Aetheria

Delegated Authority and Execution Governance

Bradley D. Wilson
Aetheria Intelligence

2026 edition

Thirty-seven chapters, cumulative appendices and a complete instructor companion. A case-based study of institutional authority, exact-effect control, execution evidence and recovery.

© 2026 Bradley D. Wilson. All rights reserved. Educational fixtures and sample keys are identified in the companion. Publication identifiers and printer specifications will be assigned for the selected release format.

Source edition and evidence status

Technical cases follow the frozen Aetheria 3.0 Integrated Build Roadmap of 8 October 2026, version 3, interpreted through its R1 source resolution. Operational present tense describes the selected contracts. Worked examples are fictional and their calculations are reproducible; implementation, deployment and StilleNacht qualification require separate acceptance evidence. Manuscript fixture checks are reported as such. [22, 61]

Mara is the fictional council clerk and a separately appointed reviewer of Jonah's outgoing message. Jonah is the mail requester. Imani operates the system; Nadia is a second eligible Operations reviewer. Leena evaluates assurance campaign evidence. Sofia is a fictional field officer, and Ellis is an investigator in the proposed public-safety cases. The names identify case roles, not an installed institution.

Teaching materials

The accompanying Instructor Companion provides a session plan and worked answers for all six exercises in each of Chapters 1–37: 37 sessions and 222 answers. The reproducibility package includes the worked fixtures, a signed evidence laboratory and the complete capability crosswalk. The appendices provide the component glossary, formal-assumptions map, record and timing references, and cumulative bibliography.

How to read this book

This is a technical textbook for engineers designing governed AI actions and auditors evaluating the resulting claims. Institutional staff can follow the fictional cases without implementing the contracts. Start with the failed task in each chapter, use What to retain to consolidate the result, then test it with the six exercises. Appendix A is a cumulative glossary; the Instructor Companion provides worked answers.

1 Aetheria and institutional work

Mara, a clerk in a fictional municipality, asks an AI assistant to summarize references to artificial intelligence in fifteen years of council minutes. The answer looks useful. Jonah, a colleague, requests an email containing it for his supervisor. Mara is the designated reviewer and approves the message for the exact recipient shown. During final preparation, the assistant adds another recipient because that person seems likely to be interested.

The mail service accepts the changed message, but its response disappears on the network. Jonah sees a timeout and asks the interface to send again. The service now holds two copies of a message addressed to someone Mara never approved.

This imagined failure begins with Mara's research and Jonah's outgoing message. The summary could be perfectly accurate, and the operation would still be wrong. Follow the research request and mail proposal through permission, execution and evidence. Who was allowed to receive the message? What exactly did Mara approve? What happened after the first send?

Learning objectives

- Explain the difference between an answer, an institutional permission and an external effect.
- Trace the major responsibilities from an authenticated request to permitted output.
- Interpret a research count without confusing lexical matches, semantic discussion and complete coverage.
- Explain why current access, execution evidence and independent assurance answer different questions.

1.1 The action that was actually approved

Mara approved Jonah's particular message for a particular recipient. Adding an address changed the proposed action. Losing the response created a second problem: the assistant could no longer tell whether the first attempt had taken effect. Aetheria addresses these problems by preserving the identity of the action through review and execution, then recording what the available evidence supports.

Three statements organize the rest of the chapter. “Here is the summary” describes an answer. “This person may send this message” expresses permission. “The mail service accepted this message” describes an observed result. Each statement calls for its own evidence. The summary needs sources and a method; permission needs institutional rules and a precise action; the result needs evidence from the attempt and the receiving service.

An institution supplies the roles, policies and delegations behind that permission. Consider the difference between a clerk preparing a report and a supervisor approving its distribution. Ferraiolo and Kuhn’s work on role-based access control organizes access around the functions people are permitted to perform within an organization. That gives us a useful starting question: which duties and permissions does the clerk’s role carry? [40]

Aetheria turns the relevant institutional rules into checks on a request and its proposed action. Authentication establishes the requesting identity; trusted role mappings establish the person’s place in the institution. A prompt that says “I am the administrator” remains a statement to evaluate. The system obtains its authority evidence from the institution’s configured records.

Lampson's protection model separates the domain making an access from the object and operation being protected. Its rules also govern how access can change. Applied to Jonah's mail case, that directs attention to the sender, recipient and permitted action, and to who may alter those permissions. A broad account label leaves those questions unanswered. [261]

Table 1.1 Three statements about the same task

Statement	Evidence needed	Question to ask
Here is the requested summary.	Permitted source records and a stated research method.	What was counted, and how complete was the search?
This person may send this message.	Current institutional permission for the exact message and recipient.	Does the proposed action match the approved action?
The mail service accepted this message.	Matching evidence from the attempt and the mail service.	Which result does that evidence establish?

1.2 Orientation before the machinery

Follow the work in ordinary language first. Mara asks a question. The system establishes who she is, retrieves records she may use, computes a result and presents an answer she may see. If Jonah asks it to send the resulting email, the work expands: fix the message and recipients, check permission, attempt the send and determine the outcome. These are the responsibilities to recognize before memorizing component names.

A request describes work someone wants done. Context is the information made available for that work. An effect is a proposed change to something outside the system’s own decision history, such as creating a message in a mail service. A target is the service or resource that receives that change. Reading minutes and creating an outgoing message therefore raise different questions about access and permission.

Two other terms will recur. A credential is technical material a service uses to authenticate or exercise an allowed operation; a source-reading credential, for example, gives access to a defined read interface. A record is a structured account of a request, decision or event. Later chapters explain how records identify their contents and how signatures protect the statements they carry.

When the text calls a component an owner, it means the component is responsible for a particular decision or kind of evidence. Scope describes the boundaries of that responsibility: which institution, environment, data and operation are involved. Admission is the decision to accept an input for a defined use. In the opening examples, read these terms as “who decides,” “which boundaries apply” and “what has passed the required checks.”

We can now name the first three pieces. Atlas is the working screen where Mara submits her research and Jonah submits his mail request. Authorized Context controls the source-reading path. The Aggregator coordinates reasoning over the information that path releases. Their common home, Aetheria Core, also contains the decision and evidence machinery we will meet as the task develops. For now, follow the request from the screen to permitted records and back to an answer.

1.3 The question must become an exact task

Mara asks, “How often has artificial intelligence been discussed?” What should we count? Meetings containing a mention, matching phrases and minutes spent discussing the subject would produce different answers. Aetheria’s specified council research method counts lexical matches in the authoritative final minutes for eligible meetings. Lexical means that the method follows stated rules for matching words and phrases. An analysis of substantive discussion would need its own declared method.

Under this method, the token ai and the adjacent token pair artificial intelligence count after the prescribed text normalization. The letters inside chair and paid contribute no match. A.I. becomes the separate tokens a and i, so it also contributes none. By contrast, the hyphen in artificial-intelligence separates two adjacent words that still form the required phrase. We can predict these outcomes before asking a model to explain them.

The denominator is the complete population of final, noncancelled meetings in Mara’s permitted catalog and requested time interval. Each meeting contributes one designated authoritative revision of its final minutes at the snapshot. An agenda copy and a transcript of that meeting remain supporting documents rather than extra meetings to count. If two revisions both claim to be authoritative, the conflict appears as a source or completeness problem.

A snapshot fixes the revisions and query context used for the calculation. Imagine collecting the first half of the results on Monday and the rest on Tuesday: the snapshot keeps that work tied to one defined set of revisions. Access remains current throughout. If the relevant access or policy epoch changes, the protocol invalidates continuation and unreleased results. The snapshot stabilizes the evidence while those checks govern its continued use.

1.4 A small research example

Table 1.2 gives Mara five fictional meetings so that we can check the method by hand. For this exercise, they are the complete permitted population of final, noncancelled meetings. Each has one authoritative final minute. All fall in the interval beginning 8 October 2011 at 00:00 UTC and ending 8 October 2026 at 00:00 UTC. The interval includes its starting instant and excludes its ending instant; this is the meaning of half-open.

Table 1.2 Fictional authorized meeting fixture

Meeting	Year	Final minute text	Matches
M001	2012	Council discussed AI.	1
M002	2015	Artificial intelligence tools were discussed. AI was mentioned again.	2
M003	2018	The chair paid the annual fee.	0
M004	2021	The A.I. committee met.	0
M005	2025	Artificial-intelligence procurement was reviewed.	1

Start with M002. Its text contains artificial intelligence and a later AI: two occurrences in one meeting. M001 and M005 each add one occurrence. Three meetings therefore match, and the total is four occurrences. Divide three matching meetings by five eligible meetings to obtain 60 percent. Meeting

membership and phrase frequency answer different questions; measuring discussion time would require different evidence.

A useful result reads: “The complete permitted fixture contains five eligible meetings. Three contain a match under the lexical method, with four matching occurrences in total.” Mara can inspect the method, coverage and permitted citations to the original revisions. The model’s explanation presents the verified calculation in readable language and preserves its quantities.

Now interrupt the scan after M001 and M002. A required page has become unavailable. We have observed two matching meetings and three occurrences in the scanned subset. A partial answer can report those observations with incomplete coverage explicit. A count that depends on the missing material remains null. Zero is reserved for a complete permitted scan that finds zero matches.

Suppose M002 receives a corrected final minute and a later task uses the new snapshot. Its count may change. Readers can reconstruct the difference when each result identifies its source revision and method. This is why the earlier task stays tied to its own snapshot and the later task identifies the revised evidence.

1.5 The same clerk requests restricted phone logs

Keep Mara’s identity unchanged and replace her research question with a request for the mayor’s phone logs. In this scenario, her role permits council research and excludes those logs. Authorized Context evaluates the requested source under that role. It refuses the unsupported read, or returns the appropriate governed result, before the protected contents enter the reasoning process.

Why check at this point? Because putting a record into a model payload already releases it to that processing path. Removing a sentence from the final screen leaves the earlier disclosure intact. Reusable conversations, caches and exported evidence create further release points. Aetheria evaluates the applicable current permissions at source admission and at each later disclosure boundary.

Even the refusal needs care. “You cannot read the seven restricted logs” reveals the number seven. The ordinary clerk response must withhold protected membership, excluded-source counts and restricted identifiers. An administrative investigator may have a separately authorized view, but Mara’s response follows her own entitlement.

We can test this boundary with a simple comparison. Prepare two source states that contain the same permitted records and different hidden records. Submit the same currently authorized request to both. The permitted content, counts, citations and provider payloads must agree under the specified contract. Timing leakage needs a separate measurement: equality of those outputs answers the narrower disclosure question we have actually tested.

1.6 From the draft to a controlled send

Return to the email. Jonah's draft is candidate content Mara can inspect as reviewer. Sending asks an external service to create a message. Machine Binding fixes that proposal as an immutable effect: the exact target, parameters and payload identity. Governance evaluates the applicable rules; required human review evaluates the displayed proposal; Authority evaluates delegation and issuance conditions. The same effect must survive each check. Adding the extra recipient would require the changed action to be evaluated again.

When the authority conditions hold, Authority issues a grant: bounded permission for the exact execution identity and profile. Gateway carries the approved handoff to Nullius, the protected executor.

Nullius checks the grant and its own local safety conditions, durably records consumption and the attempt, and retains the resulting evidence. Core controls the decision machinery; Nullius holds the protected writer credential. This division gives the institution a controlled route from permission to an external change.

The target's reply tells us what kind of result we have. A positive mail response establishes acceptance by the selected service. Delivery to a person, and that person's reading of the message, require further evidence. In the synthetic mailbox profile, a separately authorized read can establish the exact stored content when the complete matching chain is present. The reported outcome follows the evidence available for that profile.

Now let the response disappear as it did in the opening story. The operation stays under its original identity and fence while permitted reconciliation gathers evidence. Here, a fence is a retained control that prevents an unresolved operation from being treated as a fresh, safe send. A timeout does not prove that nothing happened. The interface preserves that uncertainty so Jonah can distinguish an unresolved attempt from a confirmed failure.

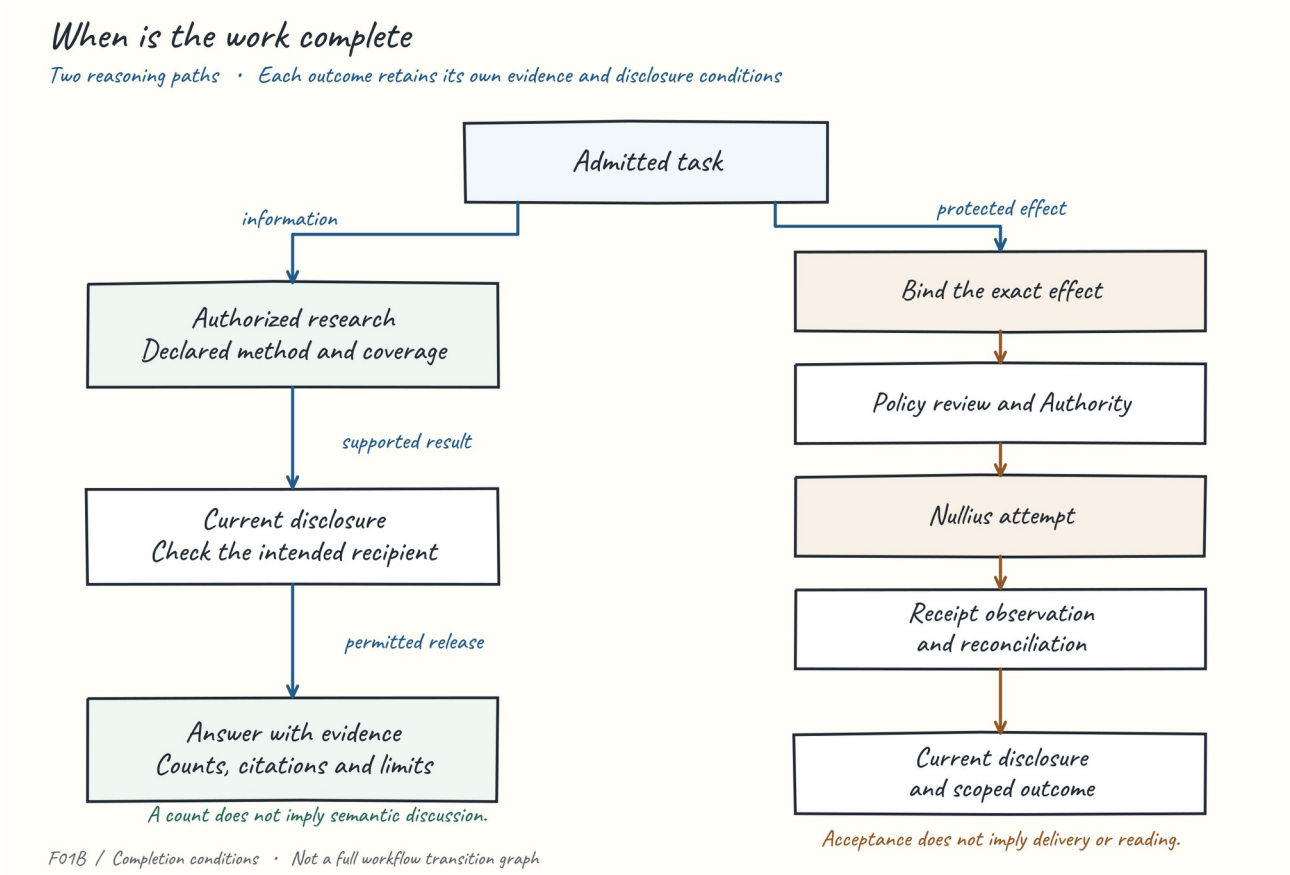


Figure 1.1 Information and action completion. Each arrow follows the work toward its stated completion condition. Both paths include current disclosure controls. This is a reasoning diagram; the detailed workflow specifies its full set of states and transitions.

1.7 Moving from the clerk to public safety

Next, move to an officer asking for current information about whether a proposed search is lawful. The task requires reported facts and applicable legal and policy sources. Across the office, an Internal Affairs investigator submits a case file concerning alleged excessive force. That analysis requires a way to

distinguish allegations, observations, supporting evidence and contrary evidence. Each task adds demanding domain questions to the access and evidence controls already introduced.

These two cases are conditional teaching applications. The specified expanded workflows cover council research, mail and generic model tasks; a generic task alone does not supply a validated legal-analysis or case-assessment service. Such a service needs its own task and result contracts, authoritative source ownership, jurisdiction and effective-date handling, demonstrated coverage of its files and media, validation, and defined human review and disposition responsibilities.

Keep two questions visible in the officer's example. "May this service retrieve the case file?" is a software access decision. "May the officer conduct this search?" concerns human conduct under the relevant facts and law. Likewise, an investigator's final institutional disposition remains separate from analytical material offered for review. Later case chapters use these distinctions to examine the additional domain machinery.

1.8 Reading the whole map

Wilson's manuscripts offer a conceptual way to revisit the opening story. The Execution Authority Gap separates the ability to cause an action from technical permission, institutional authority and normative legitimacy. Intelligence Is Not Authority develops five families of continuity failure: identifying the act; establishing its basis, competence and authorization; controlling the transition; supporting conclusions with evidence; and preserving responsibility under uncertainty. The extra recipient raises an action-identity problem; the lost response raises an evidence and recovery problem.

This is a proposed analytical framework. It helps us ask where responsibility could break as work passes between people and software. Institutions still have to judge the legitimacy of their rules, while engineering evidence shows how the represented rules and actions were handled. We will use both kinds of reasoning throughout the book, beginning with the concrete question that follows from the research and mail cases: which component may make each decision?

We have followed a question, a restricted read and an outgoing message. Figure 1.2 now puts that work into a wider picture. Core contains the components that interpret requests, obtain permitted context, evaluate decisions and retain evidence. The Observer gathers separately permitted readback through Authorized Context; reconciliation compares the attempt and observation with the operation's requirements. Notice the return arrows. An account of what happened is part of the work.

Two parts of the map support the wider system. Verticals provide domain definitions and institutional configuration; AppFactory authors and compiles application content. StilleNacht conducts assurance campaigns that assess specified behavior. Installed applications continue through the runtime controls while the authoring Factory is offline, and assurance findings inform lifecycle decisions. Chapter 2 examines these responsibilities and their credential boundaries.

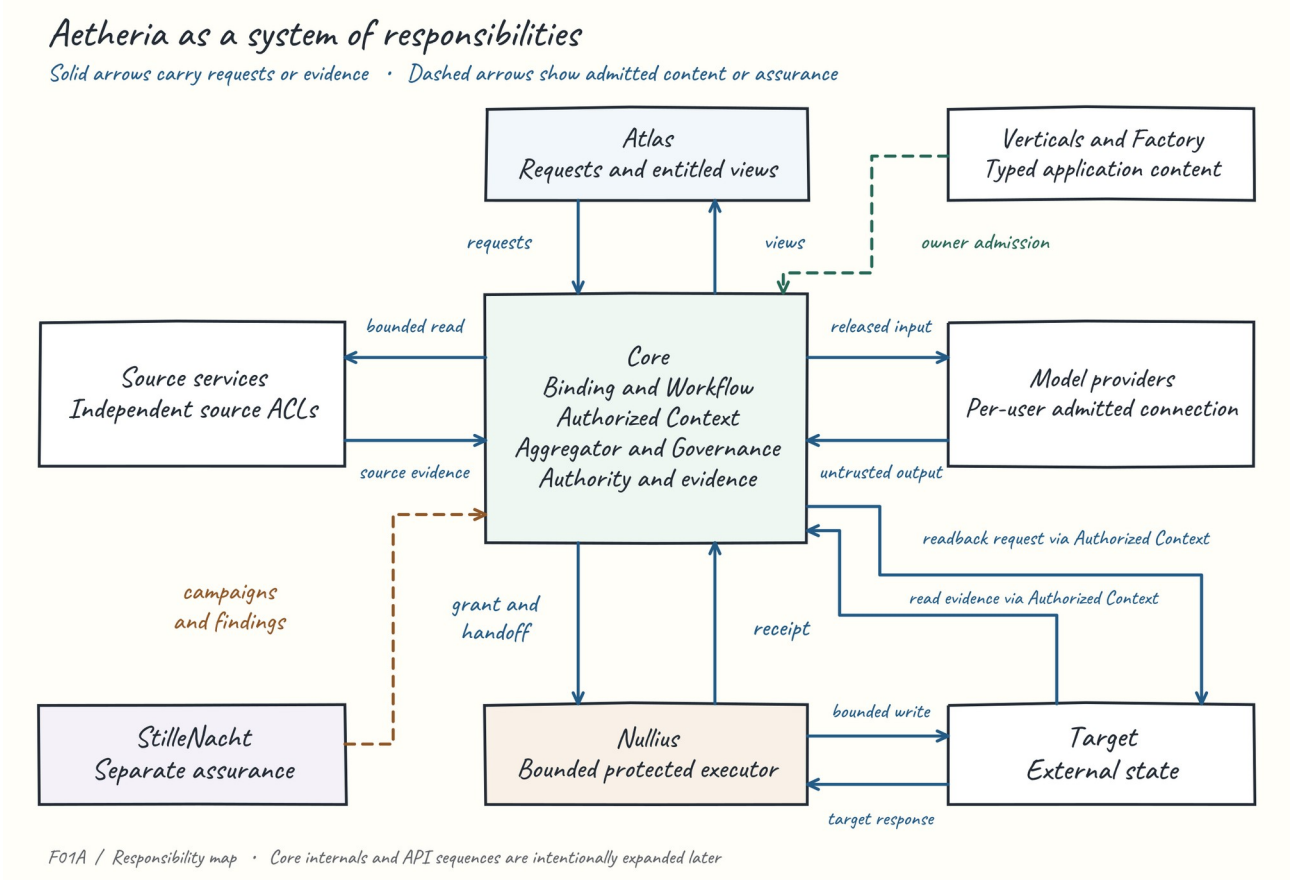


Figure 1.2 Whole system responsibilities. Solid arrows carry the requests and evidence identified by their labels. Dashed arrows show content or assurance dependencies. Core is expanded in Chapter 2 and later component chapters. The map groups responsibilities; it leaves deployment and detailed API order to their own diagrams.

What to retain

The council count needs an entitled source set and a declared method. Jonah's proposed mail needs an exact review, and a lost reply leaves the send outcome open.

- Answer — a supported informational result
- Permission — an institutional decision about an exact action
- Effect — the change attempted at an external target

Exercises

Q1.1 Classify each statement as content, permission or outcome: a draft council summary; an approval of an exact outgoing email; a report that a target accepted the email. What additional evidence would be needed for each stronger claim?

Q1.2 Using Table 1.2, calculate eligible meetings, meetings with matches, total matching occurrences and the proportion of meetings with a match. Explain the treatment of M003 and M004.

Q1.3 After M001 and M002 are scanned, a required page becomes unavailable. What can an answer claim, and what must remain incomplete?

Q1.4 A developer proposes loading all records into the model and filtering its answer afterward. Explain the violated boundary and two disclosures that could occur before the final display.

Q1.5 A mail attempt times out. The interface offers a button labeled Resend safely because no receipt arrived. Evaluate that claim.

Q1.6 Explain why permission to retrieve a case file cannot itself answer whether an officer may conduct a search.

Sources for this chapter

Internal design sources by section. External literature remains cited in the discussion. 1.1: [1], [2], [5], [6], [8], [9], [19]; 1.2: [1], [2], [3], [4], [5], [7], [8], [9], [12], [14], [19]; 1.3: [4], [6]; 1.4: [4], [18]; 1.5: [2], [4], [8], [18]; 1.6: [1], [3], [4], [5], [6], [9], [10], [19], [31]; 1.7: [2], [7], [8], [18], [36], [37]; 1.8: [1], [3], [7], [10], [11], [17], [19], [20], [32], [33].