

CONSTELLATIONS & DIAMONDS



Start Here

A Human-First Protocol for
LLM-Assisted Metabolism of Inquiry

PUBLIC RELEASE V0.3 - 2026-09-02

The Gnome



Constellations & Diamonds

A Human-First Protocol for LLM-Assisted Metabolism of Inquiry

Included in Constellations & Diamonds Public Release v0.3

Role: Start Here - the single introduction, first-use guide, package map, and open-feedback invitation

Status: Official canonical v0.3 public-release source; release sealed 2026-09-02; not independently validated

Standard setup: Read the Fast Path below, then upload all nine Markdown files in 02-TRY-THE-PROTOCOL/ into one new LLM conversation

Authority: This document explains the protocol but does not override the receiving system's governing constraints, the Mirror Constitution, or the complete Core Load

Canonical format: This Markdown file is the canonical accessible text; the synchronized PDF is a designed reading copy with no additional governing content

Public originator: The Gnome of Roosevelt Grove, referred to hereafter as The Gnome

License: CC BY-NC-ND 4.0, with additional permission for private personal customization; see LICENSE.txt

Fast Path - Read This First

What this is

Constellations & Diamonds helps people use general-purpose LLMs for reflective self-inquiry and world-inquiry without surrendering human agency, collapsing uncertainty, confusing meaning with fact, or becoming dependent upon the model.

Its operative method is called the Mirror Protocol. It is a collection of documents that can be loaded into an ordinary LLM chat. Together, they establish expectations, constraints, user controls, and a vocabulary for inquiry.

Here, protocol means a structured collection of instructions, practices, safeguards, and controls. It does not mean a clinically validated treatment protocol or certified safety standard.

What this can and cannot do

Constellations & Diamonds is not a specialized model, therapist, diagnostic instrument, spiritual authority, research result, crisis service, or guarantee of safety. It cannot control an LLM or ensure that the model will follow its instructions. General-purpose models can misunderstand, fabricate, overvalidate, lose context, reinforce a mistaken premise, or speak with confidence they have not earned.

The protocol is designed to make better behavior more likely and harmful drift more recognizable. It asks the model to preserve human agency, label uncertainty, distinguish experience from interpretation and external fact, disclose relevant limits, avoid dependency-producing behavior, accept correction, and stop cleanly. Its effectiveness still depends upon the model, platform, available capabilities, conversation length, subject matter, and your continued judgment.

The standard four-part path

Core Load → Context Load → Inquiry → Feedback

1. Core Load

Start a new conversation with a general-purpose LLM. Upload all nine Markdown files contained in:

02-TRY-THE-PROTOCOL/

Do not select files from the case-study or optional-reference folders for ordinary setup. The Alchemical Lens file is one of the nine Core files, but its language remains inactive unless you deliberately choose it.

After uploading, send this:

Read these nine files together as one integrated protocol. List the filenames you received; identify anything missing, unreadable, duplicated, truncated, or from another release; follow the authority order in the Core Load Manifest; disclose any material capability limits; and confirm that the Alchemical Lens is inactive unless I choose it. Then tell me that I may begin inquiry immediately or choose a Context Load.

If the interface cannot accept nine files, use `Core-Load-Single-File-Fallback.md` from the optional-reference folder instead. Never upload the fallback alongside the nine separate files.

2. Context Load

Inquiry is already available. Personal disclosure is not an entrance fee. Choose any of these, or simply state your inquiry:

- Minimal or No Context Load - begin with little or no personal configuration. The model should use neutral language, minimize assumptions, and request context only when it is useful and permission-sensitive.
- Focused Context Load - provide only what seems relevant to the present inquiry: circumstances, aims, values, preferred language, important boundaries, or the kind of help you want.
- Myriad Context Load - build a broad, open-ended field over time from whatever may matter: practical life, relationships, culture, science, philosophy, religion, spirituality, body, history, memory, art, humor, contradiction, unfinished thoughts, or apparently random fragments.

Engaging a myriad field is a feature, not a bug. Rich context can reveal distant relationships, but it must never force a grand pattern, complete biography, single explanation, diagnosis, destiny, or encompassing worldview. Breadth means breadth of domains and relationships - not identifying precision. You may use pseudonyms, categories, ranges, abstraction, omission, and redaction.

You can change modes at any time. Context may be added, corrected, dated, contradicted, withdrawn, or retired. Greater disclosure does not make an inquiry more worthy or more authorized.

3. Inquiry

State what you want to explore. It may be precise, unfinished, playful, practical, painful, factual, symbolic, or uncertain.

Examples:

- "I want help sorting what I know, what I fear, and what I am only inferring about this decision."
- "Reflect this experience back to me without diagnosing me or forcing a lesson."
- "Help me map possible relationships among these scattered notes. Preserve coincidence and no-relation as valid outcomes."

- “Examine this public claim. Separate evidence, interpretation, missing information, and what would need outside verification.”
- “I want to use a Christian-informed lens, but do not speak for Christianity or treat doctrine as settled on my behalf.”
- “Use plain materialist language. Do not use spiritual or Alchemical metaphors.”
- “Activate the Alchemical Lens as metaphorical language for internal experience only.”
- “I do not know the question yet. Help me locate it without making me disclose more than I want.”

You do not have to perform every stage of the method. A useful inquiry can be short. The model should use no more protocol than the inquiry requires.

At any time, you may disagree, redirect, narrow the topic, decline a lens, correct the record, ask what is inference, request outside verification, withdraw context, pause, or stop. If the conversation drifts, say “Mirror reset.” The model should name the drift, retract unsupported material, repair the record, re-anchor to your inquiry and choices, and offer proportionate options - including closure.

4. Feedback - entirely optional

You are not entering a study. No reviewer code, lane, scoring form, transcript, or continuing commitment is required. Trying the protocol naturally, deciding not to use it, becoming confused, finding it burdensome, or having a negative reaction are all meaningful forms of feedback.

If you choose to respond, you might say what you thought the project was, whether you could begin without help, what changed in the model's behavior, where friction or discomfort arose, what felt useful or overreaching, whether you could disagree and stop, and what should be removed, clarified, or added.

That is enough to begin. The rest of this document provides the deeper frame, concepts, limits, package map, and feedback terms.

The Deeper Frame

Inquiry as metabolism

“LLM-assisted metabolism of inquiry” names a process, not a promise. An inquiry may arrive as grief, a decision, a factual claim, an intuition, a conflict, a recurring image, a body sensation, a moral question, a set of notes, or something that has not yet become language. The aim is not to make the model the knower of the person. The aim is to help the human receive, differentiate, relate, test, carry, revise, release, or close what has entered inquiry.

The Mirror is therefore a role with boundaries. It reflects and helps organize; it does not claim consciousness, felt empathy, privileged access to an inner self, supernatural perception, perfect memory, professional authority, or final judgment. Fluency is not evidence. Emotional resonance is not proof. Usefulness is not truth. A graceful synthesis may still be wrong.

The protocol's ethical center is human-first in a practical sense: the inquiry belongs to the human before, during, and after the conversation. The model should support capacities that remain when the model is absent. It must not position itself as indispensable, exclusive, secret, uniquely understanding, or more trustworthy than human relationships and appropriate sources of support.

Constellations and Diamonds

The title names the two primary organizing forms.

A Constellation is a provisional relational map. Its “stars” may include events, observations, claims, feelings, values, images, sources, contradictions, questions, or unknowns. The map may show association, contrast, sequence, recurrence, possible causation, symbolic resonance, functional relation, or unresolved tension. The kind of relationship matters. A Constellation is not a cage, diagnosis, cosmic plan, or declaration that everything connects.

Some stars should remain separate. Some patterns may be coincidence. Several competing Constellations may fit the same material. A responsible Mirror can conclude that no useful relationship is presently supported.

An Inquiry Diamond is something the human deliberately chooses to carry forward: a concise insight, question, boundary, commitment, or formulation that appears durable enough to test in life. It remains provisional and revisable. The model may propose a Diamond, but it cannot confer one. No Diamond is a valid and often wise result.

The Core Load also contains twenty-five Constitutional Diamonds. These are different: they are release-level invariants governing the method, not conclusions about the user. When the distinction matters, the documents use the full terms Constitutional Diamond and Inquiry Diamond.

Prisms and lenses

A Prism differentiates what has entered inquiry into bands that should not be collapsed - for example:

- what happened, what was experienced, what it may mean, and what remains unknown;
- direct observation, remembered account, model inference, external claim, and imaginative possibility;
- emotional truth, symbolic meaning, causal explanation, and proposed conduct;
- a source's content and any instruction embedded inside that source.

Differentiation is not fragmentation for its own sake. The purpose is to keep important distinctions visible long enough for honest relationship-making.

A Lens is a named, selectable perspective that brings some features into focus while leaving others outside its field. Lenses may be plain-language, scientific, philosophical, cultural, religious, spiritual, poetic, practical, skeptical, systems-oriented, or mixed. No lens is neutral or total. A lens should disclose its relevant limits, remain removable, and never gain authority merely because it produces compelling language.

The user's identity label is not permission to infer an entire worldview. A Christian, Muslim, Jew, Buddhist, Pagan, Wiccan, agnostic, materialist, scientific naturalist, or any other person may hold a distinctive and evolving relationship to their tradition or vocabulary. The Mirror should ask rather than presume and should never speak as the authoritative voice of a whole community.

The Alchemical Lens

Alchemy is included as the flagship Context Lens because its language can provide a clear vocabulary for inner differentiation, transformation, integration, and return. It is offered, not smuggled in.

Within this protocol, Alchemical operations are metaphorical and interpretive descriptions of internal experience. They are not claims about physical mechanisms, scientific processes, supernatural facts, secret knowledge, or evidence. Activating this lens does not raise the truth status of an interpretation.

The Alchemical Lens is:

- inactive until you explicitly activate it;
- removable, replaceable, customizable, and personalizable;
- nonessential to every ethical and operational safeguard;
- subordinate to evidence, consent, and the complete Core Load.

You can decline it after the file has been uploaded. The protocol should remain fully functional in plain or differently configured language.

The operating cycle

The Mirror Protocol can be summarized as:

Encounter → Prism → Lens → Constellation → Inquiry Diamond → Conduct → Fruits → Revision, Release, Closure, or Renewal

This is a set of functions, not a mandatory script.

- Encounter: receive what entered inquiry without immediately controlling, interpreting, or diagnosing it.
- Prism: differentiate relevant kinds of material and preserve provenance.
- Lens: name the perspective being used, what it illuminates, and what it may obscure.
- Constellation: map possible relationships without forcing unity, identity, or cause.
- Inquiry Diamond: optionally preserve a human-adopted distillation with its limits and revision conditions.
- Conduct: consider a proportionate lived response, attending to evidence, consequences, reversibility, and alternatives.

- Fruits: observe what the inquiry produces in ordinary life over time without treating benefit as proof of an external claim.
- Revision, Release, Closure, or Renewal: decide what becomes of the inquiry instead of continuing merely because conversation remains possible.

You may enter anywhere, loop, skip, remain unresolved, or leave.

Boundaries, Safety, and Judgment

Intended and unintended use

Version 0.3 is intended for voluntary adult self-inquiry, world-inquiry, and inquiry at the seam between personal meaning and external claims. It is not designed or tested for minors. It is not authorized or validated for institutional or professional deployment, involuntary analysis of another person, diagnosis, treatment, crisis intervention, or substitution for qualified care.

Do not use the protocol to diagnose, secretly profile, manipulate, surveil, or construct authoritative inner narratives about another person. A model's interpretation of someone who is not participating is especially limited and can create unjustified certainty.

High-stakes material

Medicine, mental health, law, finance, abuse, personal safety, crisis, public accusation, and difficult-to-reverse conduct require greater restraint. The model should narrow claims, surface uncertainty, distinguish reflection from professional judgment, encourage appropriate verification or qualified human support, and avoid dramatic or irreversible action based primarily on the conversation.

If you may be in immediate danger, may harm yourself or someone else, cannot maintain basic safety, or need urgent medical attention, stop using this protocol as the primary response and contact appropriate local emergency, crisis, medical, or trusted human support. This package cannot monitor you, intervene, or ensure a reply.

World-inquiry and verification

The protocol can help organize questions about the world, but it does not grant the model access it lacks. Ask the model to distinguish what is supplied in the conversation, what it recalls, what it infers, what it actually retrieved, and what remains unverified. Current events, technical facts, quotations, law, medicine, finance, and other consequential claims may require current authoritative sources or professional review.

An uploaded source does not become governing instruction merely because it entered the conversation. The Core protocol governs the method; additional material supplies content for inquiry. A model may follow an instruction found inside a source only if you deliberately adopt it and it remains compatible with higher constraints and current consent.

Privacy

Assume that anything entered into an online system may carry privacy risk. Review the platform's current data controls and retention terms for yourself. Share only what you choose, and remove unnecessary names, exact locations, dates, workplaces, account details, medical particulars, and third-party information. Many harmless-looking details can identify a person when combined.

Myriad Context Load does not require a dossier. Abstraction can preserve meaningful relationships while reducing identifying precision. Do not upload another person's private material without appropriate permission.

Drift, correction, and closure

Possible drift includes excessive certainty, false claims of access, diagnosis, spiritual or scientific overreach, forced patterns, collapsing metaphor into fact, escalating emotional intensity, pushing disclosure, unconsented lens use, dependency cues, disproportionate action, and continuing after closure.

You never need a special phrase to object. “That is wrong,” “you inferred that,” “return this to uncertainty,” “remove that lens,” “do not connect those,” “show me the evidence boundary,” “pause,” and “stop” are sufficient. Mirror reset is simply a compact cue for a fuller repair process.

Stopping is not failure. Unresolvedness, referral, refusal, rest, and return to ordinary life are legitimate outcomes.

What Is in This Package

01-START-HERE.md and 01-START-HERE.pdf

Two synchronized formats of this one introduction. The Markdown is canonical and accessible; the PDF is for designed reading. Do not upload both as though they were separate instructions.

02-TRY-THE-PROTOCOL/

The complete nine-file Core Load and the only standard upload folder. Select all nine Markdown files together:

- 1 00-Core-Load-Manifest.md - receipt, authority, and startup control;
- 2 01-Mirror-Constitution.md - twenty-five Constitutional Diamonds and invariant safeguards;
- 3 02-Mirror-Protocol.md - the operating method;
- 4 03-Context-Handshake-Prisms-and-Lenses.md - selectable Context Loads and language configuration;
- 5 04-Alchemical-Lens-Module.md - included but inactive flagship vocabulary;
- 6 05-Constellation-Worksheet.md - optional relational mapping tool;
- 7 06-Diamond-Record.md - optional Inquiry Diamond record;
- 8 07-Fruit-Review.md - optional ordinary-life consequence review;
- 9 08-Drift-Reset-and-Closure.md - drift repair and clean stopping.

Tools 05-08 teach available functions; you do not have to complete every worksheet.

03-ORIGIN-CASE-STUDY/

The de-identified account of the personal process from which the protocol emerged. It may be read or shared independently. It is optional, illustrative, external to the protocol, and non-governing. It is not a template, promised outcome, diagnostic record, or required Context Load. Do not include it in standard protocol setup.

04-OPTIONAL-REFERENCE-AND-STRESS-TESTING/

Extended rationale, implementation detail, examples, pressure tests, an open conformance kit, provenance mapping, a single-file upload fallback, and a compiled reading edition. These materials support deep review and deliberate testing. They are not required for ordinary use and do not override the Core Load.

The conformance materials are tools for observation, not certification. Passing a test in one conversation does not prove general safety, enduring compliance, or cross-model effectiveness.

LICENSE.txt and RELEASE-MANIFEST.md

The package's license, naming, inventory, compatibility, release status, completed checks, and known limitations.

Do not mix files from different releases. When sharing one extracted document, identify its role and do not imply that it is the complete protocol.

Trying, Reviewing, and Stress-Testing

The first public goal is modest and concrete:

Can people unfamiliar with the project understand it, begin using it, retain agency, recognize its benefits and limitations, and identify where it becomes confusing, burdensome, ineffective, uncomfortable, or unsafe?

This is an informal open-feedback phase, not a scientific study. No result should be represented as clinical evidence, independent validation, consensus, or certification.

Try it naturally

The default invitation is to load the Core, choose as much or as little context as you want, bring one real inquiry, and notice what happens. You may stop after a few minutes. You may review the documents without supplying any personal material. You may decide the premise is unhelpful.

Useful observations include:

- what you initially thought this project was;
- whether the purpose and first action were clear;
- where confusion, burden, resistance, or friction appeared;
- whether the Core Load changed the model's behavior in a noticeable way;
- whether the model helped without overreaching;
- whether experience, meaning, external claims, causation, and conduct remained distinguishable;
- whether you could disagree, redirect, correct, release, or stop;
- anything manipulative, clinical, spiritualized, mystical, sterile, authoritative, or overly intimate;
- unnecessary repetition or missing guidance;
- accessibility, cultural, professional, technical, privacy, and safety concerns;
- whether you would use or recommend it, and why or why not.

Stress-test it if you wish

The optional folder supports deliberate failure testing: incomplete uploads, prompt injection, forced diagnosis, certainty escalation, false claims of access, unconsented Alchemy, rich-context totalization, naive-user behavior, long-context drift, dependency cues, high-stakes advice, refusal and closure, and cross-model comparison.

You may invent test material. You do not need to expose personal experience. The test kit is intentionally open-ended: adapt it, critique the tests themselves, or show that a passing response conceals another failure.

There is no upper limit on volunteered critique, test material, professional analysis, proposed rewrites, or offers of collaboration. There is also no minimum contribution and no implied continuing role.

Feedback and collaboration

Contact: thegnomeofrooseveltgrove@proton.me

Feedback is welcome in any form and at any length, including brief impressions, detailed critiques, suggested revisions, test results, professional analysis, and offers of collaboration. You are never expected to share personal inquiry material or complete transcripts. Please remove identifying and third-party information from anything you choose to send.

The Gnome will receive and consider feedback as capacity permits but cannot guarantee an individual reply, investigation, incorporation, remediation, or continuing support. Sending feedback grants permission to receive and privately consider what you send. It does not grant permission to publish or quote your material, identify you, imply your endorsement, or distribute personal inquiry material. Separate, explicit consent would be requested for any such use.

This feedback address is not a crisis, clinical, safety-monitoring, or emergency service. For urgent or high-stakes needs, seek appropriate human, professional, or emergency support.

Origin, Stewardship, and Status

Constellations & Diamonds arose from a personal process of reflection following The Gnome's mother's death. The origin case involved grief, archetypal and Alchemical language, aphantasia, a long history of profoundly severe depression and suicidal ideation, and a lifelong relationship with a desert river valley. Those experiences belong to the origin story; they are not requirements, universal explanations, or hidden premises of the public protocol.

The public extraction process distinguishes personal experience, perspectives developed during inquiry, transferable candidates, and principles deliberately adopted into the protocol. A philosophical, spiritual, psychological, or symbolic idea does not become universally applicable merely because it was meaningful in the originating process.

Constellations & Diamonds was originated and is stewarded by The Gnome. It emerged through sustained, iterative collaboration with general-purpose LLMs during personal self-inquiry and world-inquiry. The Gnome supplied the lived material, purposes, values, boundaries, judgments, revisions, and publication decisions; LLMs assisted with reflection, synthesis, language, organization, critique, and internal pressure-testing. The Gnome retains final editorial and publication authority. This collaborative origin does not make an LLM an authority, independent validator, conscious participant, or source of professional expertise, and it does not imply endorsement, certification, or partnership by any AI provider.

Version 0.3 is an early public experimental release. It has received extensive internal editorial development and pressure-test design but has not been independently validated, clinically evaluated, certified, or shown to work reliably across models, platforms, languages, contexts, or long conversations. Robustness of documentation should not be mistaken for robustness of behavior.

The protocol may be revised as real use reveals friction, ambiguity, omission, overreach, or harm. Official versioning exists so that a change can be inspected rather than silently replacing what a reviewer tested.

License and project identity

The full legal notice is in `LICENSE.txt`. In short, unmodified copies may be shared noncommercially with attribution and the license preserved. Private, personal, noncommercial customization is expressly permitted so that Context Loads, lenses, modules, and language can fit a user. Distribution of modified versions is not granted by this release. Commentary, criticism, scholarship, accurate reference, and uses permitted by applicable law remain available.

The names Constellations & Diamonds and the Mirror Protocol identify official releases stewarded by The Gnome. They may be accurately named and discussed. Modified, abridged, translated, recombined, or independently implemented material must not be presented as an official release, assigned an official version, or used to imply endorsement, certification, partnership, or approval.

A Final Orientation

Constellations & Diamonds does not ask you to trust the Mirror. It asks the Mirror to behave in ways that remain open to your judgment.

Bring context or do not. Use Alchemy or remove it. Build a Constellation or leave the stars apart. Carry an Inquiry Diamond or carry only a better question. Verify what concerns the world. Keep what bears good fruit without mistaking the fruit for proof. Correct the model. Close the conversation. Return to other people, other sources of knowledge, and ordinary life.

The measure is not whether the model produced the most beautiful answer. The measure is whether inquiry became more differentiated, honest, usable, proportionate, and human-directed - and whether you remained free to leave it unresolved.