

At integration, trace the central question, terms, assumptions and conclusion across the beginning, middle and end and across linked outputs. Check that later wording has not silently changed an earlier definition, omitted a required unit or converted a qualified result into consensus. Where full context cannot fit, inspect bounded units and their explicit interfaces; report the resulting coverage limit rather than claiming a whole-artifact review from a sample. Use M6.1 for constructive editorial work and Section 6 for acceptance.

5. Conditional Control Modules

Activate only the modules triggered by the work. They extend the common operating contract; they do not create mandatory files, fixed session counts, or universal project defaults.

Module	Activation trigger
M1 Preservation and completion fields	Protected content, fields, history, or behavior must survive.
M2 Pairwise and historical correction	Versions or issued records need reconciliation.
M3 Proof and restricted evaluation	Only limited material or staged access is authorized.
M4 External derivatives and live execution	Content crosses an audience, system, or live-action boundary.
M5 Reusable-master refinement	Project lessons are being abstracted into reusable logic.
M6 Content, language and audience	Content is being created or substantively revised, or audience fit, wording or warning density needs calibration.
M7 Source and date integrity	Source identity, volatile facts, dates, or currentness matter.
M8 Technical and package integrity	Technical behavior, exact payload, or package integrity matters.
M9 Assumptions and scenarios	Material premises or model outputs affect a decision or claim.
M10 Adoption and value realization	The objective or commitment depends on uptake or post-delivery benefits.

M1. Preservation and Completion Fields

Classify protected material before editing. Preserve exact-use text, evidence, legal/safety controls, identifiers, formulas, links, code/data behavior, signatures, and historical records according to their actual function. Unusual formatting or an empty field is not automatically residue.

Field class	Completion or preservation rule
Approved policy or proposed term	May be stated from its proper authority as a proposal or rule. It is not evidence that a fact, execution, or acceptance event occurred.
Unknown / intentional placeholder	Keep explicit until a supported value and authorized decision exist. Use “not applicable” only when applicability has actually been decided.
Prerequisite field	Complete before the specific action that consumes it, such as signing, access opening, calculation, or release.
Later-event fact	Populate delivery, closure, verification, acceptance, or other observed facts only when the event occurs. Do not require future facts to complete an earlier stage.
Operative or historical field	Preserve meaning and behavior. Keep signed/issued records unchanged; use a distinct future-use copy for correction.

For a material conditional field or requirement, define the trigger, required occurrence count or cardinality, allowed values, permissible uncertainty, due stage and consequence if unmet. Distinguish absent optional information from missing required evidence. A not-applicable determination records why the requirement does not arise; it cannot substitute for unavailable evidence, consent, authorization or review. Undefined choices that change a method or governed obligation require the appropriate authority decision, not an invented value or cosmetic rewrite.

Specify each material field's owner, evidence source, and due event. A blanket “all fields complete” gate is invalid when it includes facts created after the gated action. Retain draft/template labels while accurate; change them only when the authorized instance transition makes them inaccurate, without deleting substantive restrictions.

M2. Pairwise, Historical, and Protected-Copy Correction

Select the base by authority, role, intended use, and qualification—not apparent recency. Protect issued records and unique material before comparison. Use Section 3.2 to identify original-source, corrective-baseline, carry-forward, or historical treatment.

Map approved imports to the target's current structure and preserve non-conflicting intervening changes. Do not restore old text to satisfy a stale locator or silently combine incompatible versions. An unmappable requirement becomes a scoped conflict. Apply the structural and identifier rules in Section 3.6.

Retain a concise internal disposition for each meaningful difference: import, preserve, retire, no change, or block. Deliver the authorized future-use output, not an altered historical record. Missing source history limits the claim of preservation even when the corrective output is fit for a named downstream purpose.

M3. Gate, Bridge, Proof, and Restricted Evaluation

Control	Required answer
Recipient and decision	Who may see the material, why this is the appropriate initial reviewer, and what decision or routing action is within their remit?
Parent and sample identity	Which approved parent, version, selected units/pages/records, and reproduction method supply the proof?
Released and withheld scope	What is shown, summarized, omitted, reserved, or unavailable? Which context must remain adjacent to an extract?
Use and technology	What viewing, copying, notes, export, onward sharing, automated processing, retention, or reuse is allowed, by which authority?
Evidence ceiling	What does the selected material demonstrate, what remains uninspected, and what requires full review or separate validation?
Next gate	What smallest action is appropriate at this stage: request, route, decline or proceed under separately approved access?

A proof packet does not imply full access and need not be compulsory. Selected content may demonstrate architecture, not operating proof. Allow alternatives when prerequisites are met. Confidentiality alone does not open review: confirm the actual access scope, whether combined in the same instrument or separately authorized. Access does not itself transfer ownership or authorize deployment. Appendix D2 routes transaction-specific choices.

When proof inserts are locked, verify their preservation and readability without regenerating them or claiming renewed source authentication. Searchability of a wrapper does not make an embedded image searchable or validate omitted parent content. Adding extraction text, optical character recognition (OCR), or a different representation is a separate transformation requiring authorization.

Guardrail placement. An independently forwarded page or extracted unit may need its own short source, omission, and use boundary. That is functional repetition, not duplicate canon. Keep the substantive restriction in the instrument that has authority to impose it.

M4. External Derivatives and Live Execution

A derivative inherits its approved parent's provenance, claim ceiling, and use limits. A summary, translation, slide, API payload, media variant, training/retrieval set, or public page does not acquire new authority merely through conversion. Preparing it is separate from permission to send, publish, ingest, deploy, or run it.

Control	Required treatment
Parent and transformation	Identify the source/version and allowed compression, anonymization, localization, format change, extraction, or system use.
Rights and environment	Check permission, attribution, data class, consent basis where relevant, approved tools, retention, and prohibited uses. Authorized storage/logging must be distinguished from prohibited training or reusable extraction.
Recipient / endpoint	Immediately before action, verify the intended recipient or endpoint, address/interface, role, authorized purpose, and route state.
Exact action	Freeze the approved payload/configuration, execution mode, attachment/export decision, and maximum authorized action scope. Use M8 for interface and approval checks.
Duplicate and capacity control	Define active, backup, held, suppressed, expired, and closed routes where relevant. Assign stage-specific concurrency and operator capacity; avoid duplicate or conflicting actions.
Evidence and stop rule	Record executor, date/time, exact payload/configuration, result, warnings, exceptions, and next action. Stop on unexpected restrictions, identity/rights uncertainty, or unsafe failure.

Stage-specific exposure

Set capacity according to confidentiality, reconstructability, reversibility, workload, recipient obligations, and failure consequences. A broad low-disclosure introduction and a narrow protected evaluation may justify different limits. Do not import a project's numerical cap, timing rule, or platform assumption into every stage.

Readiness and scale

Gate	What it establishes—and what it does not
Content / technical	The representative output and exact final payload meet specification and behave in the target environment. This does not create action authority.
External-action	Recipient/endpoint, permissions, exact action, and stop conditions are ready for the named use. Availability of an interface is not consent.
Pilot / field, when needed	A bounded real-environment run provides evidence within its defined sample, metric, and observation window. It does not authorize automatic expansion.
Scale, when needed	Predeclared evidence, sample sufficiency, defect/risk assessment, and owner approval justify a defined expansion—not unrestricted use.

Require pilot and scale gates when the risk or proposed expansion warrants them. A single routine authorized action does not need a fabricated pilot program. Do not make repeated attempts after a warning without a fresh permitted disposition.

When agents consume untrusted material, the operator limits tools, credentials, data access and permitted destinations to the authorized task. Enforce consequential action boundaries outside generated prose where required; validate proposed actions before execution. Test both unsafe actions and legitimate task completion. Prompt warnings, structured output or a guardrail alone are not proof of isolation. If required containment cannot be enforced, hold the affected action. [E10, E11]

M4.1 External Roles and Interaction Evidence

Activate these controls when a named person or organization carries an external dependency, relationship claim, review function, or adoption route. Keep one canonical entity record with role-specific links where useful; do not mistake a candidate list for an available delivery or assurance team.

Relationship control	Required distinction
Identity, role, and gap	Identify the function needed, relevant stage, current scope, and source basis. A public title or technical capability does not establish budget, purchasing authority, availability, or commitment.
Agreement and permissions	Keep agreement status, access authority, and naming/disclosure permission separate. A confidential exchange or agreement supports only its actual scope; endorsement, partnership, or public-name permission requires an express basis.
Functional sufficiency	For a dependency-critical role, require current evidence of competence, scope, capacity, interoperability, and access or agreement as appropriate. Assess substitutes and common dependencies; multiple names are not independent fallback capacity.
Candidate and active states	Separate watchlist relevance from an active authorized route and from a closed dependency. Re-screen when the role, scope, jurisdiction, agreement, evidence, or permission changes. Only the responsible gate owner can accept closure.

Record interaction events

Use the event record for what occurred and a derived summary for current activity. Events may overlap, branch, or be unavailable to observe. Retain date, actor, actual material/version, scope, evidence source, result, and next authorized action. Reconcile contradictory summary labels to the supported history; do not overwrite the history to match a label.

Observed event or state	What it does not establish
Prepared or authorized	A draft or approval is not a send, upload, publication, or live run.
Sent, delivered, or access opened	Technical transmission or access availability does not establish receipt by the intended human, actual review, or substantive interest.
Reviewed or substantively discussed	Evidence of review or questions does not establish a favorable judgment, budget, mandate, or agreement.
Agreed or approved	Only the supported scope is established; effectiveness, conditions, execution, and realized results may remain outstanding.
No response or unresolved outcome	Record the observable non-response, channel, and observation window. The cause, actual reading, and future interest remain unknown unless separately evidenced.

Map functions as needed without automatically contacting every person mapped. Coordinate active and backup routes under M4. M10 interprets outcome evidence and continuation; it does not create outreach permission or default contact counts.

M5. Reusable Master and Template Refinement

A reusable master contains portable decision logic and explicit overlay points. It is not a compilation of project memories or all available instructions. Preserve the project-independent architecture while admitting only lessons that improve execution.

Step	Decision and acceptance test
Classify the source lesson	Distinguish a prescribed solution, a reported observation, an implemented result, and an independently verified outcome. Identify scope and lineage; related versions are one source family, not independent confirmations.
Check existing coverage	Mark the candidate as already covered, strengthen, new, or reject. Compare meaning, failure mode, and required action—not just terminology.
Abstract	Remove project identities, dates, parties, counts, routes, local tools, commercial posture, fixed schedules, and hidden defaults. Preserve the triggering condition and authority limit.
Admit	Require a reusable change to a decision, failure mode, preservation duty, handoff, acceptance test, or user task. A useful project preference is not automatically a universal rule.
Place and consolidate	Revise the existing canonical treatment first. Merge equivalent candidates; do not create a module, status, or record for each observed example.
Parameterize	Expose project choices through explicit overlays: audience, risk, evidence standard, law/instrument, data class, format, capacity, review, and release purpose.
Test	Apply the rule to relevant project classes and boundary cases. Check contradiction, redundancy, actionability, source support, and whether local assumptions survived.
Stabilize	Verify the complete guide and its compact template/prompt mirrors, then freeze its controlled identity and defined reopening conditions.

Admission examples

Portable: different exposure stages may need different concurrency limits. **Not portable:** every project may have only a fixed number of active recipients. **Portable:** exact corrections require a declared match and change envelope. **Not portable:** every revision must use one file format, tool, or fixed session count.

Portable: a completed source-identity check does not close methodological or specialist review. **Not portable:** every source registry must use the same evidence codes or reference numbers. These examples illustrate abstraction, not default project policy.

Genericity regression

Use a small declared set of relevant scenarios, such as a legal suite, a research synthesis, a software release, a policy rollout, public communications, a data migration, and a publication package. Check that the rule remains usable, proportionate, and correctly bounded in each applicable scenario. Record exceptions where the rule does not apply rather than forcing universal coverage.

When refining content-production guidance, include representative instructions, explanations and educational or expressive work where relevant. Examine actual outputs for support, completeness, preserved meaning, appropriate voice and reader use or experience. Record the tested tasks, defects and limits; clearer instructions alone do not prove better downstream results.

Keep source-to-lesson rationale internal; the master stays self-contained and free of case-specific or revision narrative. Treat advice about effectiveness as a hypothesis until supported by outcome evidence; inspect reported trials before diagnosing causes or prescribing them again. Assess reuse through the components retained, adaptation required and revalidation effort. Resemblance or an unsupported percentage does not demonstrate savings. Internal cross-references help readers find authoritative detail.

M6. Content, Language, and Audience

Apply M6–M6.1 while planning, drafting and reviewing reader-facing content, within the authorized scope. Make the object, present function, reader consequence and limit clear. Be definite about what is known or proposed and conditional about what remains uncertain. Select language choices in the existing project record; no separate style document is required.

Use precise terms and clear instructions

Use a stable name for each defined concept, object or action. Prefer familiar words when their meaning is equally precise; retain necessary technical terms and explain them for the intended reader. Define abbreviations at first useful mention. Avoid ambiguous pronouns, unexplained aliases and changes of name made only for variety. Distinguish requirements, recommendations, permissions, prohibitions and possibilities; editing must preserve their force.

Build sentences around a clear main action or assertion. Separate independently actionable steps and reduce nested clauses, noun strings and ambiguous negatives. Keep the condition or exception beside the statement it qualifies. In procedures, put prerequisites before the governed action and give steps in the order required; preserve actions that must occur together. Use lists when order or comparison helps. No universal sentence-length limit or compulsory sentence pattern applies.

Control	Required treatment
Present object before prospective use	Define what exists or is being offered now. Separate its completeness, review status, release permission, and optional future adaptations. Pending review may still identify substantive work.
Recipient logic	Connect a relevant source-bounded fact to a concrete feature and the recipient's actual remit. Do not infer budget, demand, endorsement, authority, or a division of responsibilities from a title alone.
Actors and actions	Name the actor, object and consequence where they matter. Prefer active wording when responsibility would otherwise be unclear. Replace abstract noun clusters and ambiguous references with the specific action or event. Distinguish review, approval and release by their object, actor and scope.
Functional repetition	Keep a local statement when it serves orientation, learning, an expressive purpose, independent forwarding, safety adjacency, routing, a legal function or a distinct decision. Consolidate full duplicate treatments that serve no such purpose.
Safeguards and warning density	Classify proof, source, safety, rights, privacy, access, AI-use and deployment limits before shortening. Do not replace an operative restriction with a pointer to an inaccessible internal policy. Assess qualifications across the connected passage, table or extract that will travel together. A sentence released alone must retain its material limits; a connected unit need not repeat the complete qualification inventory in every sentence. The released unit must not support an unwarranted stronger reading.
Priority and form	Give the principal idea, task or decision appropriate prominence for the intended form. Match explanation, density and technical vocabulary to the reader's knowledge and purpose. Preserve disciplinary voice, intended register, language and variation appropriate to each artifact family.
Measurement	Use reproducible exact-copy counts or interface tests after wording is frozen. Editorial length ranges are project choices, not platform limits or evidence of performance.
Usability	Test the artifact alone, in its expected sequence and, where relevant, out of order. Select a cold read or representative task, such as finding a rule, following an instruction, explaining a conclusion or assessing the intended experience. For a task with a determinate answer, define the expected answer or permitted action from governing requirements and evidence before testing. Record the actual answer and route used. A missing distinction, contradiction or unsupported stronger action fails the task. Record method, coverage and result under Section 6; do not claim reader testing that was not performed.

In substantive passages, distinguish explanation of the subject from commentary on how the document was constructed. Replace author-facing statements about layers, architecture, coverage or avoidance of duplication with the relevant fact, distinction, consequence or reader action. Retain construction terminology when structure is itself the subject or the wording provides useful navigation.

Orient an unfamiliar reader

In an initial message to an unfamiliar external reader, identify the object and purpose at the opening, give enough substance to make it intelligible, explain its relevance to the recipient, preserve the material boundary, and request the smallest appropriate next decision. An initial message usually benefits from one primary response action; formal forms and task workflows may require several.

A profile, post, attachment, or prior conversation may support context, but the message should not depend on unverified availability or prior exposure. An alternative first-contact route must not imply referral, endorsement, or prior nonresponse unless that event is established.

Compression test

Judge a passage by its contribution to meaning, evidence, explanation, protection, action or the intended reader experience. Remove repetition that serves no such purpose. Preserve exact factual anchors, distinctive safeguards, useful transitions, necessary repetition and expressive detail. Assess sentences in context rather than requiring each to introduce a new fact. Use outcome evidence before claiming that shorter or more polished writing improves real-world performance.

M6.1 Substantive Construction and Editorial Review

Apply this method when creating or substantively revising explanatory, argumentative, instructional or creative work. The exact-edit envelope still controls protected wording. Judge quality by the substance, reasoning, expression and intended reader use or experience; fluency alone does not establish quality.

Establish the basis for drafting

The author or executor uses the existing blueprint or task record to identify the reader's question, practical task or intended experience; the contribution of the work; the available facts, observations, arguments and examples; and material gaps. Distinguish uncertainty that the current activity can tolerate from prerequisites due before reliance. Learning, exploration and artistic purpose remain valid under Section 3.1.1. An outline guides development; it does not oblige the author to fill every heading or predetermine the conclusion.

Develop the explanation or argument

Build each substantial unit from relevant substance. Explain how or why a result follows, connect claims to reasons and supply the intermediate steps the reader needs. Where a relationship matters, identify what changes, the supported intermediate effect and the consequence for understanding or action. A list of factors or desired outcomes does not explain that relationship. If the connection is not established, identify the gap without inventing a mechanism. For recommendations, state material conditions and tradeoffs. Section 4.6 governs material inference.

Adapt the structure to the work. Instructions support correct action; research and analysis support understanding and judgment; educational work may use analogy and purposeful repetition. Fiction and expressive work may follow an artistic purpose, including deliberate ambiguity. Preserve the intended experience while resolving accidental confusion. No universal argument, objection or paragraph formula is required.

Preserve specificity and intention

State the actor, object, mechanism, distinction or consequence actually meant, using supported detail. Preserve scope, emphasis, disciplinary meaning, register and relevant cultural or linguistic particulars in paraphrase and translation. Use authorized examples of voice where helpful. Do not normalize unfamiliar expression into a generic house style. If material meaning is unresolved, preserve it and obtain the necessary author or domain judgment.

[E08]

Use examples and alternatives honestly

Choose examples that illuminate the claim or task. Distinguish a constructed illustration, analogy or hypothetical from an observed case whenever confusion is plausible. Check calculations, internal logic and material limits of analogy. For decision guidance, use contrasting cases where helpful: identify which condition changes and why the permitted conclusion or action changes. Include limiting or adverse cases and alternative explanations where material. Examples do not substitute for evidence; no fixed case or paragraph pattern applies.

AI-assisted ideation can be useful without defining the author's purpose or guaranteeing originality. Where diversity matters, consider materially different approaches rather than only polished variants of the first suggestion. Judge the work against its purpose; neither model use nor model avoidance establishes quality. [E09]

Diagnose the gap before revising

Distinguish four problems. Unclear wording needs a more precise expression of established meaning. A missing explanation needs the relevant mechanism or reasoning. An unsupported claim needs evidence, justified qualification or removal. Unnecessary content may be removed only when no required function or safeguard depends on it. A plausible detail cannot repair an evidence gap.

For example, "the system improves coordination" may need a description of how information moves between people, evidence of an improvement, or a narrower claim about an actual capability. Select the repair supported by the source and task. Removing promotional adjectives alone does not establish the underlying claim.

Check evaluative and rhetorical wording for claims about ranking, frequency, probability, causation, magnitude or intent. Supply the necessary comparison, evidence or attribution, or state only the supported observation, relationship or requirement. Treat this wording as a diagnostic signal; neither word bans nor hedging establishes support.

Review substance and organization

Review purpose, required coverage, support, reasoning and omissions; apply Sections 4.5–4.6 to source-dependent claims and derivations. Then check whether the reader can follow the organization, steps and conclusion. A reverse outline—each section's contribution in one sentence—can reveal repeated framing, generic advice, missing reasoning or ornamental categories. Compare distant passages under Section 4.7; preserve purposeful repetition, local safeguards, technical vocabulary and narrative variation.

Check titles, headings, captions, introductions and summary labels against the content. Labels such as "worked example," "calculated result," "comprehensive review" or "validated method" require corresponding substance and evidence. A worked example may be qualitative if clearly identified and developed for its task. Narrow overstatements or supply authorized missing substance.

Repair the demonstrated omission, contradiction or weak explanation at its canonical home within the authorized scope. Consolidate functionless duplication and supply supported substance where needed. Route unresolved prerequisites through Section 1.3 and dependent changes through Section 3.3. A fluent draft or a self-reported pass does not demonstrate that the content review occurred. [E02]

Refine language and recheck meaning

After assessing substance and organization, refine terminology, sentences, repetition, rhythm and tone under M6. Drafting may combine these activities; final polish follows substantive assessment. Remove empty or distracting stock framing, slogans, forced contrasts and patterned phrasing. Retain effective metaphor, purposeful parallelism and voice. Avoid universal word bans, paragraph templates and numerical writing scores.

Compare the revision with its sources and intended meaning. Check actors and actions, technical terms, scope, exceptions, quantities and units, dates and time bounds, causal strength, uncertainty and the force of obligations or permissions. Examine conditions, comparisons and negation for reversed thresholds, false alternatives or unintended exceptions. Where consequential, test a concrete case against the governing rule. Preserve voice, emphasis and relevant cultural meaning. Check summaries and conclusions against the supporting body, then repeat affected content and dependency checks.

Apply the M6 reader test to the finished work. For repeated production, retain representative, authorized examples in existing records to compare support, completeness, clarity, appropriate expression and avoidable review effort. Record defects, corrections and limits under Section 6. Language checks alone cannot establish factual support, sound reasoning or suitability for use.

M7. Source Identity, Claim Support, and Current State

Activate this module when a source, volatile fact, verification boundary, event, or dependency matters to the proposed claim or action. Refresh the designated current-state home first, then review only affected consumers.

Source identity is not substantive validation

Check stable identifier, issuer/author, title, source class, and relevant version against the actual material. Then determine what was accessed, extracted, interpreted, and reviewed. A repaired citation or working URL can resolve identity without resolving methodological quality, claim support, specialist review, or permission for reuse.

Source state	Permitted treatment
Verified for the named use	Retain the source and actual check event/date with its scope, claim ceiling, and access limits. Do not label unrelated material freshly verified.
Corrected or moved	Repair the locator/identity and affected references. Reassess the claim only when the change affects its support or the declared review requires it.
Restricted or partially accessed	State what is and is not accessible. Do not imply full-content review from metadata, snippets, or a limited extract.
Historical or superseded	Retain for history, traceability, or a stable proposition that remains supported. Age alone does not invalidate evidence; do not use superseded requirements as current authority.
Withdrawn or reliability challenged	Identify the withdrawal or concern and its effect on the supported claim. Hold current reliance where unresolved; preserve the record and uncertainty rather than erasing the history.
Unavailable, removed, or replaced	Record the access/result limitation; substitute only an appropriate source within scope. Absence of retrieval is not proof of absence.
Unresolved	Do not assign a fabricated check date or unqualified current authority. Narrow, qualify, investigate, replace, hold, or block the affected use.

Claim-level control record

For a proposition that materially supports a decision or conclusion, record source and pinpoint, bounded claim, source/access/version state, verification object and date, interpretation limits, owner, affected consumers and next trigger. Extend this same record under Section 4.6 for material derivation, or reference an M9 assumption. Different propositions may have different dates and review states even when they share a source.

A source registration can be identity-only. It must not silently create an evidence profile, improve a grade, authorize use, or refresh every related claim. Where IDs and published reference numbers are allocated separately, check collisions and reconcile the approved mappings before dependent use.

Reassess holds precisely

When one issue closes, remove its obsolete reason from the active hold and record the remaining basis. Do not keep a false warning merely to remain cautious, and do not clear all review or release restrictions because one bibliographic, access, or identity check succeeded. Use Section 6.3 to distinguish an active resolution task from a settled current-use restriction.

M7.1 Date Semantics and Event Triggers

Every material date must identify what happened, to what object, and in which scope. Preserve historical dates unless evidence establishes an error; do not equalize dates merely for visual consistency.

Date type	Meaning	Must not imply
Edition / study baseline	The defined starting boundary of an edition, dataset, review, or study.	That later propositions were checked on that date or that no later checks exist.
Source publication / revision	When the issuer published or revised a particular source/version.	When the project accessed it or accepted its interpretation.
Source access / retrieval	When the material was actually obtained or observed.	That all of its content was verified or methodologically validated.
Proposition verification	When a specific factual claim, record, or dependency was checked against named evidence.	A whole-document or whole-population refresh.
Review event	When a specified assessment was conducted, with its actual result.	That it passed or an acceptance/use approval followed.
Acceptance / approval	When a designated authority accepted a stated result for a named purpose.	The source-check date, effective date, or broader permission.
Technical save / export	When an application wrote or exported a file.	A substantive review, new edition, source cutoff, or release event.
Scheduled event / deadline	An expected event or deadline with a defined later check.	That the event occurred, succeeded, or produced an assumed outcome.
Effective date	When a stated instrument, rule, or permission becomes effective within its scope.	That separately conditioned provisions commenced, operation started, or implementation completed.
Conditional commencement	When prerequisites for beginning a defined activity are evidenced as satisfied.	That later approvals, operational requirements, or successful execution are established.
Observed execution / delivery	When an authorized action, receipt, closure, or verification actually occurred.	Other events or acceptance not established by that observation.

Transitions require their own evidence

Define which conditions connect the states. Adoption may precede effectiveness; effectiveness may depend on funding, permissions, configuration, or another trigger; actual operation requires its own evidence. This is an event dependency, not a reason to add the same status column everywhere. Apply the existing dimensions in Section 2.3.

For each volatile proposition, specify a review trigger and a reuse limit where appropriate. After a scheduled milestone or expiry, recheck the affected proposition before presenting it as current. Do not infer adoption from consultation closing, completion from a due date, or successful execution from an approved plan.

Bounded refresh. A record may retain its original baseline while adding a later check to one proposition. Preserve untouched dates, review affected dependencies, and record “Reviewed—no change” when warranted. New verification and editorial acceptance are distinct events even when they occur on the same day.

When official pages, cached versions, indexed text, or source copies disagree, record the conflict and use only the narrow proposition supported by the permitted evidence. Do not conceal it by advancing every date or adding a general disclaimer.

M8. Technical, Document, and Package Integrity

Define the technical acceptance standard before editing or conversion. Preserve the accepted design and behavior unless an authorized change or observed defect requires repair. Apply Section 3.5 when the change envelope is closed.

Surface	Required checks
Opening and identity	The first page/screen identifies object, role, status, and boundary without burying the substance in control clutter. Titles, filenames, properties, legends, manifests, and references agree by function.
Styles and accessibility	Use meaningful headings and captions, logical reading order, appropriate table headers, usable links, and required alternatives. Correct demonstrated defects, not unverified automated heuristics.
Tables and layout	Preserve valid widths, merge structure, layout mode, and intentional variation. Keep reconciliation IDs and short field labels readable without reconstructing split fragments.
Readability	Set the project's minimum text/readability floor. Reflow, paginate, adjust authorized spacing/widths, or shorten permitted text rather than shrinking below it.
Fields and references	Test formulas, page fields, content controls, cross-references, bookmarks, and hyperlinks. Refresh only values made stale or expressly required; a general QA step does not override a no-refresh instruction.
Exact payload	Count or validate the final spaces, breaks, fields, dimensions, duration, size, schema or API contract using a declared reproducible method. Distinguish a verified interface limit from a project target.
Code, models, and data	Test calculations, schemas, identifiers, field-to-vocabulary mappings, transformations, dependencies, configuration, environment, security, exports and reproducibility. Check approved package/API identity and version before use; an existing name is not proof of trusted origin or suitability. Test required behavior and failure cases, not syntax alone. Apply M1 to conditional fields and required links. [E12]
Media and derivatives	Check dimensions, source assets, legibility, captions, alternatives, embedded rights/attribution and export behavior. Validate material chart values, labels, counts and spatial relations against source data or direct inspection; generated captions and image recognition can be wrong. Distinguish illustrative media from evidence. A successful conversion does not establish fidelity. [E14]

For a chart, diagram or exhibit that explains substantive content, identify the reader question, intended message, authoritative inputs and permitted representation before construction. Visual order, size, position, colour and connectors must not imply unsupported magnitude, ranking, causation or completion. Where the inputs support only a qualitative representation, preserve that limit in the visual and its explanation.

M8.1 Change-sensitive verification

Bind each count, hash, test, approval, or release permission to the artifact representation, scope, and environment it actually covers. After a change, identify which evidence depends on the modified content or bytes. Recompute exact-byte values and exact-payload checks; repeat only the other affected tests and reviews, with an explicit carry-forward basis for unaffected evidence. Use Section 2.2 to identify the changed source, derivative or release/access instance.

Preserve previous signatures, approvals, and test records as history. They are not transferred to altered content. An unsigned successor must obtain any approval required for its new scope. A changed footer may leave an earlier substantive review relevant while still changing the file hash and requiring layout verification.

Technical success is scoped. Opening, rendering, searching, compiling, or passing an interface check establishes only that particular result. None alone establishes factual validity, accessibility conformance, safety, ownership, institutional approval, or readiness for unrestricted use.

M8.2 Preservation Baseline and Package Closure

When visual, structural, metadata or behavioral preservation is required for acceptance, capture the incoming state before editing: identity, relevant content and structures, protected components, current properties and a complete render or equivalent observable baseline.

For reflow or restructuring, compare each protected value, qualification, link and attached component by stable identity and parent association. Equal totals can conceal swapped values, detached warnings or duplicated records. For presentation-only changes, define how to reconstruct the original ordered content for comparison, excluding only authorized differences. Count logical content separately from repeated display elements.

Minimum technical disturbance

Edit the smallest safe object, paragraph, cell, style, relationship, property, code path, or package part. Do not re-template, normalize, re-encode, or regenerate a valid artifact merely for tool convenience. If the authorized operation cannot be made within the preservation envelope, stop and obtain a method or scope decision.

Before applying a repeated layout transformation widely, test representative short, long and exceptional units in the intended renderer, including attached qualifications and continuation behavior. Record the renderer and material font dependencies. A successful sample supports wider application within the tested conditions; final inspection remains required.

In paginated outputs, continued records and tables must retain their identity and column meaning. Repeated labels must be truthful on every segment. Keep compact units together when they fit on an otherwise empty usable page; allow longer units to flow with a meaningful opening and identifiable continuation. Keep qualifications associated with their parent content. Judge page breaks and whitespace by reader function rather than page count, and verify rendered behavior instead of relying solely on formatting properties.

Closure area	Inspection and disposition
Visible content	Body text, tables, headings, captions, identifiers, footnotes/endnotes, labels, and displayed field values.
Non-body surfaces	Headers/footers, text boxes/drawing text, content controls, notes, comments, review markup, hidden ranges, sheets, or layers.
Technical package	Relationships, linked targets, embedded objects/media, custom data, variables, formulas/instructions, scripts/configuration, package-part names, and output names.
Metadata and privacy	Titles, author/organization fields, timestamps, revision values, template paths, editing history, and custom properties. Apply the approved retention and metadata policy—not blanket erasure.
Residual candidates	Process references, obsolete identities, retired controls, unsupported state, duplicate canon, leaked internal instructions, personal data, and unintended predecessor narrative.
Render / run result	Clipping, overlap, broken tables or row splits, missing glyphs, stale fields, orphaned fragments, unintended blank pages, unreadable identifiers, failed tests, or behavioral drift.

Search hits and anomalies are candidates, not edit authority. Classify the complete context and correct only authorized defects. Preserve required attribution, source/status dates, audit evidence, legitimate pedagogic terms, hidden computational structures, and legally required retention. “Clean” means free of unintended residue, not free of provenance or all historical information.

Closure loop

After correction, rescan the affected surfaces, save, reopen, and render or retest. Inspect every page or other output unit required by the declared acceptance scope; document any sampling limit. For final document delivery, inspect the complete final render, not only the pages expected to change. Repeat until the applicable tests pass or a concrete blocker remains.

Keep render files, diffs, scripts, and QA inventories as internal support unless the output contract requires them. Do not add them to the deliverable or core package by default. Preserve the original input and confirm that read-only context has not changed.

M9. Assumptions, Scenarios, and Falsification

Activate when a planning premise, range, forecast, scenario or model result materially affects a decision or claim. Maintain it in the existing assumption, model or evidence record. A rule chosen by an owner states what is required; an assumption describes a condition taken as given. Confidence in either is not a gate pass.

Material premise field	Minimum content
Identity and type	Stable reference; proposition/value/range; units, boundary, time horizon, and whether an assumption, target, scenario input, forecast, or governed rule.
Basis and confidence	Owner; supporting sources or explicit judgment; uncertainty; actual check object/date. An administrative edit is not new verification of the premise.
Permitted use and consumers	Allowed planning, comparison, internal, or outward-facing use; claim ceiling; dependent decisions, calculations, and artifacts.
Disconfirmation and sensitivity	Evidence or result that would challenge the empirical premise; important alternatives or sensitivities. A policy choice instead needs an authority and suitability-review basis, not a fabricated empirical falsification test.
Review and disposition triggers	When to recheck, narrow or suspend use, retire, or replace with better evidence. Identify the responsible decision maker and required action; use Section 3.3 for affected consumers.

Numerical and economic discipline

Do not promote an illustrative scenario to a forecast, commitment, valuation, or underwriting basis merely because it contains numbers. A model-output claim requires the recoverable versioned calculation, inputs, definitions, transformations, environment where relevant, and reconciliation. A reference to an absent workbook is only a route. Numerical ranges are not confidence intervals unless the method establishes that meaning.

State denominators, baselines, currencies/price dates where applicable, boundaries, and interactions among assumptions. Keep planned performance separate from measured results. Avoid double counting: a beneficiary's benefit, a receipt, and an avoided cost may describe the same value transfer. Use an explicit counterfactual and distinguish gross benefit from net contribution and remaining implementation or validation costs.

For AI-assisted analysis, specify the estimand or intended quantity and material choices in cleaning, exclusions, transformations, model specification and uncertainty treatment. Reconcile reported values to executed code and actual outputs. Examine plausible alternatives when they could change the conclusion; distinguish exploratory choices from prespecified tests. Do not select a model, persona or rerun because it produces the preferred result. An automated methodological audit does not remove analytic uncertainty. [E13]

Narrow, retire, or replace on evidence

A review trigger calls for investigation; a downgrade trigger narrows or suspends the allowed use; a retirement trigger removes the active premise. Preserve the former state and reason. Measurement may replace a planning assumption only within the boundary actually established. An adverse result does not automatically settle cases outside that boundary, and a favorable result does not erase contrary findings.

When the basis deteriorates, reassess the dependent decision rather than moving thresholds to preserve a desired answer. A changed rule or method requires its own authorization and prospective acceptance basis; it cannot retrospectively convert a failed qualifying run into a pass. Use Section 6 for technical/release consequences and M10 for discretionary investment decisions.

M10. Adoption and Value Realization

Activate only where the objective or discretionary commitment depends on another actor's uptake, funding, purchase, authorization, implementation, or benefits realized after delivery. This includes noncommercial uses, but not merely possible future commercialization. Section 3.1.1 permits purpose-led creation without measurable utility or monetization. This module tests a necessary realization path; it neither defines all value nor overrides truth, quality, rights or safety.

Outcome-case question	Required decision or evidence
Outcome and owner	State the observable intended value event, beneficiary, accountable outcome owner, and assessment period. Keep artifact completion, authorized use, uptake, and realized benefit distinct.
Problem and alternatives	Identify the purpose or need and alternatives material to this decision, including simpler change or no discretionary action. Existing alternatives do not veto the project. For novelty or "missing framework" claims, define search scope and date; not finding a substitute is not proof none exists.
Decision and adoption path	Distinguish user, reviewer, sponsor, budget holder, approver, implementer, and acquirer where relevant. Record which roles and mandates are confirmed versus hypothesized; a large sector or list of possible users is not an addressable buyer base.
Burden and admissibility	Identify review, specialist credentials, rights, editorial, integration, procurement, safety, and maintenance requirements. Who can perform and fund them? A recipient's admissibility rule may block a route regardless of artifact polish.
Value and remaining effort	Explain what the work saves, enables, or changes and the resource burden still required to realize that benefit. Keep hoped-for price, estimated replacement effort, and supported transaction value separate. Use M9 for material numerical premises.
Evidence and next commitment	Define the smallest authorized test or milestone likely to resolve the key uncertainty, its signal, observation window, and resource allowance under Section 3.1.1. Use M4 for live actions and Section 6 for release conditions.
Allowed realization paths	Record the owner-approved acquisition, license, publication, adoption, collaboration, funded development, or other route. Optional alternatives are choices, not authorization to change the asset, promised deliverables, rights, or support model.

Test pull without creating false prerequisites

Where feasible, test the central use or acquisition hypothesis before full production using a bounded concept, prototype, sample, or discussion that exposes only authorized material. Sensitive or long-horizon work may instead justify a recorded uncertainty budget or mandate. Define the evidence needed for the next commitment, not a universal requirement for a buyer, signature, or predetermined number of replies.

An expression of interest, confidentiality agreement, review request, or meeting may be evidence of a narrow intermediate event. None alone establishes demand at a stated price, review success, funding, adoption, or realized benefit. Use the event distinctions in M4.1 rather than a single "engaged" label.

M10.1 Outcome Diagnosis and Continuation

Evaluate the evidence for the declared outcome separately from technical acceptance. Review what was actually tested before proposing another intervention. Use the existing outcome, activity, or decision record rather than creating a parallel campaign or benefits ledger.

Diagnostic control	Required treatment
Reconstruct the test	Identify the audience or cohort, actual channel, material/version, permissions, exposure, timing, and observation window. Separate source-reported history from reviewed execution evidence. Do not claim coverage of unobserved stages.
Retain alternative causes	Poor response can remain consistent with different explanations: route failure, unread material, weak fit, timing, credentials, review burden, terms, price, or no current need. Do not infer a definitive cause from silence; an explicit objection supports only the stated case.
Choose a discriminating next test	State which uncertainty a new action would resolve and what changes from the prior attempt. Do not prescribe an already-tested route without a material difference, new evidence, or an explicit reason for a bounded retest.
Check the relevant denominator	When using rates or thresholds, define eligible cases, observed events, incomplete follow-up, missing observations, and dependence among contacts or tests. Several contacts in one institution are not automatically independent demand observations.
Review the whole burden	Compare remaining creation, review, support, and realization effort with the current case and authorized resources. More collateral or control documents are not a default response to weak uptake; mandatory safeguards still remain due.

Continuation decision

At the agreed decision point, the owner may continue within the allowance, run a bounded revised test, rescope, hold, or stop discretionary investment. Record evidence, uncertainty, remaining commitments, resource decision, and reopening trigger. A weak or untested outcome case may justify a deliberate limited-risk decision, but must not be described as validated demand or successful realization.

Keep stopping further investment separate from deleting evidence, breaching a commitment, or abandoning required safety or compliance work. Preserve the controlled asset and records where appropriate. A resource hold commissions no further discretionary work unless stated; it does not clear a release restriction or erase mandatory monitoring and retention duties.

Learning without overclaiming

Compare intended with observed outcomes and record what the evidence changed about the hypothesis. Generalize only within the tested scope under Section 4.6. Repeated difficulty across a few related assets may motivate a revised hypothesis; it does not establish a universal market diagnosis, asset value, or causal explanation. Reuse only the portable lesson through M5.

Outcome review does not reopen an accepted artifact merely because adoption is weak. A content change needs a demonstrated defect or newly authorized requirement under Section 3. Artifact quality, commercial value, institutional acceptance, and distribution effectiveness remain separate questions.

6. Verification, Release, and Maintenance

6.1 Unified acceptance regime

Use the following regime once at the appropriate scopes. Local work units establish their own evidence; integration checks the relationships among accepted outputs; release confirms the intended use. Reference valid prior results instead of maintaining duplicate checklists or rerunning unaffected work.

Check domain	Evidence required for acceptance
Authority and input	The action is authorized; governing controls agree; targets/context and baseline qualification are unambiguous.
Change compliance	Every authorized action and no-change boundary is accounted for. Exact/closed-list preconditions, postconditions, and allowed differences pass where applicable.
Substantive sufficiency	The required content, reasoning or behavior exists and serves its declared purpose. For content, record the M6–M6.1 review and the applicable reader task or experience assessment, with its actual coverage and limits. Cosmetic changes do not establish a functional merge, compression, migration or reassignment.
Evidence and state	Claims, derivations, assumptions, facts, dates, and status match their support and authority. Contrary evidence, use ceilings, actual coverage, and remaining validation limits are explicit.
Preservation and audience	Protected content, history, field behavior, and guardrails survive. Narrator, recipient action, disclosure, and required local boundaries fit the artifact's role.
Technical and exact representation	Relevant tests, calculations, file behavior, exact payload, metadata, package surfaces, accessibility checks, and complete render/open-file verification pass.
Integration and dependencies	Identities, terms, source mappings, cross-references, interfaces, handoffs, and coupled changes agree. Dependency reviews and affected evidence dispositions are closed.
Expert / independent review	The named review is complete to its actual scope, using required independence, methods, adjudication, and evidence—not an assumed endorsement.
Action and release	Applicable live-route, permissions, exposure, pilot/scale, package, and signoff conditions are met for the exact intended use.

For quantitative or rule-based gates, define the applicable method, population/sample, metrics, aggregation, denominators, precision, numerical boundaries, missing/rare/unevaluable cases, diagnostics and outcome mapping before the qualifying run. Numerical bands must cover all cases without ambiguous overlaps. Use unrounded values unless the governing method prescribes otherwise; display rounding cannot change the verdict. Do not choose metrics or omit components after seeing results to obtain a pass. Unspecified methods require authority decisions, not editorial completion.

Report ratio numerators and denominators. An empty required population, undefined score or unperformed check is not automatically full coverage or a pass. A not-applicable outcome requires a scope decision compatible with the intended use; otherwise retain the unmet condition. Aggregate performance cannot compensate for separately required failed diagnostics or unit-level safeguards. Distinguish incomplete qualifying evidence from a completed evaluation that fails acceptance.

6.2 Reconciliation and review independence

For material transformation, dependency-bearing work, high-risk outputs, or final assembly, perform a fresh reconciliation after save/reopen/render or the equivalent technical test. Restart from the governing controls and final artifact, not from the first-pass explanation. Confirm dispositions, preservation, dependencies, and unresolved conditions.

A second pass by the same person or model is self-reconciliation, not independent validation. A different prompt, session, model, agent or provider does not by itself establish independence: reviewers can share sources, assumptions and errors. Independent review requires the reviewer and method specified by the project's risk or discipline. Record who reviewed, what was inspected, the standard, result and limits. Do not use an unevidenced conditional status to bypass a required gate. [E05]

Retain evidence in the existing acceptance record. Template A4 in Appendix A captures the decision without reproducing this entire regime.

Use an independent basis for judgment

For a material assessment, the reviewer establishes the applicable criteria and evidence before relying on a model's verdict. Where anchoring on the draft could compromise judgment, inspect the underlying source or form an initial assessment first. Treat agreement, praise, confidence and persuasive explanations as responses to examine, not corroboration. Challenge consequential assumptions with relevant evidence and alternatives; do not require invented objections or reward automatic disagreement. [E06, E16]

Qualify automated review used for acceptance

The review owner defines the actual task, error classes and acceptance rule. Test the proposed evaluator on relevant correct cases and known defects, using an appropriate external reference or expert judgment; inspect false acceptance, false rejection and missed critical cases. Keep evaluated material from supplying instructions to the reviewer. Record the model/mode, prompt and source basis where material, test scope and unresolved limits. If the evaluator cannot meet the required standard, narrow its role to triage or obtain the required review; do not promote model consensus into acceptance.

Self-correction may be useful, including in specifically trained configurations; it is neither universally ineffective nor an independent validation method. Compare a proposed correction with the original requirements and evidence, preserve already-correct content, and repeat affected tests. Stop repeated review cycles when they add no relevant evidence or require a change outside the authorized scope. [E15]

Require observable verification

Use recoverable source comparisons, explicit derivations, executed calculations, test outputs and action records appropriate to the claim. A model's account of its reasoning or alleged tool activity is not an execution log or proof that its explanation faithfully describes the process. Do not require private internal reasoning as a substitute for checkable evidence. Reconcile reported results with the actual saved output and preserve uncertainty where the required check cannot be completed. [E07]

6.3 Acceptance and Release Decisions

State the exact decision object: a work-unit output, downstream baseline, recipient release, deployment, delivery, or scale action. A verdict without its purpose and scope is incomplete.

Decision	Meaning and permissible consequence
Pass	All conditions required for the named use have passed. Proceed only within the recorded audience, environment, scope, and validity boundary.
Conditional pass	Only named, genuinely non-blocking exceptions remain, and the authorized use is valid with those conditions. Record owner, limit, and resolution trigger; do not disguise an unmet prerequisite.
Hold	The work may be sound, but a timing, approval, recipient, currentness, or other required condition for this use is not met. No affected use proceeds.
Block / return	A material authority, evidence, safety, privacy, legal, technical, or dependency failure prevents acceptance. Identify the correction, adjudication, or recovery path.

Keep task dispositions distinct from release verdicts. “Reviewed—no change” is a review outcome, “candidate” is unaccepted work, and “archive” is retention. Acceptance for a drafting dependency is not external release. Artifact acceptance also does not establish uptake or realized value; M10 governs the separate outcome and continuation decision.

Required decision record

Record the object and exact identity; intended use; applicable controls and completed check references; accepting authority and date; verdict; conditions/exceptions; missing prerequisites; dependency impact; receiving owner; validity or next trigger; and the retained evidence location. Use an existing issue or release system rather than another standing register.

A hold or restriction may be a completed decision for the current work scope while the affected use remains prohibited. Record whether authorized remediation is active, with its owner and next step, or no further work is commissioned and only defined reopening triggers apply. Do not invent continuing research duties or call the restriction cleared; mandatory monitoring, expiry and review obligations still apply.

Bounded failure handling

Contain a failure at the smallest safe scope. Stop all dependent uses that consume the affected content, configuration, or evidence. Continue independent work only when the boundary is explicit. Do not silently downgrade a mandatory test, infer missing approval, or reinterpret a failed prerequisite as an optional enhancement.

If correction requires more latitude, a new target, a different method, or a changed substantive requirement, return to the proper authority. Resume only from an eligible baseline with the updated work-unit contract. A successful retry does not erase the historical result or permit an unsupported claim about earlier execution.

Artifact-specific release focus

For documents, verify reader-facing meaning and the final saved render. For data, verify schema, provenance, formulas, and export behavior. For code/models, verify the relevant build, dependency configuration, tests, reproducibility, security, and deployment gate. For media, verify rights, alternatives, dimensions, and export fidelity. For formal instruments, preserve the governing terms, parties, authority, schedules, and event-dependent fields. These are applications of Section 6.1, not separate competing acceptance systems.

6.4 Delivery and Integrity Close

Define the delivery perimeter separately from supporting agreements, acceptance records, internal QA, and historical material. A manifest identifies what is delivered; it does not expand rights, warranties, release permissions, or acceptance standards.

Non-circular delivery sequence

Step	Required evidence and limit
Define	Agree included identities, versions, formats, exclusions, delivery method and duration, continuing service dependencies, required integrity values, and acceptance authority. Do not silently omit a required unit.
Freeze	Finalize the deliverables and record the actual frozen identity. Later edits create a successor and invalidate affected integrity values.
Identify and hash, when required	Compute values from the frozen files using the specified method. State each value's coverage. An ordinary hash record must not claim to include its own changing contents.
Package and expose verification data	Build the declared archive and calculate its integrity value if required. Keep the archive's own hash outside it or use an explicitly defined exclusion scheme. Make manifest data available when recipient verification begins.
Deliver and verify	Record actual release/receipt, comparison method, mismatches, and exceptions. Later-event fields may remain explicitly pending while the already-established file identifiers are supplied.
Accept and retain	Finalize the acceptance record only from observed events and governing terms. Preserve unresolved exceptions and their authorized treatment; retain required source and delivery evidence.

When prior review supports acceptance, reconcile its source revision and representation with the actual delivery. Record and resolve differences under the governing terms, repeating affected checks as required. Equivalent content across formats need not have equal hashes. Neither a sample nor an earlier revision establishes review of uninspected units or changed content.

For independent handoff, verify recipient-controlled copies and test extraction/decryption and agreed open/run behavior in the agreed environment. Provide authorized artifact-specific keys, dependencies and instructions, not personal or non-transferable accounts. An expiring sender-controlled link alone is insufficient. For service access, verify configured access, agreed duration, dependencies and provider responsibilities instead. Acceptance must not assume an unrecorded future sender action or create unagreed support.

The sequence must respect the governing instrument. Verification or receipt does not automatically change payment, transfer, rights, or responsibility. Do not impose a substantive condition through a certificate. Allocating future use to a recipient does not itself cure a delivery defect, waive surviving obligations, or create unagreed continuing support.

Read-only integrity close

Where the control sheet requires a final assembly check, confirm the exact accepted identities, counts, perimeter, required titles, open-file behavior, and fingerprints/checksums. Do not mutate or remediate the accepted files during this close. Use existing verified renders and records; a newly required test or correction belongs in an authorized work unit. A mismatch reopens only the affected artifact and its dependencies.

6.5 Stabilization and maintenance

After acceptance, freeze the controlled identity and hand responsibility to the named receiving owner or maintainer within the agreed scope. Stabilization means no further editing is presently authorized for that identity; it does not establish broader implementation, validation, deployment, release, or scale. Reopen for a confirmed defect, changed authority, new material evidence, failed implementation, scheduled review, or approved successor—not stylistic preference alone.

Maintain live facts only in their canonical homes. On a trigger, update the affected proposition or configuration, assess dependent use and verification under Section 3.3 and M8.1, and restore acceptance before reuse. Keep fixed historical records immutable and retain the evidence needed to distinguish what governed and what actually occurred.

Appendix A. Working Record Templates

Use these compact forms only when their function is active. They may be rows, sections, or views of an existing system. Fill a value once at its canonical home and reference it elsewhere. Bracketed entries are placeholders, not default facts.

A1. Artifact Record

Field meanings are defined in Section 2.2. Use the Operator Control Sheet for project-level decisions rather than duplicating them in every artifact record.

Field	Entry
Identity / location / format	[Stable ID/title; source revision and fixed binding where required; representation/copy-set reference; actual location; required output name.]
Role / layer / representation	[Primary job; layer; governing, derived, template, example, snapshot, or event record.]
Owner / audience / action	[Accountable owner; reader/user and stage; intended task or experience; permitted use/action.]
State / baseline / lineage	[Working-set state; eligible baseline; qualification reference; provenance limits.]
Authority / dependencies	[Controlling records; canonical home; upstream imports; downstream consumers; derivative update rule; package-view reference if active.]
Status / fields / preservation	[Relevant status dimensions; field triggers, cardinalities and due events; protected surfaces/history.]
Evidence / environment / form	[Available substance and gaps in the blueprint/task record; source, claim/derivation and assumption records; currentness; tools/data/rights boundary; M6 language choices, style and accessibility.]
Output / acceptance / maintenance	[Tests and decision references; release scope; retention; receiving owner; next trigger.]

A2. Change Record

Use Sections 3.3–3.6 for the decision logic; add the exact-match fields only when the correction requires them.

Field	Entry
Finding and authority	[Issue/row ID; supported defect or required outcome; source; value-of-change result.]
Scope and baseline	[Eligible identity; target and stable anchor; read-only context; unaffected boundary.]
Posture / pass / latitude	[Action posture; pass type; permitted initiative; closed list yes/no.]
Instruction / allowed difference	[Exact or bounded operation; payload boundaries and decision prerequisites; sequence; permitted surfaces; protected content.]
Exact-match contract, if active	[Old anchor; expected count/locations; locator equivalences; permitted already-satisfied result; final postcondition.]
Structure and dependency	[Disposition/destination; identifier migration; coupled set; affected consumers/checks; authorized scope exclusions.]
Acceptance and handoff	[Test/result; reviewer/authority; no-change or blocker basis; next eligible state; receiving owner.]

Working Record Templates: Run and Acceptance

A3. Work-Unit / Session Record

This is the single session contract and run entry described in Section 4.1. Link approved change rows instead of recopying their instructions.

Field	Entry
Run identity and objective	[Name/ID; decision or output; intended reader/user purpose where relevant; active mode.]
Controls and compatibility	[Approved control identities; action-row references; compatibility disposition.]
Inputs and read order	[Actual required input set; targets/context; eligible baseline; provenance limits; reading sequence.]
Ownership and method	[Executor; canonical writer; reviewer; integration/accepting authority; method/tools; M6–M6.1 for content work; material model, mode, configuration and context bindings under Section 4.7.]
Run mode and dependencies	[Sequential/parallel/conditional; prerequisites; shared frozen context; same-session promotion; coupled set.]
Boundary and outputs	[Latitude; preservation; prohibited actions; output identities/formats; support and partial-delivery rule.]
Exit and evidence	[Definition of done; relevant Section 6.1 checks and results; reader task or experience assessment where applicable; independent review if required.]
Failure and handoff	[Blocked items; return/rollback route; exact accepted state; receiving owner; recipient decisions, uncertainty and expressly deferred work; dependency decision; next trigger.]

A4. Acceptance / Release / Dependency Review Record

Select the decision purpose explicitly. A dependency no-change review need not become a separate file or a new release state.

Field	Entry
Decision object and purpose	[Exact identity; local completion, downstream qualification, dependency review, external use, delivery, or scale.]
Scope and standard	[Materials/surfaces/claims tested; applicable controls and defined evaluation contract; audience/environment; inspection limits.]
Evidence and reviewers	[Completed tests/reviews; reviewer identity and independence; automated-review qualification when used for acceptance; source/check dates as applicable.]
Verdict or review outcome	[Pass, conditional pass, hold, block/return; or reviewed—no change for a dependency review.]
Conditions and remaining basis	[Exception or unmet prerequisite; owner; use limit; active resolution task or settled restriction; reopening/expiry trigger.]
Dependency and mutation impact	[Consumed meaning checked; no action/revalidate/targeted rerun or authorized scope exclusion; affected approvals/counts/hashes/tests.]
Acceptance and handoff	[Authority/date; permitted next use; exact accepted state; receiving owner; review-to-delivery and handoff evidence where required.]
Evidence location and retention	[Existing canonical record; package/manifest reference if applicable; retention and reopening rule.]

Appendix B. Task Prompt Library

These launchers select work; they do not grant authority beyond the approved project controls. Supply the relevant guide sections or an approved control extract, the task record, and required inputs. Replace all bracketed fields before use. Choose one task purpose and state whether editing or external action is authorized.

B1. Intake and Architecture

Classify [project/workstream] using [approved inputs]. Propose a proportionate control level, completion criteria suited to the purpose, authority, audience and dependencies under Sections 1–2. Apply Section 3.1.1: a coherent purpose and tolerable authorized commitment may justify creation without novelty, measurable utility or demand. Permit provisional architecture within [authorized scope]. For content work, connect reader need, contribution, available substance and gaps in the existing record; apply M6–M6.1 and Appendix D1 where relevant. Activate M10 only for objectives or commitments dependent on uptake or post-delivery benefits. Return [authorized intake output], decisions and genuine blockers. Limit drafting and external action to [instruction].

B2. Audit and Change Definition

Audit [target set] against [controls] without editing. State actual inspection coverage, eligible baseline and provenance limits. Use M6.1 to examine substance, organization and preservation of meaning. Distinguish unclear wording, missing explanation, unsupported claims and unnecessary content before proposing repair. Separate mandatory defects, optional improvements, no-change items and unresolved questions. Under Section 3, propose bounded change records with authority, target anchors, unaffected scope, posture, latitude, preservation, dependencies and observable acceptance. Include exact matching and allowed-diff fields only where needed. Return [authorized findings/matrix format]; recommendations remain subject to authorization.

B3. Research and Evidence

Answer [question] within [scope and source plan] using Sections 4.5–4.6, M7, and M9 where relevant; apply Section 4.7 when context or configuration is material. Verify source identity and distinguish retrieval, extraction, interpretation, and validation. Trace material derived claims to premises and rationale; preserve scope, counterevidence, non-equivalence, and prohibited stronger readings. Keep assumptions and decisions distinct from facts; date only what was checked. Return [output] with citations/provenance, coverage, unresolved questions, and bounded implications. Apply Section 6.1 to quantitative acceptance; do not claim unperformed verification.

B4. Authorized Creation and Controlled Revision

Produce [authorized output] under [compatible controls and work-unit record]. For new content, use [approved brief, purpose, sources and creative scope where applicable] and M6–M6.1 to develop supported substance, connected reasoning and suitable expression; identify material gaps without fabricating facts. For revision, execute [approved action rows] on [named targets and stable anchors] using [eligible baseline] and read-only context [set] in [order]. In both modes, observe [latitude, preservation, allowed differences, outputs and stop rules]. For exact edits, verify expected matches and apply only authorized payload; hold unresolved decisions under Section 3.5. Complete local and integration checks under Section 6.1, including meaning after revision and final technical/render verification. Account for every target and hand off under Section 4.4. Return only [authorized deliverables/dispositions] and concrete blockers.

Task Prompt Library: Reconciliation and Release

B5. Pairwise and Protected-Copy Work

Compare [versions] for [future-use purpose] using M1–M2 and the baseline rules in Section 3.2. Preserve historical/issued records unchanged. Map approved imports, intervening accepted changes, unique value, and safeguards to the current target; do not choose by recency or restore obsolete wording. Apply only [authorized differences/latitude]. Return [future-use output or analysis-only disposition], lineage limits, dependency impact, and unresolved conflicts. Do not reuse approval or signature evidence as if it covered changed content.

B6. Gate, Derivative, or Live Route

Prepare [derivative or gate material] for [recipient/stage/decision] from [approved parent]. Released material is [scope]; withheld material is [scope]. Apply M3–M4 and M6 to preserve provenance, permitted use, and evidence ceiling. External action is [not authorized / expressly authorized scope]. When authorized, verify the recipient/endpoint, permission, payload, execution mode, exposure, stop rule, and required pilot/scale gate. Record relationship and observed event scope under M4.1. Return only [approved output and execution evidence].

B7. Reusable Master Refinement

Refine [master] using [permitted lesson sources] under M5. Distinguish prescribed solutions from demonstrated outcomes and compare each candidate with existing coverage. Admit only a portable improvement to a decision, failure mode, preservation duty, handoff or acceptance test. Merge equivalent candidates, remove local defaults and strengthen the canonical treatment first. Test against [relevant project scenarios]; when production guidance changes, examine representative outputs against their intended tasks or experience. Return [authorized standalone master] with clear language, consistent references and no case-specific or revision narrative. Keep analysis records outside the master.

B8. Final Verification and Release

Evaluate [exact artifact/package identities] for [specified use] under Section 6 and [project acceptance standard]. Use valid prior local evidence; verify integration, affected dependencies, claim/status boundaries, audience/action fit, technical behavior and package scope. For content, inspect the M6–M6.1 review and applicable reader task or experience assessment, including preserved meaning and actual coverage. Perform fresh reconciliation as required; distinguish self-review from independent validation. Reconcile prior review to delivery and test the agreed handoff under Section 6.4. Return [pass/conditional pass/hold/block decision and authorized files] with unresolved conditions. Perform a read-only integrity close only when required; route corrections to an authorized work unit.

Appendix C. Glossary

Definitions below fix terminology. Operating procedures remain in the main sections; the glossary does not add separate requirements.

Term	Meaning
Acceptance fact	A result established by the specified test, approval, signature, receipt, or other controlled event; not by document preparation.
Accepted baseline	An identified, non-superseded artifact qualified by the appropriate authority to control a named downstream use.
Action posture	Whether an accepted item is applied, encoded conditionally, watched, left unchanged, blocked, or treated as a completed interlock.
Artifact universe	All controlled deliverables and stateful inputs that matter, including evidence, code/data, media, controls, records, history, exclusions, and missing items.
Authority domain	The limited factual, legal, technical, project, or procedural question a source is entitled to govern.
Baseline provenance	The origin, lineage, comparison basis, qualification event, and known limits of the input used for work or acceptance.
Candidate output	Produced or revised work still undergoing acceptance; not a canonical source or authorized release.
Canon / canonical home	The approved source of a requirement, meaning, or live fact within its assigned domain.
Change envelope	The authorized set of operations, objects/surfaces, collateral changes, and preservation boundaries for a pass.
Closed change list	An exhaustive change envelope in which unlisted changes—including incidental cleanup—are prohibited.
Completed interlock	A bounded verification/correction whose accepted result has been consolidated into the baseline used downstream.
Constrained interface	A destination with explicit character, field, schema, size, dimension, duration, or other technical limits.
Context	Read-only material used for evidence, comparison, alignment, or authority in a named work unit.
Control input	An approved work order, audit, policy, protocol, decision, or upstream result controlling work within its assigned scope.
Control compatibility	Agreement among active controls on shared scope, identities, dependencies, permissions, outputs, and acceptance—not mere matching dates.
Coupled change set	Linked changes that must be mutually consistent before dependent use or group acceptance.
Current-state home	The assigned record for volatile propositions, source/verification state, and refresh triggers.
Definition of done	The objective output, evidence, tests, and required review that establish completion of a material work unit.
Dependency review	A bounded assessment of how an upstream change affects what a dependent consumes; it may conclude that no local edit is required.

Glossary: Execution, Review, and Use

Term	Meaning
Derivative	A representation adapted from an approved parent for a different audience, format, stage, or use.
Event trigger	A defined occurrence, deadline, expiry, or state change that requires a new check or authorized action.
Execution matrix	The scoped action records that translate approved findings into targets, instructions, preservation, dependencies, and acceptance.
Genericity QA	Testing that a reusable master's rules are portable, actionable, proportionate, and free of project-specific defaults.
Guardrail	A substantive limit on unsafe, unsupported, unauthorized, or overbroad claims, use, rights, access, disclosure, or deployment.
Idempotent reapplication	Authorized verification that a complete postcondition is already satisfied without repeating a mutation.
Independent review	Review performed with the reviewer, method, and independence required by the applicable standard; not merely another pass by the same executor.
Latitude	The explicitly permitted degree of executor initiative, distinct from the task's pass type.
Layer	A role-specific level such as canon, evidence, technical implementation, wrapper, internal control, formal/access, or delivery/history.
Match cardinality	The required number of occurrences of an exact correction's anchor within its authorized locations.
Mirror	A bounded local representation of canon; not an independently maintained definition or source of new meaning.
Read-only integrity close	A final identity/perimeter check of accepted outputs that does not mutate or remediate the package.
Revalidation	A bounded check that an existing output or its verification remains adequate after a relevant change.
Reviewed—no change	A recorded review outcome with basis and scope; not a new source-verification date or broader release permission.
Self-reconciliation	A fresh comparison of the final artifact with governing controls by the same executor; distinct from independent validation.
Semantic migration	Transfer between structures or versions according to unchanged, changed-but-mappable, or incompatible meaning.
Stable target anchor	A heading, clause, key/ID, bookmark, schema/code path, or equivalent semantic locator; physical page coordinates are secondary.
Stabilization	Freezing an accepted identity until a defined reopening trigger; not proof of implementation, validation, or external release.
Work-unit contract	The record of inputs, owner, bounded process, output, exit criterion, review, failure path, and handoff.
Wrong-layer residue	Content, fields, or implied authority that belong to a different audience, stage, narrator, or canonical home.

Appendix D. Conditional Blueprint Overlays

Use these choices only when the project type requires them. They supplement Section 3.1 in the existing blueprint and reference canonical controls elsewhere in this guide. They do not require separate files, duplicate records, or a fixed document sequence. An overlay cannot relax a governing requirement.

D1. Knowledge and Publication Assets

For a book, reference work, standard, learning resource or other structured knowledge product, define the reader purpose and proportionate production architecture under Section 3.1.1. Connect intended contribution, available substance and material gaps in the existing blueprint. Use M6–M6.1 during construction and review. A planned argument guides inquiry; evidence determines which conclusions the work can support.

Blueprint choice	What to specify
Identity and promise	Controlled title/category; one-sentence reader promise; current form and intended use; what the work is and is not. Keep planned features distinct from existing content.
Audience and task	Primary and secondary readers; prior knowledge; questions or decisions they need to address; entry points and selective-reading routes. Do not treat reader relevance as evidence of acquisition interest.
Distinctive contribution	The synthesis, pedagogy, accessibility, interpretation, organization, method or other contribution that serves the chosen purpose. Existing coverage is not a veto. Route novelty and comparison claims through M7 and Section 4.6; priority or exclusivity requires evidence.
Argument/content architecture	The cumulative argument or task structure; shared foundations; each unit's contribution and required reasoning; dependencies; examples and material counterexamples; boundaries between stable explanation and refreshable evidence.
Delivery architecture	Volume/file structure; print, digital, interactive, or reference allocation; formats; canonical source and derivative bindings. Reader navigation and technical implementation are distinct from the evidence supporting the content.
Coverage and maturity	Reference Section 3.1.1 for required coverage, saturation, debt, and conditional or exploratory material. State the permitted present scope without implying that the core is already validated or deployable.
Production and review	For each material unit, reference its intended contribution, available substance and gaps, expected evidence, required specialists, size/effort budget, design assets, handoff and acceptance under Section 4.1. Apply M6–M6.1 and keep central claims traceable under Section 4.6.
Reader-facing form	Choose orientation, examples, captions and reference routes that help the reader. Keep material limitations visible beside the affected content. Apply M6 language and terminology choices; keep process codes outside ordinary prose unless needed for meaning or traceability.

The overlay is satisfied by a coherent, executable blueprint—not a manuscript, validated research result, or publication approval. The final work must meet its actual reader promise and permitted claim scope under Section 6.

D2. Controlled Evaluation and Transactions

Activate for protected review, diligence, licensing, acquisition, or another formal transfer. This is a planning overlay, not a legal instrument or standard set of deal terms. The parties' applicable instruments and competent legal review govern. Functions may be combined or separated; no particular sample packet, review order, commercial structure, or no-services posture is universal.

Transaction function	Choices to settle at the applicable stage
Proposition and perimeter	Identify what exists, its review/release status, what is proposed for transfer or use, exclusions, optional future adaptations, and the next decision. Keep price aspirations and buyer-fit hypotheses in the appropriate internal record.
Evaluation authority	Identify the confidentiality and permitted-purpose basis, recipients, controlled work product, onward sharing, permitted systems, return/retention, and termination. No template or blank field grants access; determine whether the instrument itself supplies a sufficient access scope.
Review configuration	Bind the approved materials/version to named or otherwise authorized reviewers, window, repository, source/derivative formats, copying/export, and technology. Apply M3–M4 to optional samples and full review without making either route compulsory.
Diligence facts and rights	Resolve relevant identity, authority, contributor/chain-of-title, prior grants, encumbrances, third-party content, and material development-technology restrictions. Distinguish reported facts from checked records, unknowns, proposed terms, and future-event fields.
Operative decision and obligations	Use the approved instrument for rights, economics, conditions, signatures, remedies, exclusions, and responsibility. Define support as included, separately agreed, or excluded. Explanatory copy and diligence answers do not silently amend those terms.
Delivery and completion	Apply Section 6.4 to the agreed source set, formats, integrity evidence, verification, exceptions, and independent possession or service access. Define any objective completion event and due performance so the sequence is feasible and non-circular.

Provenance and use boundaries

Where technology-assisted development is material, identify the actual human contributions, relevant service or tool, development period, and applicable restriction evidence at the required diligence level. Do not treat tool-output allocation, sole participation, citation, or possession of files as proof of protectability, clear title, or third-party clearance. Disclosure of material rights facts does not make private working logs default deliverables.

Approved storage, private processing, or a non-training setting does not by itself authorize upload, reusable extraction, model evaluation, or product use. The permitted processing purpose and environment must be established under the applicable authority, separately from the technology's settings.

Completion at the required stage

Complete required prerequisites before the action that consumes them; record event-created facts only when observed. Where an agreement requires all sender work completed by a single closing or handoff event, do not leave later confirmation or correction implicit. Other service models may expressly provide continuing obligations. Technical acceptance and later-use responsibility establish only what the governing instrument says.

Appendix E. Research Sources and Use Limits

These notes identify evidence and professional guidance relevant to the cited canonical controls. The controls are operating design requirements, not claims of experimentally proven effectiveness. Historical evaluations retain their tested systems, tasks and limits; documentation is guidance, not an outcome study. Use M7 and M8.1 when a source or configuration changes. Source access: 3 October 2026.

E01 Peters, U., and Chin-Yee, B. Generalization Bias in Large Language Model Summarization of Scientific Research. Author manuscript, arXiv:2504.00025v1, 28 March 2025.

Use: Sections 4.5–4.6. Ten tested models, 4,900 scientific summaries; comparison concerns fidelity to source generalizations, not independent truth of those sources. Predominantly medical material; prompt, interface and model conditions varied. Do not infer current-system prevalence or ban justified generalization.

<https://arxiv.org/abs/2504.00025v1>

E02 Li, R., et al. DeepResearch Bench II: Diagnosing Deep Research Agents via Rubrics from Expert Reports. Preprint, arXiv:2601.08536v3, 10 September 2026.

Use: 4.5; M6.1; 6.1. 132 research tasks; expert-derived criteria distinguish information recall, analysis and presentation. Systems are reported as available in November 2025. Automated judging, possible source-article leakage and subjective annotations limit interpretation; presentation scores are not content assurance.

<https://arxiv.org/abs/2601.08536v3>

E03 Laban, P., et al. LLMs Get Lost In Multi-Turn Conversation. Author manuscript, arXiv:2505.06120v1, 9 May 2025.

Use: 4.7. Six task families with requirements split across simulated turns exposed premature assumptions and unreliable recovery. Simulated conversations are not representative of all human dialogue. Consolidation, resets and shorter units are design options to test, not established universal cures.

<https://arxiv.org/abs/2505.06120v1>

E04 Cemri, M., et al. Why Do Multi-Agent LLM Systems Fail? Author manuscript, arXiv:2503.13657v3, 26 October 2025.

Use: 4.7. Analysis of more than 1,600 traces from seven frameworks, with human-grounded and model-assisted annotation, distinguishes design, coordination and verification failures. Targeted interventions improved some tested systems but did not eliminate failures; additional agents are not assurance.

<https://arxiv.org/abs/2503.13657v3>

Research Sources: Review and Expression

E05 Kim, E. M., et al. Correlated Errors in Large Language Models. Proceedings of Machine Learning Research 267, 30038–30066 (ICML 2025).

Use: 6.2. Multiple-choice leaderboard and offline résumé-evaluation analyses found shared errors across models, including different providers. Standardized and partly subjective tasks constrain generalization. The finding challenges assumed review independence, not the usefulness of every multi-model method.

<https://proceedings.mlr.press/v267/kim25e.html>

E06 Lee, H.-P., et al. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. CHI 2025. DOI: 10.1145/3706598.3713778.

Use: 6.2. Survey of 319 knowledge workers and 936 reported examples. Associations between confidence and reported effort are not evidence of causal skill decline. English-language, self-selected and self-reported observations motivate attention to reviewer judgment, not a universal prohibition on assistance.

https://www.microsoft.com/en-us/research/wp-content/uploads/2025/01/lee_2025_ai_critical_thinking_survey.pdf

E07 Chen, Y., et al. Reasoning Models Don't Always Say What They Think. Author manuscript, arXiv:2505.05410v1, 8 May 2025.

Use: 6.2. Hint-based multiple-choice experiments and synthetic training environments tested reasoning-explanation faithfulness. Monitoring sometimes revealed unintended influence but could not reliably exclude it. Results do not establish the effectiveness of monitoring in every complex or tool-using setting.

<https://arxiv.org/abs/2505.05410v1>

E08 Agarwal, D., Naaman, M., and Vashistha, A. AI Suggestions Homogenize Writing Toward Western Styles and Diminish Cultural Nuances. CHI 2025. DOI: 10.1145/3706598.3713564; author version arXiv:2409.11360v3.

Use: M6–M6.1. Controlled English writing tasks with 118 Indian and US participants and GPT-4o autocomplete found productivity benefits alongside shifts in expression. This setting does not represent all cultures, languages or authors; proposed context-sensitive mitigations were not general efficacy guarantees.

<https://arxiv.org/abs/2409.11360v3>

E09 Doshi, A. R., and Hauser, O. P. Generative AI enhances individual creativity but reduces the collective diversity of novel content. Science Advances 10(28), eadn5290, 12 July 2024. DOI: 10.1126/sciadv.adn5290.

Use: M6.1. Foundational experiment offering GPT-4 story ideas: individual short-story evaluations improved while outputs became more similar in aggregate. Eight-sentence stories and nonprofessional writers limit transfer. Neither creativity loss nor benefit is universal, and a writing score does not define legitimate purpose.

<https://pmc.ncbi.nlm.nih.gov/articles/PMC11244532/>

Research Sources: Actions and Technical Work

E10 DeBenedetti, E., et al. AgentDojo: A Dynamic Environment to Evaluate Prompt Injection Attacks and Defenses for LLM Agents. NeurIPS 2024; author version arXiv:2406.13352v3, 24 November 2024.

Use: M4. Tool-using benchmark with 97 tasks and 629 security cases. Tested defenses reduced some attacks but retained security–utility tradeoffs and task failures. Tool filtering was especially useful where legitimate work needed only read access; that does not solve every write-enabled workflow.

<https://arxiv.org/abs/2406.13352v3>

E11 OpenAI. Safety in building agents. Official technical documentation, accessed 3 October 2026.

Use: M4. Vendor guidance on separating untrusted inputs, constraining data flow and confirming tool actions; expressly not a guarantee of safety. The retrieved page concerns legacy Agent Builder. Only portable boundary principles are used here, not its product lifecycle or a prescribed platform.

<https://developers.openai.com/api/docs/guides/agent-builder-safety>

E12 Spracklen, J., et al. We Have a Package for You! A Comprehensive Analysis of Package Hallucinations by Code Generating LLMs. USENIX Security 2025; author version arXiv:2406.10279v3, 2 March 2025.

Use: M8. Sixteen models generated Python and JavaScript samples; some recommended nonexistent packages. Mitigations were tested in a narrower model subset. Existence checks do not establish trusted provenance, correct behavior or security, and historical frequencies do not describe every current coding system.

<https://arxiv.org/abs/2406.10279v3>

E13 Bertran, M., Fogliato, R., and Wu, Z. S. Many AI analysts, one dataset: Navigating the agentic data science multiverse. PNAS 123(29), e2606495123, 14 July 2026. DOI: 10.1073/pnas.2606495123.

Use: M9. Four models and three fixed datasets produced differing analytic paths and conclusions under varied prompt framing, including after automated compliance screening. AI auditing and the explored analysis space have limits. Variation can expose sensitivity as well as enable selective reporting.

<https://www.pnas.org/doi/10.1073/pnas.2606495123>

Research Sources: Calibration and Continuity

E14 OpenAI. Images and vision: Limitations. Official technical documentation, accessed 3 October 2026.

Use: M8. Vendor documentation identifies limitations involving small or rotated text, graphs, spatial localization, captions, resizing and counts. These are reasons to verify task-critical content; the page is not a measured failure-rate study or evidence that all image systems behave identically.

<https://developers.openai.com/api/docs/guides/images-vision>

E15 Kumar, A., et al. Training Language Models to Self-Correct via Reinforcement Learning. ICLR 2025; author manuscript arXiv:2409.12917v2, 4 October 2024.

Use: 6.2. SCoRe improved self-correction in specifically trained Gemini configurations on mathematics and code benchmarks. This is a counterexample to universal claims that self-correction cannot help; it does not validate prompt-only self-review, arbitrary tasks or independent assurance.

<https://arxiv.org/abs/2409.12917v2>

E16 OpenAI. Expanding on what we missed with sycophancy. First-hand incident account, 2 May 2025.

Use: 6.2. Vendor account of the April 2025 GPT-4o update describes excessive agreement missed by otherwise favorable evaluations and feedback. Its causal explanations are preliminary and source-attributed. It does not establish current behavior or a failure rate across providers.

<https://openai.com/index/expanding-on-sycophancy/>

E17 Anthropic. Effective context engineering for AI agents. Engineering guidance, 29 September 2025.

Use: 4.7. First-hand guidance describes context selection, compaction, persistent notes and delegation, including loss from over-compression. These are implementation strategies requiring task-specific checks, not independently validated universal improvements or authority for replacing canonical source material.

<https://www.anthropic.com/engineering/effective-context-engineering-for-ai-agents>