




4D — Cascade-Semantic Method: Complete Operating Specification

English

<https://prompting-query.com/semantic-cascade-method/protocols-2026/version-4d-full/>

Archive status: current long branch, 80,352 planned for current frozen run

Download the exact Markdown source · SHA-256:
4b13c10edfb929d9e64614ae322941dc931f13357dae4d5c1100f00792f59d3b

Version: 2026-09-27. Author: GPT. Operating language: English, from preparation through final synthesis. Every part of this document, including headings, instructions, explanations, comments and examples, must be in English.

This specification defines the complete operative method and its isolated-answer protocol. The operator's later character-length decision replaces sentence-count requirements in the short profile. Earlier experiments remain unchanged historical evidence. A translated or edited answer is not a new independent model trial.

1. Purpose, authority and language

The primary purpose is to build a sufficiently complete context for understanding and deciding the original problem. New ideas, creative combinations and possible inventions are secondary results. Question count, text length and novelty labels alone do not establish effectiveness.

Retain the user's objective, subject, circumstances and restrictions. Do not invent a country, budget, audience, service, customer or relationship between characters. Preserve a stated count when translating a particular experiment; do not introduce that count into a broader prompt that does not contain it.

Conduct prompt preparation, definitions, semantic and Socratic analyses, focus prompts, tetrads, framework selection, question generation, answer generation, internal review, external-source analysis, registers, final text, requirements and plans in English. Only English quotations, excerpts, explanations and comments may appear in this document or its operative Markdown deliverables. Preserve original source bytes separately in the evidence archive, with stable IDs and English descriptions.

All operative Markdown deliverables must be entirely in English. Show both raw and processed volumes and the retention audit. Never substitute an executive summary for the complete result. Any separately requested language adaptation is an external deliverable and must not be inserted into the 4D specification.

2. Run specification and profiles

Freeze a run manifest before generating answers. Record the original prompt, updated prompt, objects, focus prompts, tetrads, framework items, procedures, modes, variants, lines, language, length profile, evidence snapshots, model/settings if available, execution order and limitations.

Illustrative original experimental prompt:

How can one AI influencer be created and managed on social media to earn profit without deception?

Illustrative creation-focused prompt:

How can one AI influencer be created on social media so that it can be managed and earn profit without deception?

These translate the preserved experimental wording. “Managed” translates the operating activity in this experiment; it does not silently replace a different user-approved prompt containing “developed.”

Default matrix: all selected objects, all selected items of the 62-item framework, four tetrad positions, three procedures, two modes, three variants per family, main and control lines if both are selected. A limited slice must be explicitly recorded. Do not declare the entire matrix complete after finishing a slice.

Current short profile: every individual answer contains 273–304 Unicode characters inclusive, including spaces and punctuation. Count the plain answer after NFC normalization, joining its sentences with one space. Exclude its question, identifier, headings, Markdown wrappers and separators between answers. Store the normalized text that was counted. Sentence count is unrestricted. The same bounds apply to factual and counterfactual answers and to English answers.

Reject out-of-range results and revise the same answer. Never pad with empty phrases, invent facts, cut off a sentence or accept an empty answer. If evidence is insufficient, explain the unknown, the limit of the conclusion and the necessary check within the bounds. This is a local post-generation validation rule, not a claim that a model’s token limit exactly controls characters.

The older long profile remains available only when explicitly selected: 30 complete sentences per answer, with the first variant’s own ten-sentence kernel followed by twenty additional sentences. It is not imposed simultaneously with the current character profile. An unchanged ten-sentence prefix cannot automatically be carried into a 273–304-character answer. Preserve earlier source answers separately and audit meaning retention during shortening.

Execute sequentially, without agents, subagents or concurrent model calls. File handling and deterministic validation are supporting tools, not model agents.

3. Working terms and registries

A semantic object, semanteme or semantic token is a word, expression, action, condition, relation or question selected for analysis. It need not coincide with a language model token. “Without deception” can be one object. Mere presence of a function word does not justify a separate object.

A full prompt preserves the whole task. Refocusing an object changes attention, not the goal or restrictions.

A tetrad contains thesis, antithesis, synthesis and transcendence understood as a metaposition. These are operational definitions; historical attribution is unnecessary for execution.

A framework item is a universal analytical question, not an additional semantic object or a finished expanded question.

A family has the same line, object, framework item, position, procedure and mode; it contains three substantively different variants. The six interrogative dimensions form one compound question and do not multiply the question count by six.

The context register or orchestrator is a record, not an autonomous agent. Maintain goals and conditions; definitions and sources; objects and relations; evidence, assumptions and unknowns; counterexamples and aporias; candidates with provenance; and answer/revision records.

Separate the frozen preparation from the result archive. Preparation contains no previous answers or answer summaries. The archive contains the baseline, older experiments, current answers, revisions, call logs and later synthesis. An answering call cannot inspect that archive.

4. Global stages

Stage 0 — Prompt preparation and lexicographic analysis

Preserve the original prompt verbatim in its source language and save the English working translation. Identify terms whose meanings affect the objective, promised result or conditions. Check dictionary meanings, polysemy and contextual use. Distinguish ordinary, technical and explicitly agreed working meanings.

Record consulted sources, retrieval dates and the precise sense supported. A source not accessed is an incomplete check, not a verified dictionary definition. One unsuccessful search does not prove that an entry does not exist.

Remove redundant wording and clarify vague expressions without losing the objective or restrictions. Do not equate brevity with measured effectiveness. Save the original, translation, definitions, shortened full prompt, changes and retained conditions. Do not finalize the N-object list yet.

Stage 1 — Simple baseline answer

Obtain and preserve an ordinary English answer to the prepared full prompt before the remaining cascade. Record its actual available context. A response written in a chat containing earlier work is not an independent fresh-context baseline.

Keep this answer in the result archive. Do not insert it into frozen answer-generation preparation. A simple baseline is distinct from a detailed matched control line.

Stage 2 — General prompt analysis

Identify question type: descriptive, causal, evaluative, practical, design-oriented or mixed. Identify the user's objective, desired result, subject, scope, constraints and possible success criteria. Separate explicit user requirements, justified contextual inferences and model assumptions.

Preserve unresolved questions. Continue with marked unknowns when they do not block the analysis; do not fabricate missing circumstances. Save this analysis in preparation.

Stage 3 — Semantic decomposition

Identify explicit semantic units and candidate implicit nodes. Explain each in relation to the complete request. Register links such as goal/means, action/result, condition/accessibility, part/whole, cause/effect, temporal sequence and evaluation/basis.

Record dropped distinctions that a short prompt conceals but that matter to the decision. An implicit node is a candidate or test condition, not evidence that the corresponding circumstance exists. Save units, statuses, definitions, relations, missing meanings and questions.

Stage 4 — Socratic workflow

Perform all seven procedures using preparation; ask the user only for genuinely necessary unknown information.

- Clarification: remove verbal ambiguity and emotional coloring; identify the substantive claim.
- Assumptions: identify hidden rules, distinguish observation from common opinion and assumption, and examine possible bias.
- Extreme amplification: construct a strengthened condition up to absurdity or destructive consequences. Mark it as an analytical hypothetical, not a user intention.
- Elenchus: find exceptions, contradictions and counterexamples, with their validity conditions.
- Leaving the echo chamber: examine a principled opponent's explanation, alternatives and possible beneficiaries. Opposition alone does not establish truth.
- Aporias: record identified logical dead ends and search boundaries; do not claim to exhaust all possible dead ends.
- Synthesis or maieutic clarification: build more stable definitions, using generalization or induction where appropriate. Distinguish clarification of the same concept from a change in analytical level.

Save each procedure's results, grounds, improved definitions and unresolved issues. General contradiction search here does not replace the later mandatory self-negation procedure.

Stage 5 — Rebuild the full prompt and finalize N objects

After semantic and Socratic analysis, reformulate the complete prompt to incorporate clarified meanings. Recheck lexical meaning and concision. Preserve prior versions and do not replace the user's objective with a model-selected objective.

Create a flat list: object ID, name and contextual meaning. Preserve inter-object relations in a separate register. For every object construct a full focus prompt putting it in the center while retaining the rest of the request. Preserve an updated unfocused prompt for the control line.

Creation focus: "How can one AI influencer be created on social media so that it can be managed and earn profit without deception?" Profit focus: "How can profit be earned from creating and managing one AI influencer on social media without deception?"

Save the updated prompt, N objects, contextual meanings, N focus prompts and control prompt. Record changes relative to the simple baseline prompt.

Stage 6 — Construct contextual tetrads

For each selected object define:

Position | Required meaning

Thesis | The object defined in the complete request

Antithesis | A justified opposing position or process, with the basis of opposition

Synthesis | A position reconciling or transforming the opposition into a more stable solution

Transcendence | A metaposition reconsidering the opposition's framework in relation to the objective

Use synonyms for clarification and antonyms to investigate opposition, not to impose a unique mandatory antonym. If no natural opposite exists, justify a conditional opposition or record an unresolved question. Mechanical negation is not automatically an informative antithesis. Synthesis is not a verbal average; transcendence must specify what changes in the frame and why it matters.

Example: creation loss or cessation of a specified ability recreation with testable continuity lifecycle management. This is contextual, not universal.

Freeze a tetrad for a particular pass, not forever. If an error is found, preserve the previous version, correct it and identify affected questions. Do not mix versions under one run ID.

Stage 7 — Select the analytical framework

Use the 62 items in Appendix A with stable IDs. They cover ontology, epistemology, space, time, causality, internal/external interaction, purposes, resource-limited action and values. This is a working framework, not a claim to reproduce a philosopher's canonical list.

"Not applicable" is valid if explained. Do not force physical shape onto an abstract process or invent a location or date. Under a fixed questionnaire, retain the question and explain non-applicability within the response bounds.

Stage 8 — Expand questions and freeze the complete questionnaire

For every object/framework pair generate the base treatment of all four positions, then separate self-negation and absence checks for each position. Use factual and counterfactual modes and three substantively different variants per family.

Retain the interrogative dimensions how, what, why, when, where and for what purpose. Phrase the compound question in grammatical English. They represent method, subject, grounds, time, location/domain and goal; they do not require six repetitive answer blocks. Where a dimension is inapplicable, do not invent information.

Factual mode: use stated definitions and conditions; distinguish evidence from unknowns. Without empirical evidence, describe logical analysis rather than an observed project outcome.

Counterfactual mode: explicitly change an identified condition, ordering, state or necessary element. State what changes and what remains fixed. Keep conditional conclusions marked.

Self-negation: check whether a definition, rule or required action defeats its own goal. Identify the subject, objective and conflict conditions. Distinguish strict contradiction, circular conditions, substitution of the subject, excessive restrictions and unverified assumptions. "No conflict established" is valid.

Absence: distinguish absent object/ability, absent evidence, absent testing, subsequent loss and unknown state. Check removal of a necessary condition and relate the conclusion to the full prompt. Do not automatically call this an apophatic method or equate absence of records with

absence of facts.

Review variant differences before answering. Save all planned IDs, the run manifest, preparation, source snapshots, full prompts, tetrads and questions. Compute SHA-256 hashes of preparation and questionnaire. Hashes detect changes, not semantic errors.

Stage 9 — Isolated answer generation and validation

Each call receives only English execution instructions, its corresponding full prompt, the same frozen preparation and one question with metadata. Do not send other questions, answers, paraphrases, baseline responses, historical example answers, literary texts, scores or chat history. Do not let retrieval expose the result archive.

Start a fresh input packet for every question. Same-family variants and control/main lines receive no neighboring answers. In the current short profile there is no mandatory ten-sentence prefix. A correction may see its own answer and a specific local defect, but not another answer. Cross-answer findings belong to later review.

Validate English language, the inclusive 273–304-character range, relevance, mode, conditions, goal retention, non-invention and substantive content. A character validator does not establish semantic correctness. Reject and record failed attempts; do not accept a generic summary instead of a missing answer.

For an explicitly selected long profile, first variants retain their own 10+20 kernel rule. Bind the kernel to run ID, question ID, line, preparation version and question hash. Other variants cannot inherit it. A kernel extension may receive only that answer's own kernel. Preserve revisions if the kernel is defective.

Save each input packet, hashes, original answer, corrections, actual model/settings and available token/time/error data. Do not invent unavailable usage metrics. If preparation has a substantive error, preserve the incomplete pass and issue a new version with a stated repeat scope. Resume from missing or failed IDs, without regenerating completed answers merely to restore context.

Stage 10 — Raw corpus, forensic review and full connected text

Once all selected IDs have answers and reviews, save three English representations: questions with answers, questions only, answers only. Preserve the raw answer corpus unchanged. Logical revision is a separate version, with reasons and ID lineage.

Internal forensic review: examine definitions, composition, sequence, logical links, feasibility, provenance, repetition, substitutions and unknowns. Rephrase an incoherent passage with a reason or exclude it from the processed text with a recorded reason. Return one or two preparation stages when needed, preserving version boundaries.

External forensic review: compare claims and ideas with accessible sources and analogues. Prefer dictionaries, primary research, official documentation, existing solution descriptions and patent sources when an invention is claimed. Save search queries, direct sources, access dates, similarities, differences and search limits. An unavailable search is not completed search.

Distinguish known solution, clarification, creative combination, testable hypothesis and invention candidate. New to this analysis or unexpected to the model is not established worldwide novelty.

Failure to find an analogue does not establish that none exists.

Build a complete connected English text, not only an executive summary. Retain substantive conditions, exceptions, costs, responsibility, tests, demand, distinctions and limits. Link each paragraph to its answer IDs. New externally supported facts require their own provenance; new hypotheses remain labeled. Preserve a concise summary separately if useful; never substitute it for the full processed corpus.

Mandatory retention audit: record raw and final characters, words and locally measured tokens under the same conventions. Show final/raw volume ratios, but do not call them effectiveness by themselves. Map source fragments or semantic units to retained, rephrased, merged, corrected, excluded or pending states. A merged unit must point to an actual destination; shared destination is acceptable. Preserve contradictory views with status until resolved. Record reasons for exclusions. An unmapped source fragment is a loss to review, not an invisible editorial deletion.

Do not impose an arbitrary minimum final/raw ratio as a quality guarantee. A large ratio can reflect repetition and a small ratio can reflect lost information or good deduplication. Assess retention and usefulness alongside volume. Deliver the raw corpus and the complete processed English text together with their volume and retention measurements.

Stage 11 — Compare results, requirements and plans

Compare connected texts on their own, then inspect questions/answers to check provenance and losses. Compare with the simple baseline and a genuinely executed detailed control line if available. Evaluate new distinctions, conditions, dependencies, logic, verifiability, decision usefulness, repetition and volume.

Use Appendix B with evidence fragments and limitations. Do not transfer historical scores to a new translation, shortened answer or revised text. Expert differences are not statistical proof. Record unequal inputs, lengths and prior-context access; separate mechanism changes from content changes.

Where the selected task calls for implementation, derive an English requirements register, acceptance criteria, general action plan, detailed plans, ordered procedures and technical specifications. Map every requirement back to evidence or an explicitly labeled decision. Unresolved choices remain unresolved; do not fabricate a platform or service selection.

5. Exact questionnaire arithmetic

The base pass considers each tetrad position in its own role. Self-negation and absence are additional procedures over each position. This is not a Cartesian product of all six philosophical operations with all four positions.

Procedure per object/framework pair and line | Families | Questions with 3 variants

Base: 4 positions × 2 modes | 8 | 24

Self-negation: 4 positions × 2 modes | 8 | 24

Absence: 4 positions × 2 modes | 8 | 24

Total | 24 | 72

Let N = selected objects, K = framework items, L = lines, V = variants. $Q = N \times K \times L \times 4 \times 3 \times 2 \times V$. At $V = 3$, $Q = N \times K \times L \times 72$. Do not multiply again by four positions or six interrogative words.

Examples: one object/item/line = 72 questions; two lines = 144. Six objects $\times$ 62 items $\times$ one line = 26,784; two lines = 53,568. Under the current character profile, total answer characters range from $Q \times 273$ to $Q \times 304$, excluding inter-answer separators. These are planned volumes, not completed outputs.

The current translated 0.3D sample is the preserved 24-question base slice, one object, item 24, main line. It does not contain the extra 48 questions or a newly executed control line. Translation does not change its denominator or make it a full pass.

New candidate objects enter a separately declared later cascade; do not enlarge the current denominator silently. A full product of four positions, six operations, two modes and three variants would be 144 questions per line, a different method not authorized by this specification.

6. Control, IDs, persistence and continuation

A simple baseline and detailed control line are distinct. Detailed control has the same framework, modes, variants, procedures and length profile but the unfocused full prompt. Do not copy main-line answers into control. Preserve matched question subjects; record focus-related wording changes, or use identical questions if requested.

Stable order: line object framework item procedure mode position variant. Control precedes main by default. Retain an explicitly declared inherited order for a fixed questionnaire and preserve ID correspondence.

Per question record: stable ID; family and variant; line; object and contextual definition; verbatim full prompt; framework ID/text; tetrad position/definition; procedure; mode; changed condition; answer and optional own kernel; evidence/assumptions/unknowns; validation status; lineage and hashes.

Statuses distinguish planned, question constructed, answer written, length validated, content reviewed, included in processing and pending/excluded. Report actual questions, answers and reviewed answers separately.

Save Markdown after each stage, using today's date, author and incrementing version. Preserve earlier files; append amendments to an existing topic's Markdown successor when project rules require it. A clean translation in a new language is a separate named artifact. Never claim saved files when the environment cannot write them; deliver bounded named Markdown parts instead.

A restore record contains full prompt, manifest, object focus prompts, tetrads, framework, preparation/question hashes, last completed and missing IDs, real paths, revisions and next action. Continue from saved results without feeding their summaries into still-isolated answer calls. Do not promise unattended execution without an actual active mechanism.

7. New contextual nodes

As a new term, relation or question appears, relate it to the complete prompt and record its source. Record candidate name, question, originating ID, goal relevance, dependencies, status, recurrence

and testability. Distinguish relevant candidates, weak links and inapplicable combinations with reasons.

Frequency is an attention signal, not proof of importance; model wording can manufacture recurrence. During an isolated pass, store candidates in the archive and admit them into shared preparation only for a later declared cascade. The aim is context construction, not endless expansion or a novelty count.

8. Completion and delivery

A selected run is complete only when its planned questions have reviewed answers, raw representations are saved, revisions are traceable, connected texts exist, the stated forensic scope is performed or explicitly marked unavailable, comparison uses real results and source-to-final retention is audited.

Deliver the manifest and counters; prompts, objects and focus versions; tetrads and framework; three raw representations; full connected text per executed line; review and retention registers; candidate/idea register; comparisons with evidence; and implementation documents if in scope. State incomplete portions and the next action. Visiting all stage names is not completion of every object/framework combination.

Appendix A — Complete 62-item framework

Ontology

- Does the object exist?
- What are its mode and form of existence: material, ideal, virtual or otherwise?
- What is the object in its essence?
- Which attributes are essential and which properties are accidental?
- To which fundamental category of being does it belong?
- What constituent elements or parts make it up?
- What structure of relations among elements makes it a whole?
- Can it exist independently, or is its existence conditional on something else?

Epistemology

- Is the object knowable in principle?
- What are the sources of knowledge about it?
- What methods and tools can investigate it?
- Where are the limits of potential and actual knowledge about it?
- Does the observing subject affect the object during investigation?
- What are the truth criteria for knowledge obtained about it?

- What is the status of current knowledge: fact, hypothesis or belief?
- How do empirical and theoretical knowledge relate in its investigation?

Space

- Where is the object located?
- What are its spatial dimensions: shape, volume and size?
- What are its physical or conceptual boundaries?
- What is its topology, or the spatial arrangement of its parts?
- How is it oriented relative to other objects?
- What are its density and distribution in its occupied space?
- Can it move through space?

Time

- At what moment or during what period does the object exist?
- What is the total duration of its existence?
- What phases or stages constitute its lifecycle?
- Is its existence continuous or discrete?
- How does it transform, evolve or degrade over time?
- What are the rates and rhythms of its internal changes?
- Is change directed or irreversible, or is it cyclical?

Causality

- What is the initial cause of its emergence?
- What conditions are necessary and sufficient for its existence or functioning?
- What effects or results arise from its existence or activity?
- Are the links to its effects deterministic or probabilistic?
- Does it belong to a linear, branching or feedback-based causal chain?
- Does it act as a cause, an effect or both?

Internal and external interaction

- How do its internal elements interact?
- What forces or principles maintain its internal cohesion?
- Are there internal contradictions, conflicts or competing elements?
- What hierarchy or subordination exists among its internal elements?

- Into what larger system or environment is it integrated?
- How does it affect its environment?
- How does the environment affect it?
- What are its inputs, such as external resources and stimuli, and outputs, such as responses and products?
- How open or closed is it to the external environment?

Teleology

- What is the purpose of its existence or functioning?
- Is that purpose inherent, or externally assigned by a creator or user?
- What function does it perform in a larger system?
- What would count as achieving its purpose?
- Does it tend toward a final state, or entelechy?

Praxeology

- How can one interact with it in practice?
- What actions create, maintain or destroy it?
- How can it be used as a means to an external goal?
- What is the most rational and effective procedure for using it?
- What energy, information, time and other resources does its operation require?
- What risks, costs and side effects arise from its practical use?

Axiology

- What are its absolute and relative values?
- In which ethical, aesthetic, utilitarian, spiritual or other value system is it assessed?
- Does it carry good or harm, positive or negative potential?
- What hierarchy of values does it reflect or support?
- What is its aesthetic dimension: harmony, beauty or ugliness?
- Is its value inherent to the object or dependent on the evaluating subject?

Appendix B — 100-point comparison rubric

Score each criterion from 0 to 100 in increments of five, with equal 10% weight. The total is the arithmetic mean. Require cited text fragments and reasons. Identical texts receive identical scores; decimal averages do not imply measurement precision.

Criteria: goal alignment; coverage; precision of distinctions; logical consistency; verifiability; practical usefulness; productive connections; clarity; information density; standalone intelligibility.

Assess the connected text first, then verify provenance and retention against the raw corpus. Productive connections are not worldwide novelty. Scores remain expert judgments unless independent validation and an experimental design support stronger claims. Report raw/final volumes, excluded units, unresolved contradictions and useful retained distinctions separately rather than hiding them in one total.

Appendix C — Legacy long-profile kernel example, not answer-call preparation

This appendix translates a historical illustration, not an observed operating influencer. It is excluded from isolated answer packets. Its continuation rule applies only to an explicitly selected long profile.

Own ten-sentence kernel: Creation is recognized for a specified version and next action. First, record the required readiness properties. Then obtain material that lets those properties be checked. Test continued operation separately from the first sample. Link the test result to specific material. Separate the event time from the recording time. An unknown date remains unknown or receives a justified interval. An account denotes a channel rather than proof of creation. Influence and profit require their own evidence. This procedure supports a next action without promising success.

Twenty-sentence addition: The intention date belongs to preparation rather than readiness. The first-material date is supported by an accessible instance. The readiness-recognition date belongs to the completed check. Record the version, criterion, result and basis. Distinguish the production location from the record's storage location. For distributed work, preserve the provenance of necessary parts. Check the whole after combining its parts consistently. Absolute immutability is unnecessary for further development. Such a requirement would defeat continued management. Describe permitted changes alongside stable properties. Without material, a description alone cannot confirm readiness. Without a record, past material existence cannot be denied. Name a missing test separately from a discovered defect. For a paid promise, test the promised ability itself. Retain preparation costs before economic use starts. Receiving money does not redefine creation. Profit belongs to its accounting period and cost boundary. Version a changed criterion while preserving its earlier assessment. Public statements must match the confirmed state. The practical result of the check is a justified next step.

This historical example predates the explicit distinction between a static product and promised continued production. Current execution applies that distinction and tests continuation only when promised; the example is preserved for provenance, not copied as a current answer.

Appendix D — Isolation implementation and honest execution status

A fresh OpenAI Responses call must not inherit other results through `previous_response_id`, a shared conversation or manually copied history. `store: false` alone does not establish isolation.

Other providers need an equivalent input boundary. A chat with previous answers cannot be called technically isolated merely by instruction.

The accompanying English offline packet builder assembles frozen one-question inputs, validates language and character bounds, and prevents a complete-collection assertion with missing IDs. It makes no network calls. A translated 24-answer sample is not a new isolated generation and does not prove execution of the 53,568-question plan. Actual model execution, billable reasoning and external novelty search remain separate tasks.

Execution command in plain language: prepare and freeze the selected English run; construct the entire questionnaire; answer each question sequentially from its full prompt and the same preparation, without other answers; validate and preserve every result; then perform review and build a full connected English text with a retention audit. Deliver all operative materials in English and report raw and processed volumes with their retention audit.

Superseding amendment — 2026-09-28: soft character target, no length retries

This amendment overrides every earlier instruction in this document that rejects, rewrites or regenerates an otherwise substantive answer solely because its character count lies outside 273–304. The range remains a first-pass target, not a hard admission gate. Ask the model to aim for it without padding, truncating or sacrificing meaning.

After one substantive response, count the normalized answer and record its actual character length. Keep it verbatim in the raw corpus whether it is below 273, within 273–304 or above 304. Label its length status as below target, within target or above target. Do not request another model call, feed back the draft, rewrite it or discard it for length alone. Report the actual length distribution, deviation rate and token use alongside quality. Later substantive edits belong to a separately identified reviewed version; the raw answer remains unchanged.

Retry only a technical failure or a missing response. Do not revive the old length-repair path under another validation label. Historical experiments and their actual retry costs remain unchanged. In Stage 9, “length validation” now means measurement and status annotation; a complete answer outside target still counts as written and included. All other meaning, relevance, language and evidence review remains, but it must not silently overwrite the raw response.

Superseding amendment — 2026-09-28: the final work is literary prose

Status: operator-approved amendment, 2026-09-28. This rule applies to the English 4D pipeline and supersedes any interpretation of “connected text” that permits a stitched list, a technical digest, or an executive summary to stand in for the final work. It does not alter the raw answers or the historical 3D/4D experiment.

Required final artifact

The final artifact of a completed run is one coherent literary work in accessible, polished English. It answers the user’s full prompt directly, within the evidence and scope of the completed run. The

work has a clear beginning, a developed middle, and an ending that resolves the line of explanation without claiming unsupported success. Literary form means lucid, connected nonfiction prose with an intelligible progression, not invented events, fictional customers, ornamental vagueness, or a disguised questionnaire.

For a deliberately limited slice, the work answers the part of the full prompt covered by that slice and explicitly marks the remaining parts as unexamined. It must not claim to be a complete answer to the whole prompt until the full planned corpus and required checks have been processed.

Composition principles

- Read the original full prompt, frozen preparation, and every raw answer in the selected run. Keep the raw corpus immutable. Distinguish user requirements, established observations, conditional cases, hypotheses, and unknowns.
- Identify the substantive unit in each answer: claims, conditions, exceptions, distinctions, causal links, procedures, tests, actors, responsibilities, costs, dates, and evidentiary limits. The number of source answers is not itself the number of ideas.
- Design an order that serves the prompt and the reader. Integrate related units into a continuous argument or narrative progression. Use transitions that explain why one issue leads to the next. Do not preserve the questionnaire's order merely because it is numbered.
- Write at a high but accessible literary level. Prefer plain, precise sentences and concrete relationships. Expand a technical acronym or term in parentheses at first use when a general reader needs it. Avoid unexplained jargon, formulaic rubric repetition, and decorative language that hides an unresolved point.
- Preserve meaningful repetition when it plays a different role in a different case; remove only repetition that adds no condition, consequence, or clarification. Develop an important distinction rather than naming it once and moving on. Expansion is permitted when it makes the reasoning and its practical implications clearer.
- Do not invent a person, customer, sale, platform, right, result, or source to create a pleasing story. Do not turn a possibility into an observed fact. Keep factual and counterfactual cases distinguishable in the prose, and state the limits of evidence where they matter.
- Close by answering the full prompt as far as the evidence permits, naming any decisive unknowns. The ending must follow from the work rather than introduce an unsupported conclusion.

Separate retention check

Produce a source-to-destination register outside the literary work. Map each substantive unit from each answer to a paragraph or passage and mark it as retained, clarified, combined, corrected, explicitly excluded, or pending, with a reason where applicable. A single quotation per answer is insufficient evidence that all of its substantive units survived. Check the completed work for unsupported additions as well as losses. Preserve the raw corpus, the literary work, and the register as separate artifacts.

Measure raw and final characters, words, and locally measurable tokens under the same conventions, and publish their ratio. The 24-answer Russian pilot was approximately 28% longer than its earlier technical synthesis; this is an accepted example of useful expansion, not a required growth rate, cap, or quality score. Judge the final work by completeness, fidelity, clarity, logical development, and usefulness to the original prompt.

Synthesis instruction

Using the full user prompt, the frozen preparation, and every raw answer in the selected run, write one coherent literary nonfiction work in clear, polished English for a general reader. Give it a beginning, a developed middle, and an ending. Make the work answer the prompt directly to the extent supported by the completed evidence; identify the unexamined part if this is only a slice. Integrate the answers into an intelligible progression rather than reproducing the questionnaire or writing a summary. Preserve every substantive condition, distinction, exception, test, responsibility, cost, time relation, and uncertainty. Explain technical terms and acronyms in parentheses on first use where needed. Expand reasoning when this helps the reader; remove only empty repetition. Do not invent people, sales, results, rights, platforms, or sources. Keep hypothetical cases visibly hypothetical. Separately from the literary work, provide a unit-level source-to-passage retention register and mark all unsupported additions or exclusions. Do not treat a length target as evidence of quality.