

The Zemi Method

An evidentiary methodology for investigation and assurance

Version	1.1
Status	Published. Pressure-test-driven author revision; not peer reviewed or empirically validated.
Published	July 25, 2026
Canonical home	https://zemimethod.com
Copyright	© 2026 Kevin V. Watson. All rights reserved.

This document is a formatted rendition of the canonical Markdown source of The Zemi Method, Version 1.1. The authoritative text and its integrity hash are maintained at the canonical home above and in the method's version repository. Where this rendition and the canonical source differ, the canonical source governs.

Premise

Every important organizational decision is an evidentiary decision.

Courts, incident response, internal investigations, AI governance, board decisions, compliance, fraud, and digital forensics all eventually ask the same question:

Can this claim be trusted, and what decision can be defended because of it?

Most failures of evidence-based decision-making do not begin with a lack of tools. They begin because an assumption went untested. A system was trusted to record what happened. A log was read as complete. A control was believed to be working because it was documented as working. A summary was treated as the record instead of a path back to the record.

The Zemi Method starts from the opposite position. Every claim that matters to a decision, including claims produced by an organization's own systems, is a hypothesis until evidence from an independent source supports it.

This is not caution for its own sake. It comes from a specific and repeatable failure. An audit record can appear authoritative and still be incomplete in a way that only shows under corroboration. An application-layer log that captures most access will present as the full account, while a database-layer record shows access paths the application never wrote down. The application record is not false. It is partial, and partial in a direction no one flagged. The method exists to catch that failure before it becomes the basis for a decision.

Scope

The Zemi Method applies wherever a decision depends on digital or machine-generated evidence. Its first applications are digital forensic examination, incident reconstruction, expert evidence, litigation support, AI governance, AI incident examination, and synthetic media assessment. Future applications may change as evidence changes, but the reasoning discipline remains the same.

The method does not replace engagement-specific procedures, legal advice, lawful authority, tool validation, chain-of-custody requirements, peer review, or professional judgment. It governs the reasoning above them: how claims are framed, how sources are tested, how automation is controlled, how conclusions are bounded, and how findings are reported. The depth of verification applied in an engagement is proportionate to the stakes of the decision it serves; where constraints of time, cost, authority, or access limit verification, that limitation is justified in the report, stated as a limitation, and reflected in the classification of any affected finding.

The method may be used in three modes. The mode is determined by the engagement's purpose, decision context, and required form of opinion, not by the examiner's preference. The selected mode and the reason for selecting it are stated in the report.

- **Investigative mode**, where the purpose is to determine what the available evidence can establish about an event, system, claim, or sequence.
- **Preliminary evaluative mode**, where competing propositions are considered and the practitioner

expresses what the available evidence appears to support, subject to stated limits and without claiming full evaluative finality.

- **Fully evaluative mode**, where the practitioner is asked to evaluate evidence against formal competing propositions for a legal, regulatory, or adjudicative decision. Where a formal likelihood-ratio, probability, or discipline-specific evaluative framework is required, that framework governs the expression of evidential strength. The Zemi Method may still govern the framing, preservation, corroboration, automation control, classification, and reporting discipline around that work, but it does not replace the required evaluative framework.

The public method states the philosophy and control structure. Implementation detail, including tools, queries, templates, checklists, and internal review procedures, remains engagement-specific and proprietary.

For Zemi North, the method is not a service line. It is the governing discipline from which service lines, reports, assurance frameworks, software, training, and future research should derive.

Principles

Five principles govern every engagement, on both the investigation side and the assurance side.

1. Evidence over assertion. A conclusion stands on what can be shown, not on what a system or a person states. A statement of fact from a stakeholder, a vendor, or a monitoring tool is an input to be tested, not a finding to be recorded.

2. Corroboration and completeness. A single untested assertion is a lead. A source-derived proposition becomes a finding only when the source, extraction, and failure mode have been tested to the proportionate limit of the engagement. Two logs drawn from the same pipeline are not independent. A log and the artifact it claims to describe are. A source is also tested for whether the extraction taken from it is complete, because an incomplete extraction misrepresents the source as surely as the wrong source would. Where independent corroboration is available, it is required. Where independent corroboration is unavailable, the finding is reported with that limitation attached rather than hidden.

3. Bounded claims. State what the evidence supports and no more. Where the evidence runs out, say where. Every material conclusion is separated into what is known, what is assumed, and what remains undetermined. These classifications describe the epistemic status of a proposition in the report, not the strength of support on a probability scale. A finding that overreaches is weaker than a narrower finding that holds, because the overreach is the part an opposing examiner will break first.

4. AI-assisted, never AI-decided. Automated systems retrieve, correlate, summarize, propose patterns, and draft. They do not adjudicate. A person owns every conclusion and can trace it to source. This is a line, not a preference.

5. Defensibility by construction. Build the case so a reviewer, an opposing expert, or a court can follow every step from raw evidence to conclusion without taking any part on trust. Defensibility is designed in from the authority and scope gate onward, not added at the end.

Finding classifications

The method uses three classifications at the human adjudication gate. They are categories of kind, not degrees of confidence.

Known means the proposition is supported by evidence within the stated scope and limits of the engagement. It does not mean absolute certainty, and it does not mean the support is necessarily strong in a probabilistic sense. It means the report may state the proposition as established for the decision at hand, subject to the limits recorded. For each known finding, the report states the support basis: whether the finding rests on direct evidence, indirect evidence, single independent corroboration, multiple independent corroboration, inference from a bounded record, or a tool-derived measurement with stated processing and uncertainty limits. The support basis explains how the finding is supported without converting the classification into a probability or verbal strength scale.

Assumed means the proposition is relied on for a stated and limited purpose without being directly established by the available evidence. An assumption is not hidden inside a finding. It is identified, bounded, and kept separate from what is known.

Undetermined means the proposition cannot be resolved on the available evidence. The correct outcome may be that evidence is unavailable, conflicting, incomplete, outside authority, outside scope, or insufficiently independent. Undetermined is not a failure to decide. It is a finding that the record does not support a stronger conclusion. An undetermined matter is reported with a resolution status: **pending**, where identified evidence or work remains outstanding and resolution is still being pursued, or **exhausted**, where the available and obtainable evidence has been pursued to the proportionate limit of the engagement and the matter remains unresolved.

The investigative lifecycle

Seven phases carry an engagement from authority to finding. The role of automation is stated at each phase, and it narrows deliberately as the work approaches a conclusion.

0. Authority and Scope.

Before the question is framed, determine whether the examiner is legally, contractually, ethically, and professionally permitted to ask the question, obtain the evidence, and perform the work. Identify the client, decision-maker, mandate, jurisdictional or policy limits, conflicts, excluded systems, access constraints, privileged material, required standards, and any limits on release.

Authority to perform the work is separate from authority to disclose evidence, findings, or reports. If authority to perform the work is absent, the engagement stops and no Zemi-governed finding is issued. If authority exists but scope, access, mandate, or disclosure authority is limited, the engagement may proceed only if the limitation is justified, stated in the report, and reflected in the classification of any affected finding.

The method does not create authority. It requires authority to be identified before evidence work is treated as defensible, and disclosure authority to be identified before findings or supporting material leave the practice.

1. Frame.

Define the question. Identify the mode of work: investigative, preliminary evaluative, or fully evaluative. Surface the assumptions the question carries. Write down what would have to be true for the working theory to hold, and what evidence would confirm it or break it.

Where the evidence admits more than one explanation, competing explanations are framed and tested against the same evidence together, not sequentially, so that the surviving explanation is the one the evidence failed to break rather than the one considered first. A theory with no stated breaking condition is not yet an investigation.

2. Preserve.

Acquire and protect evidence with integrity intact. Chain of custody, hashing, work from verified copies, and record the condition and source of each artifact.

Before an extracted result, query output, parsed artifact, export, or summary is relied on, assess whether the extraction appears complete for the source and question at issue. Prior reports, prior examiner findings, and prior investigative summaries are preserved as source material, not inherited as established findings, unless their underlying evidence and reasoning are independently inspected or re-performed.

Material filters, date ranges, permissions, pagination, parser limitations, excluded fields, failed records, timezone handling, and known retention limits are identified where they may affect the result.

Where evidence or forensic copies are packaged, encrypted, transferred, or delivered, the record identifies the packaging or encryption method, transfer path, authorized recipient, integrity verification, and receipt confirmation. Where the engagement requires evidence, devices, temporary media, workspaces, or data to remain within a defined location or custody boundary, the record identifies that boundary, who observed or controlled it, what material remained inside it, any temporary media created, and how non-transfer or retained custody was verified.

Where automated systems are used, preserve the inputs provided to them, the outputs they return, and the tool and version records needed to test their contribution later. Nothing downstream is defensible if this phase is weak, so it is not rushed.

3. Corroborate.

Test each claim against independent sources wherever they are available, and report the limitation where they are not. Corroboration includes testing whether the result taken from a source is a complete and faithful extraction of that source for the question being asked; an incomplete or unvalidated extraction is treated as a limitation, not as the source itself.

A prior report may identify a lead, a claimed finding, or a gap, but it does not corroborate itself. Where a prior finding is material, the practitioner tests the underlying evidence, reasoning, and failure mode rather than relying on the prior conclusion as evidence of the conclusion.

Where independent sources conflict, the conflict is treated as evidence in its own right and investi-

gated, not averaged and not resolved toward the working theory; a conflict that cannot be resolved is classified as undetermined, with both sources shown.

Automation earns its place here by pulling and aligning records across systems at a speed a person cannot match. The person decides what agreement means, where the sources fail to be independent, and whether the extraction from each source can be trusted.

4. Analyze.

Reconstruct the sequence of events, evaluate competing explanations, and identify what the evidence supports, contradicts, leaves open, or requires as an assumption.

This is where automation carries the most weight and where the discipline matters most. It proposes patterns. The investigator confirms or rejects each against the underlying evidence rather than against the summary.

5. Adjudicate.

This is the human decision gate. Conclusions are set here, by a person, bounded to what the evidence supports. Findings are classified as known, assumed, or undetermined before they are relied on.

Known findings state their support basis. Undetermined findings state whether their resolution status is pending or exhausted. The classification is separate from any statement about evidential strength.

No automated system crosses this line. This is principle four made procedural.

6. Report.

State each finding, the evidence behind it, and its limits. Material evidence inconsistent with a finding is reported alongside it, not omitted. For each material finding, the reasoning that connects the evidence to the conclusion is stated, not left for the reader to reconstruct.

A reader should be able to see where certainty ends and inference begins, and should never have to guess which is which.

Reports identify the authority and scope relied on, the mode of work and why it was selected, the depth of verification applied and why it was proportionate, the sources relied on, the completeness limits of material extractions, the support basis for known findings, the resolution status of undetermined findings, the assumptions made, the material limitations, the role of automation where automation contributed to the work, and the AI role classification where an AI system is material to the engagement.

Where evidence narrows a suspect, actor, account, device, system, or source population without uniquely identifying the responsible person, system, or source, the report states the remaining candidate pool, the criteria used to narrow it, the evidence supporting those criteria, the level of attribution actually supported, and the evidence still required for stronger attribution.

Where a finding depends on layered infrastructure, intermediaries, platforms, CDN services, cloud services, domains, accounts, hosting, payment records, or financial records, the report distinguishes technical delivery, account ownership, operational control, legal responsibility, and financial benefit

where material, and states which layer each item of evidence supports.

Where a finding depends on compliance with, alignment to, or deviation from an external standard, framework, policy, or benchmark, the report identifies the standard and version where material, the specific requirement or control relied on, why it applies, the evidence mapped to it, and any interpretive limits.

Where a finding depends on a tool-derived measurement, enhancement, comparison, calculation, or transformation, the report states the tool and version where material, the processing or measurement basis, material settings or steps, validation checks, and any known error bounds or uncertainty limits.

Where a finding relies on an excerpt, compilation, edited video, demonstrative exhibit, visualization, transcript, timeline, or other derivative presentation of source evidence, the report preserves or identifies the underlying source, states the selection or editing criteria, links each excerpt or segment to its source location and timestamp or equivalent locator, identifies material excluded context where relevant, and states that the derivative presentation is not a substitute for the source evidence.

Before a report, finding, or supporting evidence is released, the report record identifies who is authorized to receive it; any privilege, confidentiality, statutory, contractual, or policy limits on disclosure; where material is transferred, the packaging or encryption method, transfer path, integrity verification, and receipt confirmation; and where material is deliberately retained within a custody boundary, the boundary, observer or controller, retained material, temporary media handling, and verification of non-transfer.

The assurance practice

The same discipline applies to a different subject. On the investigation side the method tests claims about what happened inside a system. On the assurance side it tests claims about an AI system itself: that it does what it is said to do, within stated limits, with a person accountable for its outputs.

The five principles carry over without change. A model's own reporting about its behavior is a claim to be corroborated, not a finding to be accepted. A vendor's stated accuracy is an assertion until independent testing supports it. A control described in documentation is a hypothesis until an artifact shows it operating. The subject changes. The evidentiary standard does not.

For assurance work, the method does not ask only whether a system appears to perform. It asks whether the evidence supports the claims made about the system: what it does, where it fails, who remains accountable, and whether the controls around it can be shown to operate in practice.

Authority and scope matter with particular force in assurance work. An examiner must know whether the engagement permits testing against live systems, synthetic datasets, production records, protected classes, model outputs, prompt records, audit logs, configuration files, or vendor-controlled evidence. Because assurance testing may reveal new evidence sources, system boundaries, or access constraints only after work begins, authority and scope are set at the outset and revisited when access changes. Where an assurance conclusion depends on evidence the examiner cannot access, that limitation is reported rather than absorbed into the finding.

AI role classification

Where an AI system is material to an engagement, its role is classified before conclusions are drawn. The classification describes what the evidence shows the system did. It does not assign legal responsibility, and it does not accept labels, logs, prompts, or vendor statements as proof of the system's role without corroboration.

- **AI present:** an AI system exists in the environment, artifact set, or workflow, but the evidence does not show that it materially affected the activity under examination.
- **AI-assisted:** an AI system supports a human actor through drafting, troubleshooting, summarization, coding, search, orchestration, or analysis, but the evidence does not show the system independently selecting objectives or executing material actions.
- **AI-orchestrated:** an AI system coordinates or triggers material workflow steps under human-configured rules, prompts, agents, tools, or integrations, while accountability remains with the human or organization controlling the system.
- **AI-autonomous claimed:** a claim that an AI system independently selected objectives, made decisions, or executed material actions without human control. This classification is not accepted from system labels, logs, prompts, vendor statements, or actor claims alone and requires independent evidence.

The role classification is reported with its evidence and limits. Where the evidence shows only assistance, the report does not describe the conduct as autonomous. Where autonomy is claimed but not independently supported, the claim remains undetermined.

Why the principles exist

The method is not theoretical. Each principle exists because real investigations and assurance work fail in predictable ways when the principle is absent.

- **Evidence over assertion:** a stakeholder's confident account can be sincere, useful, and still wrong. The method treats that account as a lead until the record supports it.
- **Corroboration and completeness:** an application log may show the ordinary route through a system while a lower-layer record shows a second path the application never captured. The first record is not false. It is incomplete.
- **Bounded claims:** a narrow conclusion with clear limits is more durable than a broad conclusion with one unsupported edge. The unsupported edge is where the finding will be attacked first.
- **AI-assisted, never AI-decided:** an automated summary can accelerate review, but it cannot explain itself under challenge, accept responsibility, or distinguish uncertainty from absence unless a person forces that distinction.
- **Defensibility by construction:** if the chain from source to conclusion is assembled only after the conclusion is written, the work is already vulnerable. The record has to be built while the work is being done.

Where automation sits, and where it does not

Automation is allowed to touch retrieval, correlation, pattern proposal, drafting, and summarization. It is barred from the conclusion, the adjudication, and the sign-off.

The reason is not distrust of the tool. It is that a conclusion has to be owned by someone who can be questioned, who can trace the finding to source, and who can be held to account for it. A summary cannot be cross-examined. A person can. The line holds because accountability requires it.

The automated contribution is preserved as well as documented. The inputs provided to automated systems, the outputs they return, and the tool and version used are retained as part of the engagement record, so that the separation between automated contribution and human judgment can be demonstrated later rather than asserted.

Peer review and validation status

Publication of the method makes it dated, citable, and open to scrutiny. It does not make it peer reviewed, empirically validated, or accepted by the field.

The method is therefore maintained in two forms at once: as a governing discipline that can be applied in practice, and as a research object that must be tested. The known validation questions include whether different practitioners classify findings consistently, whether readers understand the classifications as intended, whether the authority and scope gate improves defensibility, and whether the method changes reporting quality without imposing disproportionate burden.

Structured expert review, case-based testing, inter-rater reliability work, reader-comprehension testing, and eventual peer-reviewed publication are future validation paths. Inter-rater reliability work requires multiple practitioners or external collaborators, so it remains a future validation path rather than a test the present solo practice can complete alone. Until that work exists, the method's status is bounded: it is a published practitioner methodology under review, not a validated scientific technique.

Transition note for examples and templates

Materials prepared under Version 1.0 do not automatically satisfy Version 1.1. Worked examples, report skeletons, templates, checklists, and visual summaries must be revised before they are described as Version 1.1-conformant. At minimum, they must include the Authority and Scope phase, identify the mode of work, and justify any proportionality, access, or scope limitations relied on in the findings.

The defensibility check

Before a finding leaves the practice, it passes a standing test:

- Is lawful, contractual, ethical, and professional authority identified?
- Is disclosure authority identified before evidence, findings, or reports are released?
- Where evidence or forensic copies are transferred, are packaging or encryption, transfer path, authorized recipient, integrity verification, and receipt confirmation recorded?
- Where evidence, devices, temporary media, workspaces, or data are required to remain within a defined custody boundary, is the boundary, observer or controller, retained material, temporary media

handling, and verification of non-transfer recorded?

- Is the scope of the question, evidence access, and excluded material stated?
- Is the mode of work identified as investigative, preliminary evaluative, or fully evaluative?
- Is the selected mode justified by the engagement's purpose and stated in the report?
- Are proportionality, access, or scope limits justified rather than merely noted?
- Have material queries, exports, parses, extracted results, and summaries been assessed for completeness against the source and question at issue?
- If prior reports, prior examiner findings, or investigative summaries are material, have they been treated as source material and tested against the underlying evidence rather than inherited as findings?
- Can every conclusion be traced to a source a reviewer could inspect?
- Are the claims bounded to what the evidence supports?
- Are known facts, assumptions, and undetermined matters separated?
- Do known findings state their support basis without converting the classification into a probability or verbal strength scale?
- Where evidence narrows a suspect, actor, account, device, system, or source population without uniquely identifying responsibility, are the candidate pool, narrowing criteria, supporting evidence, attribution level, and evidence still required for stronger attribution stated?
- Where a finding depends on layered infrastructure, intermediaries, platforms, CDN services, cloud services, domains, accounts, hosting, payment records, or financial records, are technical delivery, account ownership, operational control, legal responsibility, and financial benefit distinguished where material?
- Where a finding depends on compliance with, alignment to, or deviation from an external standard, framework, policy, or benchmark, are the specific requirement, applicability basis, evidence mapping, and interpretive limits stated?
- Where a finding depends on a tool-derived measurement, enhancement, comparison, calculation, or transformation, are the processing basis, validation checks, and uncertainty limits stated?
- Where a finding relies on an excerpt, compilation, edited video, demonstrative exhibit, visualization, transcript, timeline, or other derivative presentation, are the underlying source, selection or editing criteria, source locator, material excluded context, and non-substitution limit stated?
- Do undetermined findings state whether resolution is pending or exhausted?
- Is the classification of each finding consistent with the evidence record?
- Would the finding hold if an independent examiner ran the work again from the preserved evidence?
- Is the contribution of automation documented and separable from human judgment?
- Where an AI system is material, is its role classified and bounded to what the evidence supports?
- Are material limitations stated clearly enough that a decision-maker can account for them?

A finding that fails any of these returns to the phase where the gap opened. Nothing reaches a client on the strength of the method's reputation alone. It reaches them on the strength of the record behind it.

Version discipline

The Zemi Method is maintained in dated, versioned form. Each public version records its publication date and material changes from the prior version. Engagement reports identify the version of the method in use where the method is relied on.

Versioning matters because the method is not written after the fact to fit a conclusion. It exists before the engagement, guides the work during the engagement, and constrains the finding at the end.

Version 1.1 is released on the basis of documented pressure testing rather than completed external review or engagement-based validation. That release basis is recorded here because the method's own discipline requires material judgment calls to be stated rather than hidden. External review and engagement-based validation remain open and will inform future versions.

The same discipline applies to findings. If a finding issued under the method is later shown to be unsupported, incomplete, or materially misstated against the preserved evidence, a correction is issued that identifies the finding, the reason, and the evidence, with the same visibility as the original.

The Zemi Method is maintained in dated, versioned form at its canonical home, <https://zemimethod.com>. Comments and scholarly correspondence are welcome and inform future versions.