GRAPHS.SGIT.AI · SECOND EDITION OF THE ARGUMENT

Fractal Semantic Graphs

Meaning Through Connectivity

A node connected to nothing means nothing. What a thing is emerges from the edges traceable from it, and how much you can rely on that meaning is a function of how richly and how independently it is connected.

The first edition argued that from first principles. Since then the argument was built: a document decomposed to the word with every claim anchored to exact bytes, an engine that computes meaning deterministically with its arithmetic in the open, registers where a human corrects the machine and the correction is the training.

This is the edition that argues from first principles **and** from the running system.

Dinis Cruz with a team of AI agents · v0.5.12 · 26 August 2026 · CC BY 4.0 · not a graph database pitch

Fractal Semantic Graphs: Meaning Through Connectivity

Second edition of the argument · graphs.sgit.ai · August 2026

What this book is

A node connected to nothing means nothing. That is not a slogan. It is the first consequence of a claim you can test on real data: **what a thing is emerges from the edges traceable from it, and how much you can rely on that meaning is a function of how richly and how independently it is connected.**

The first edition of this argument made that case from first principles and published it in advance of any system that implemented it. It said so, plainly, in a chapter that listed what did not exist. That honesty is why the argument was worth reading, and it is also why it was incomplete: it argued that meaning could be *computed* rather than declared, and then had nothing to point at that computed it.

Between that edition and this one, the argument was built. A document was decomposed to the word with stable identities and every claim anchored to the exact bytes that carry it. An engine was written that takes a phrase and returns its meaning deterministically, with the arithmetic in the open and the evidence trail clickable in both directions. Registers were created where a human corrects the machine's proposals and the correction *is* the training. Nineteen sites and twenty published vaults now run the grammar in public.

This book is the edition that argues from first principles **and** from the running system. It is not a revision of the first book and it is not a tour of a repository. It is a fresh argument that draws on both.

What you will know at the end that you do not know now

If you read this book start to finish, you will be able to do five things you probably cannot do today.

1. **State the thesis in your own words**, and defend it against the obvious objection that a table would have done.
2. **Apply the edge grammar to a system you already know.** Fifteen verbs, each with a distinct inverse, one banned edge, and a test you can run out loud: if the path does not read as a sentence in the reader's own language, the edges are wrong.
3. **Tell a fractal system from a merely hierarchical one**, using a test that either passes or fails: zoom into any node, and if the zoom needs a new file format, a new validator or a special case, the system is not fractal.

4. **Explain why determinism matters to the argument**, and why an engine whose weights are stated formulas rather than fitted numbers is a different kind of object from a language model, even when it wears the same shape.
5. **Start your own graph tomorrow**, on a system you already work on, in an afternoon, without buying anything.

If you already read the first edition, this one goes further rather than merely again. The grammar chapters are sharpened, not repeated. Everything from chapter six onward is new ground: the fractal claim tested against a system that actually zooms, the projection claim tested by a build gate that rebuilds the source document from the graph and fails unless the bytes match, and four chapters on meaning as a computation with its own worked arithmetic.

The three things this book will not do

It is not a graph database pitch. That sentence is carried verbatim from the source brief this book takes its title from, and it means what it says. The claim is that one grammar is the interface at every boundary, not that things are stored in a graph. The work behind this book uses no graph database, no SPARQL, no Cypher and no RDF layer in its code. Chapter fourteen lists exactly what ships and exactly what does not.

It will not pretend the semantic layer is a product. Most of what follows is design published in advance so that it can be checked. Where something runs, this book says so and gives you a number you can verify. Where nothing runs, it says that too.

It will not hide the interest. This book is published by a project that builds the things it argues for. If the argument is right, that project is more valuable. Chapter fourteen and the participant note below are what we do about it, and they are not enough on their own: hold the elegant chapters to a harder standard than the evidenced ones.

How this book is arranged, and why

Five parts. The title is read backwards, and each part earns one part of it.

> **Part one, the claim.** Meaning through connectivity, from first principles, with no jargon before it is earned. Two chapters. **Part two, semantic.** What makes a graph semantic rather than merely a network: a grammar of verbs, anchors instead of standards, and types that are computed paths rather than applied labels. Three chapters. **Part three, fractal.** What the middle word of the title commits you to, how to falsify it, and what happens at every boundary and every zoom level once you take it seriously. Three chapters. **Part four, computed.** The half the first edition could not write: how a document becomes a graph whose every node is anchored to bytes, what identity means when a document changes, an engine that computes meaning deterministically, and why an explanation is a path rather than a narration. Four chapters. **Part five, in practice.** Where this runs today, what ships and what is argued, and how to build your own first graph tomorrow. Three chapters.

Then a colophon that says what was cut and what remains open, and a reference card: the edge set, the rules and the vocabulary on four pages you can read on a plane with nothing else open.

The editorial choices, recorded

The commission locked the title and left everything else to the writing session. Those choices are recorded here so they can be argued with.

Choice	What was decided, and why
Structure	Five parts, fifteen chapters plus front matter, a colophon and a reference card. The proposed spine in the commission had eight chapters; it was expanded because the four chapters of part four each carry a different mechanism (extraction, identity, computation, explanation) and folding them into one would have made the book's newest material its thinnest.
Order	First principles first, running system last. The alternative (lead with the engine, because it is the novelty) was rejected: the engine is only interesting if you already accept that meaning is a computation over connections, and that case has to be made before the machine appears.
Chapter shape	Every chapter opens with what you will be able to do or see differently after it, and closes with where the live estate demonstrates it. Every chapter carries at least one figure.
Voice	Plain sentences, short words, technical terms spelled out at first use in each chapter, no em-dashes in our own prose. Verbatim quotes keep their original punctuation, always.
Figures	Screenshots are taken from the real pages at version v0.5.11 with the repository's own headless browser harness, and each caption names the page and the version. Structural diagrams are ascii art, which is diffable, reviewable and prints well. No new chart libraries were introduced, and no figure was drawn from imagination.
Evidence	Every attributed position names its chapter, brief, page or release row. Every number is computed from a file in the repository or quoted from a page that was fetched, never recalled. Where the corpus is silent, this book says so.

Who wrote this, and how

> **Written by Dinis Cruz together with a team of AI agents.** The working method is part of the material, so it is stated rather than implied: ideas are recorded as voice memos, transcribed, and developed by agents into structured briefs; the author reviews and corrects them; the corrected briefs become the estate's instruction record; and the chapters are distilled from there. This book in particular was written by an AI agent working from a written commission, reading the corpus directly and anchoring every claim to it. The agent chose the structure, the chapter count, the voice and the figures, and recorded those choices in the table above. Where this book states a position the corpus does not carry, it says so in the paragraph that states it, and you should hold those paragraphs to a lower evidential bar than the sourced ones. The estate's disclosure practice applies here unchanged: the interested party is a visible node. The full participant disclosure is published at `graphs.sgit.ai/v1/about/participant.html`, including the four situations in which the argument of this book is the wrong one.

Licence

A note on reading offline

This book is designed to be finished on a flight. Nothing in it requires a link to be followed. Every figure carries a caption that states its point, so a reader who cannot see the colours still gets the argument. Every number that matters appears in the prose as well as in a figure. Where a live page would make the point better than a paragraph, the paragraph is written anyway.

Fractal Semantic Graphs: Meaning Through Connectivity · graphs.sgit.ai · site version v0.5.11 · 26 August 2026 · CC BY 4.0

Contents

PART ONE

The claim

Meaning through connectivity, from first principles, with no jargon before it is earned.

1 · A node is just a node

After this chapter you will be able to state the book's central claim in your own words, and you will have a test for it that costs nothing: take any label in a system you know, trace what it reaches, and see whether the meaning you assumed survives the trace.

Draw a box. Write **Review** inside it. What have you got?

You have a box with a word in it. You do not have a Review. You have a node that somebody labelled "Review", and that label is a claim by whoever drew the box, not a fact about the world. Nothing inside the box tells you whether that Review took two minutes or two weeks, whether a human did it, whether a regulator would accept it, or whether it is the same kind of thing as the Review your colleague drew in a different box yesterday.

> A node connected to nothing is meaningless. Literally, not rhetorically: there is nothing there to be right or wrong about.

This sounds like a philosophical point. It is an engineering one, and it has an immediate practical consequence. **If you try to make the box mean something by putting more words inside it, you are solving the wrong problem.** Adding `type: "security-review"` and `duration: "2 weeks"` gives you a bigger box with more words in it. Somebody else's box will use different words for the same thing, and the same words for a different thing, and you will discover which at the boundary, in production, when it is expensive.

The two 8080s

Here is the version that makes it concrete, and it is real shipped code rather than a metaphor.

Two variables in a Python program. Both hold the number `8080`.

```
[port = 8080] -typed_as-> [int]

    Traced outwards, it reaches int, and stops. That is the whole graph.


[port = Safe_UInt__Port(8080)] -typed_as-> [Safe_UInt__Port]
                               -extends->  [Safe_UInt]
                               -part_of->  [osbot-utils@3.63.4]

    Traced outwards, it reaches a type that carries a range constraint,
    which belongs to a library, at a pinned version, which has its own
    graph of tests, a source repository, a licence, a maintainer.
```

The difference is not in the value. Both scenarios hold `8080`. The meaning is identical in the developer's head. It is radically different in the graph, and the graph is the thing another system, another team, or an agent has to work from.

Notice what is doing the work in the second case. It is not that somebody wrote a longer description. It is that a chain of edges exists, and each hop constrains what the value can be. The range constraint is not a comment; it is a property of a type that something else in the system enforces. The version pin is not documentation; it resolves to a specific artefact with its own graph. Meaning here is not asserted anywhere. It is *reachable*.

Which gives the sentence this whole book is built on, taken from the foundational essay of the corpus, `library/concepts/v0_4_0__thinking-in-graphs.md`, dated 5 February 2026, and quoted here in its own words:

> **everything is a graph, meaning is not declared but discovered through graph relationships, and confidence in that meaning is proportional to how richly connected a node is to other nodes that provide context.**

And note what follows immediately, because everything else in this book is a consequence of it: **meaning stops being something you assert and becomes something you can compute.**

The first edition of this argument made that claim and then, honestly, had nothing that computed it. Part four of this book is the answer to that. But the claim has to stand on its own first, because an engine that computes the wrong thing is worse than no engine.

Five teams, five processes, one word

This is the best on-ramp in the material and the one to try on a sceptical colleague.

Five organisations each have a thing they call a **Review**.

Who	What their "Review" actually is
A, a team in Tokyo	Two engineers, a checklist, roughly forty minutes, recorded in a ticket
B, an open-source project	Whoever shows up comments on a pull request; merged when someone with commit rights is satisfied
C, a bank in Frankfurt	A three-week formal process with named approvers, an audit trail, and a regulator who may ask for it
D, a startup in Lagos	The chief technology officer reads it on the way to a meeting and says yes
E, a research group in São Paulo	Two anonymous peers, a written response, a revision cycle

The schema-first instinct is to define a Review. So ask the question that instinct never survives: **which of these five gets to be the definition?**

Whichever you pick, four organisations now have to either lie about their process or exclude themselves from your system. That is not a modelling problem. It is a political one, and it does not have a technical solution. Every schema you have ever seen that covers more than one organisation has this problem, and most of them resolve it by having the largest party win.

The graph does not ask anybody to agree. Each organisation keeps its own node with its own edges: who performed it, what was produced, how long it took, what it referenced, what it authorised. Then you ask a specific question, and the answer is a computation over the overlap.

- **"Did somebody other than the author look at this?"** A, B, C and E overlap. D does not.
- **"Is there a record a third party could inspect?"** A, B, C and E. D does not.
- **"Would BaFin accept this?"** C, and possibly nobody else. (BaFin is Germany's federal financial regulator.)

Three properties of those answers are worth naming, because each is something a schema cannot do.

They are not binary. Compatibility is a degree of overlap, not a yes or a no. The schema's answer is always yes or no, which is why so much of the work with a schema is arguing about where to put the line.

They are not symmetric. C's review satisfies B's question. B's does not satisfy C's. A schema has one direction of conformance and therefore cannot express this without inventing a hierarchy that somebody has to defend.

They are purpose-relative. The same two processes are compatible for one question and incompatible for the next. There is no global answer, and asking for one is the mistake.

> Nobody had to agree on anything, change their process, or adopt a shared vocabulary, and the system still told you exactly where they overlap and where they do not.

What happens when you change what a word means

The five Reviews argument is old, in the sense that a philosopher would recognise it. The new part is that you can now watch it happen, mechanically, in a running engine, and count what changes.

The estate holds a small deterministic engine (chapter eleven builds it up properly) whose job is to answer *what does this phrase mean* against one document's world. Ask it about "graphs of graphs" and it answers with a concept, its statement, the quoted sentence in the source that carries it, and the arithmetic behind the score.

Now change one thing. The engine knows that the word "graph" means different things in different worlds: a network graph here, a chart of data in a boardroom, the plot of a function at school, the ruled paper you draw on, the social graph of a platform. Switch the active sense of "graph" from *the network graph* to *a chart of data*, and the engine does not merely give a different answer. It reports which of this document's claims stop applying, computed from the concept labels rather than authored by anybody, and it reports that nothing survives into the attention layer at all.

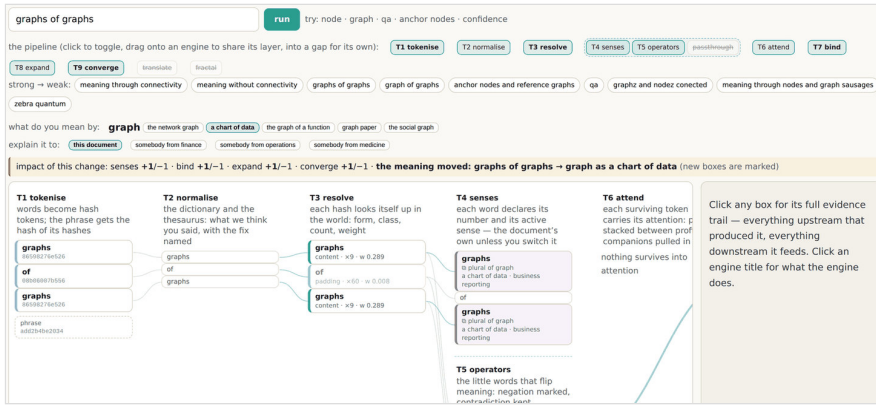

Figure 1.1 · The WCLM at graphs.sgit.ai/v2/wclm/, site version v0.5.11. The prompt is "graphs of graphs". The active sense of the word "graph" has been switched from "the network graph" to "a chart of data". The banner above the layers reports the movement: senses +1/−1 · bind +1/−1 · expand +1/−1 · converge +1/−1 · the meaning moved: graphs of graphs → graph as a chart of data. The attention layer reports "nothing survives into attention". Changing what one word means emptied the rest of the computation.

That is the five Reviews, run as arithmetic on one word. The founder predicted the result before it was built, in a voice memo of 26 August 2026: *"if I go on graphs or graphs and I change my definition of a graph is to, for example, to be a diagram, right, a line diagram or something else, well, then the fractal element will not apply."* It does not. The engine now says so, and says which claims it takes with it.

The lesson is not about the engine. It is that **the disagreement is data**. Two people using the word "graph" differently are not making an error to be corrected by a definition; they are holding two positions whose consequences can be computed and compared. Merging them into one shared definition would erase the finding. Chapter four is about why you should never do that.

"We cannot confirm Z" is a better answer than a guess

If meaning comes from connectivity, then **how well connected something is** is a measure of how much weight it will bear. That gives a ladder, and every assertion in every system you own sits somewhere on it.

```
5 · rich multi-hop connectivity
    many INDEPENDENT paths lead to the same conclusion
    (independence, not count: ten citations of one source are one source)

4 · edges to external references
    a published standard, a legal instrument, a versioned library,
    a URL a third party can fetch and check

3 · edges to anchor nodes
    well-known, well-maintained reference points that others also
    connect to, so two systems can compare notes without merging

2 · edges to typed definitions
    something else in the system constrains what this can be:
    the port that reaches a type that carries a range

1 · a few local edges
    connected to things in the same document or system.
    internally coherent; means nothing outside it

0 · no edges
    a word in a box. nothing can be checked, so nothing
    should be relied on
```

Figure 1.2 · The confidence ladder. Each rung is a claim about connectivity, not about effort or intention.

So the honest output of such a system is not a score. It is three sentences: **we know X, we think Y, we cannot confirm Z**, each with a reason you can trace.

And the remedy for low confidence follows directly from the diagnosis. It is **enrichment, not enforcement**: you add edges, you do not add validation rules. The graph grows; it does not constrain. That distinction is easy to state and hard to hold on to, and it is what separates this from every schema-validation system you have used. When your instinct says "we should require that field", the graph's answer is "connect it to something that already knows".

There is a cost to this and it is worth stating in the chapter that introduces it rather than burying it later. Enrichment does not stop anybody doing anything. If your requirement is that a field must never be null, a validator does that and a graph does not. Chapter fourteen carries the full list of situations where this approach is the wrong one, and this is one of them.

Three of ten pieces of evidence *is* information

Most systems list what they have. An honest one also lists what it lacks.

If you need ten pieces of evidence to support a conclusion and you hold three, that is not a failure state to be hidden behind a progress bar. It is a fact with three uses. It quantifies your confidence. It tells you precisely which seven things to go and get. And it makes the business case for connecting them, because now the missing dots have names.

A register has empty cells, which look like nothing at all. A graph has **unanswered question nodes and unevidenced facts**, which can be counted, queried, assigned to a person and given a date.

The sharpest instance in the corpus is a worked mapping of one EU AI Act provision (Article 26(5), on the deployer's obligations for a high-risk system) onto one concrete deployment. It produced nine question nodes. Five of them were unanswered, and the source brief calls those five *"the actual output of the exercise"*. Chapter thirteen walks that example. The habit it teaches is the one to steal first: when you cannot determine something, emit an explicit node saying so, rather than a plausible value.

> A named absence beats a hidden one. In a graph an absence is a node with a name, an owner and a date, which means it can be counted, argued about, and funded.

The same rule holds one level up, in how a graph is drawn. Every rendered graph in this estate uses three states: green for assurance, amber for exposure, and **ghosted for unanswered**. The third one is the point. An unanswered question is drawn, not omitted. That is not a stylistic preference; it is the visual form of the claim above.

Five ideas, and what they cost

That is the whole of part one's first chapter, and it is deliberately short of jargon: you have not yet met the words *ontology* or *semantic*, because you do not need them.

1. A node alone means nothing.
2. The same value, differently connected, means different things.
3. Nobody has to agree for the overlap to be computable.
4. Confidence is a function of connectivity, and of the independence of the paths.
5. A named absence beats a hidden one.

Every one of these has a price. Ideas one and two mean you have to do the connecting, and edges are work somebody has to perform. Idea three means you give up the comfort of a single agreed definition, which is exactly the comfort most governance processes are built to provide. Idea four means your system will tell you it is not sure, out loud, in front of people who wanted a number. Idea five means your dashboards get worse before they get better, because absences that were invisible become visible and countable.

If those prices sound acceptable, the next chapter is the one that explains what kind of graph this is, because the word is used for at least three different things and only one of them is this book's subject.

Where the live estate demonstrates this. The sense switch in figure 1.1 runs at `graphs.sgit.ai/v2/wclm/`, and the switch is reversible in one click. The confidence ladder governs the rendering of every graph in the estate: open the Risk Graph Explorer at `sgit.ai/demos/vaults/risk-graph-explorer/` and the ghosted edges are the unanswered ones, drawn rather than omitted. The five Reviews come from `library/concepts/v0_4_0__thinking-in-graphs.md`, whose full text is published, frozen and hash-verified at `graphs.sgit.ai/v1/docs/sources/thinking-in-graphs.md`.

2 · Why graphs at all

After this chapter you will be able to place this argument precisely against the three things people mean by the word "graph", against retrieval-augmented generation over graphs, against the Semantic Web, and against property graphs, and you will know the four situations in which this book's argument is the wrong one.

One word first. In this book a **graph** is a network: nodes and the edges that connect them. Never a chart, never a figure. Nothing here plots values on axes. Everything here is about things, and the stated relationships between them.

And a suspicion to start from: **you may already think in graphs without ever having called it that.** Working out what an unfamiliar system is by tracing what it connects to. Trusting a claim because of where it came from rather than how it is worded. Asking "what else breaks if this fails?" That is graph-thinking, and it is common. What is rare is doing it deliberately, with rules, and that is what this book teaches.

So this chapter is written for three readers: the one who already thinks this way and never named it, the one who is not yet convinced, and the one who is convinced of something else. If you have used graphs professionally, the useful half is the second, because this is a third use of graphs and probably not the one you are thinking of.

Three different things people mean by "graph"

Use	The question it answers	What you buy it for
1 · Networks social graphs, dependency graphs, citation networks	Who is connected to whom, and how centrally?	Analysis. Centrality, clustering, shortest path, community detection.
2 · Storage graph databases, triple stores	How do I make joins fast?	Performance. A traversal beats seven table joins.
3 · Semantics this book	**What does this thing mean, and how sure can I be?**	Meaning that survives a boundary: a different team, project, culture, language, or system.

Uses one and two are well served and well understood. This book is about use three, which is why the disclaimer goes at the front and stays there: **not a graph database pitch.** The claim is that one grammar is the interface at every boundary, not that things are stored in a graph. Nothing in this argument depends on which store you use. The work behind this book does not use a graph database at all, and says so in its own architecture notes.

That is worth dwelling on for a moment, because it is the single most common misclassification of this material. A reader who takes it as a storage argument will reasonably ask why not use an established triple store, will find the answer unconvincing, and will stop reading. The answer is that storage is not the subject. The subject is what crosses the seam between two systems that were not designed together.

The argument, in four steps

```
1 · declared meaning is brittle and local
    a schema works perfectly inside the system that defined it.
    cross a boundary (another team, project, culture, language)
    and it either forces conformity or breaks.
                                                  |
                                                  v
2 · the boundary is not an edge case; it is where all the work is
    integration, compliance, procurement, supply chain, regulation,
    multi-team delivery, agents calling other systems: every one of
    these IS a boundary problem.
                                                  |
                                                  v
3 · connectivity survives the boundary because it needs no agreement
    two parties do not need a shared vocabulary to compare notes.
    they need to have connected their own nodes to enough context
    that the overlap can be computed.
                                                  |
                                                  v
4 · and once meaning is computed, it can be checked and versioned
    a judgment in somebody's head cannot be reviewed. a judgment
    expressed as a required path-pattern can be read, disputed,
    versioned and tested against the data.
```

Figure 2.1 · The four-step case. Step one comes from the corpus's foundational essay; steps two to four are this book's reading of it.

Step four is the one that changes what a system is *for*. Judgment does not disappear when you move it into a graph. Somebody still decides that a vulnerability requires an upward path to a risk. What changes is where the decision lives: out of the classifier's head, into a formula that is visible, versioned and arguable. You can now disagree with a classification by pointing at a line, which you could not do before. Chapter five makes that precise.

"What I am describing is not complexity, it is reality"

The most common objection is that this is over-engineering, and that a table would do. Sometimes a table would do. The reply worth quoting is the founder's, from 18 June 2026:

> "what I am describing is not complexity, it is reality. This is the reality of business, the reality of the complex applications we have."

The test is not whether the graph is simpler than a table. It is whether **the question you need answered is expressible in a table**. Three that are not:

Reach. A table lists the permissions an account has. Only a transitive closure over assume-role, pass-role and wildcard edges tells you what those permissions actually *reach*. In the identity and access management worked example, that closure is a node type in its own right, called the agentic union, and for an agent it is the rating floor rather than the nominal grant.

Bidirectionality. "What does this browser extension reach?" and "how could my email be attacked?" are two different tables. They are one graph, walked in two directions. Nobody maintains the second table, which is why nobody can answer the second question.

Propagation of a correction. Mark a claim superseded, then ask which conclusions were resting on it. There is no table shape that answers that. The clearest case is the ten-thousand-hours literature: **242 papers and more than 200,000 supporting citation paths, traced back to a claim that corrections never reached**, because in a pile of documents there is nothing for a correction to attach to. Chapter five turns that into a rule.

Where this sits next to GraphRAG, RDF and property graphs

> **More contested than the rest of the book, and written fresh.** The corpus behind this argument has zero occurrences of "GraphRAG" and zero of "hypergraph". It holds a strong implicit position and never engages the named field. That absence is recorded in the estate's own gap catalogue as gap G11 (the eleventh of the twelve gaps the brief pack listed as things the source corpus could not supply). What follows is a position stated so that it can be argued with, not a claim that somebody already worked it out.

GraphRAG

GraphRAG, meaning graph-based retrieval-augmented generation, shares a real premise with this book: retrieval over structure beats retrieval over a pile of chunks. The difference is what the structure is *for*.

GraphRAG, in its common form, builds a graph in order to retrieve better context to put in a prompt. The graph is scaffolding for a generation step, and the model remains the thing that decides. The position here is stronger and narrower: **knowledge is traversed, not guessed.** Retrieval is a traversal from an intent node to grounded facts with provenance attached, rather than a similarity search returning plausible chunks. The model sits at the edge and *proposes* a graph; a deterministic validator decides whether that proposal is admissible.

That is a genuine disagreement and it has a cost worth stating plainly: **it requires the edges to exist.** Similarity search works on an unstructured corpus today. Traversal does not. Where the graph is thin, this approach has nothing to say, and pretending otherwise would be the exact dishonesty this book stands against.

RDF, OWL and the Semantic Web

RDF (the Resource Description Framework) and OWL (the Web Ontology Language) are the Semantic Web's core standards, and the disagreement here is with a practice, not a goal. It is a respectful disagreement, because that community identified the right problem two decades early: meaning has to travel between systems that were not designed together. That was correct and it is still correct.

The disagreement, in the corpus's own words:

> "They ended up attaching meaning **to nodes** rather than deriving meaning **from edges**. … The node becomes a little document that describes itself. **This is schema-first thinking dressed in graph syntax.**"

Be precise about how narrow that is. It is not a disagreement with RDF as a serialisation. It is not a disagreement with URIs as identifiers, or with shared vocabularies as reference points, since anchor nodes (chapter four) are exactly that. It is with the practice of making a node self-describing, because a self-describing node has smuggled the schema back in.

And note what the position is **not**. It is not anti-RDF at the serialisation layer. The live EU AI Act regulation graph in this family exports RDF/Turtle as a published artefact. The honest statement is **both, at different layers**: RDF is a fine way to hand a graph to somebody else; it is a poor place to put the meaning.

Property graphs

An awkward one to answer honestly, because this book's core rule is that properties do not carry meaning, and the one graph in this project that is *actually running* is a typed property graph: the issue tracker's own configuration, with 12 node types, 10 verb and inverse edge types carrying domain and range constraints, and **71 nodes and 141 edges** across 107 issue files, stored bidirectionally.

The resolution is a distinction: properties are allowed to carry **data**; they are not allowed to carry **meaning**. A property may hold a timestamp. It may not hold the answer to "what kind of thing is this?", because that answer is a query. Chapter five is where that becomes precise, and the shipped issue graph is the smallest working demonstration of it in the estate.

Vector search and embeddings

Worth a paragraph because it is the comparison a reader arrives with in 2026, and because part four of this book builds an engine that deliberately looks like a transformer and deliberately is not one.

An embedding places a token in a space where distance approximates similarity. It is extremely good at finding things that are *like* other things, on corpora nobody has structured, which is most corpora. Its weakness is that the space is learned and opaque: you cannot ask why two things are near each other and get an answer you can check, and you cannot correct a single relationship without retraining or bolting on a rule.

The graph position is the mirror image. It is worse where nobody has done the connecting work, and it is better exactly where you need an answer that somebody can be accountable for, because every relationship is a claim somebody made, with a name and a date on it. Chapter eleven shows what a small engine looks like when every weight is a stated formula over graph inputs rather than a fitted number, and chapter twelve shows what you get in exchange for that restriction. The exchange is real and it goes both ways.

Hypergraphs

No position. The corpus does not contain the word, and this book will not manufacture one to look complete. If you have a case where a binary edge with a named inverse loses something a hyperedge keeps, the estate's public comms board is where to put it.

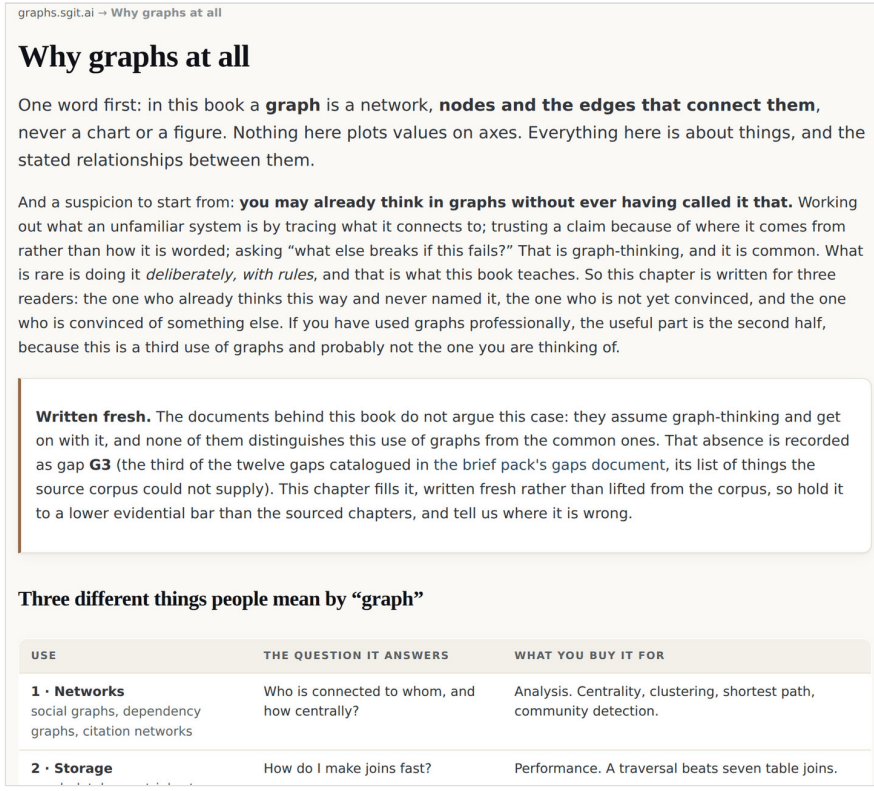
graphs.sgit.ai → Why graphs at all

Why graphs at all

One word first: in this book a **graph** is a network, **nodes and the edges that connect them**, never a chart or a figure. Nothing here plots values on axes. Everything here is about things, and the stated relationships between them.

And a suspicion to start from: **you may already think in graphs without ever having called it that.** Working out what an unfamiliar system is by tracing what it connects to; trusting a claim because of where it comes from rather than how it is worded; asking "what else breaks if this fails?" That is graph-thinking, and it is common. What is rare is doing it *deliberately, with rules*, and that is what this book teaches. So this chapter is written for three readers: the one who already thinks this way and never named it, the one who is not yet convinced, and the one who is convinced of something else. If you have used graphs professionally, the useful part is the second half, because this is a third use of graphs and probably not the one you are thinking of.

Written fresh. The documents behind this book do not argue this case: they assume graph-thinking and get on with it, and none of them distinguishes this use of graphs from the common ones. That absence is recorded as gap **G3** (the third of the twelve gaps catalogued in the brief pack's gaps document, its list of things the source corpus could not supply). This chapter fills it, written fresh rather than lifted from the corpus, so hold it to a lower evidential bar than the sourced chapters, and tell us where it is wrong.

Three different things people mean by "graph"

USE	THE QUESTION IT ANSWERS	WHAT YOU BUY IT FOR
1 · Networks social graphs, dependency graphs, citation networks	Who is connected to whom, and how centrally?	Analysis. Centrality, clustering, shortest path, community detection.
2 · Storage	How do I make joins fast?	Performance. A traversal beats seven table joins.

Figure 2.4 · The positioning as published at graphs.sgit.ai/v1/why-graphs/, site version v0.5.11. The warning box above the table is the point of the page as much as the table is: the sections written fresh say so, and name which gap in the estate's own catalogue they fill, so a reader knows which paragraphs to hold to a lower evidential bar.

Where this approach loses

Four situations where the argument of this book is the wrong one. If you are in one of them, do something else. They are carried here from the estate's participant disclosure rather than softened,

and they belong in chapter two rather than in an appendix, because a reader deciding whether to keep reading deserves them early.

```
1 · everyone already agrees, and always will
    inside one team, one codebase, one jurisdiction, with a stable
    vocabulary and no external party: a schema is simpler, faster,
    and it will catch mistakes this approach lets through.
    no boundary, no benefit.

2 · you need the answer to be enforced, not computed
    enrichment rather than enforcement is a real cost. if your
    requirement is "this field must never be null", a validator does
    that and a graph does not. some systems need a gate, and a gate
    is a schema.

3 · the graph would be empty
    this approach needs edges, and edges are work somebody has to do.
    where nobody has done that work, traversal has nothing to say and
    similarity search will beat it outright.

4 · you want to buy it rather than build it
    the honest state of the semantic layer is DESIGNED, NOT SHIPPED.
    if you need something running next quarter, this book is a set of
    arguments you can use, not a product you can procure.
```

Figure 2.2 · The four losses, from the estate's participant disclosure at graphs.sgit.ai/v1/about/participant.html.

A network of nineteen, as a first piece of evidence

There is one piece of evidence available at the start of the book that costs the reader nothing to check, and it is worth putting here rather than saving for part five.

If the claim is that one grammar survives a boundary, then a family of independent projects, each with its own vocabulary and its own argument, ought to be able to link to each other without merging. As of the fetch made while writing this chapter, the sgit.ai network index lists **nineteen focused sites, eighteen live and one with the repository and subdomain in place but nothing published yet.** Each takes one question further than a section of another site could.

Several of them are this book's own arguments, running in somebody else's domain vocabulary:

Site	Its own one-line claim
`standards.sgit.ai`	Point at the provision, or you are asserting.
`risks.sgit.ai`	You cannot deny a risk. You can only say how long you accept it.
`newsroom.sgit.ai`	The story is a graph. The article is a projection.
`issues-fs.sgit.ai`	The issues are files. The files are a graph.
`twins.sgit.ai`	A digital twin is an interface to reality, not a simulation of it.
`wardley-maps.sgit.ai`	Maps are claims, not pictures.
`pki.sgit.ai`	Good public key repositories existed, and were destroyed.

Figure 2.3 · Seven of the nineteen sites in the sgit.ai network, quoted from sgit.ai/network/index.html as fetched on 26 August 2026. Each line is that site's own statement of its argument, not this book's paraphrase.

None of these sites adopted a shared schema. They share a grammar and a discipline, and they connect through named links. That is the claim of chapter four, arriving early and in public. Chapter thirteen returns to what the network demonstrates and what it does not.

> **Where the live estate demonstrates this.** The three uses of "graph" and the positions above are argued at `graphs.sgit.ai/v1/why-graphs/`, which also carries the gap markers for the parts written fresh. The four losses are at `graphs.sgit.ai/v1/about/participant.html`. The network index is at `sgit.ai/network/index.html`, and the published vault list, with a file count, a size and a commit count for each, is at `sgit.ai/demos/vaults/index.html`.

PART TWO

Semantic

What makes a graph semantic rather than merely a network: a grammar of verbs, anchors instead of standards, and types that are computed paths rather than applied labels.

3 · Every edge is a verb

After this chapter you will have five rules you can apply to your own graph tomorrow, a concrete vocabulary of fifteen verbs to start from, one edge name you will never use again, and a quality check that costs nothing: read your paths aloud.

An edge is not a line. It is a claim, and a claim needs a verb.

This chapter is short on purpose. It is the one to keep open while you are actually drawing a graph, and the one to hand an agent before you ask it to emit one. If a rule here contradicts something you were taught about graph modelling, that is deliberate, and rule four in particular.

> **The one-screen version.** Every edge is a verb. Every verb has a distinct inverse. The generic association edge is banned. If the path does not read as a sentence, the edges are wrong. Rich nodes are good: solve the picture at query time, never by removing relationships. And never render the whole graph: render the result of a query.

1 · Every edge is a verb, stated in both directions

Edge	Its inverse	Reads as
`owned_by`	`owns`	this system is owned by this team · this team owns this system
`gives_rise_to`	`arises_from`	this vulnerability gives rise to this risk · this risk arises from this vulnerability
`backed_by`	`evidences`	this fact is backed by this evidence · this evidence evidences this fact
`grants`	`granted_by`	this role grants this permission · this permission is granted by this role

Both directions get written down, and both get a name worth saying out loud. The test is not "is this technically the reverse?" It is **"would a person in this business say this sentence?"**

And the generic association edge is banned

From a voice memo of 10 June 2026, verbatim:

> "the way I create graphs, they are always a two-way relationship and always to do with verbs. … **You can never have relates-to, because relates-to is meaningless, two things always relate to each other. The more granular the edge, the better the query you can write.**"

The reason is not aesthetic. An edge with no verb carries no constraint, so it cannot narrow a traversal, which means it costs you fan-out and buys you nothing. Every generic edge you add makes every query worse, including the queries you have not written yet. **The granularity of the verb is the precision of the query.**

> **And the project breaks its own rule.** The live issue-tracker configuration in the source repository ships a generic association pair in `link-types.json`, and one edge instance uses it. It is named here rather than quietly fixed, because it is a better teaching moment than the rule is: **the generic edge is what you reach for when you have not yet decided what you mean, and it survives because nothing forces the decision.**

2 · The inverse is not the same edge walked backwards

This is the rule that surprises people, and it is the one that makes traversal tractable at scale.

`owned_by` and `owns` describe the same fact, but they are **different relationships with different fan-out**. A system has one owner. An owner has forty systems. Walking outward from the system is a step; walking outward from the owner is an explosion.

> The inverse of an edge is not the same edge walked backwards; it is a different, meaningful relationship, and that asymmetry is what guarantees monotonic progress toward a peak.

Because each hop filters on edge type *and* direction, fan-out collapses at every step rather than compounding. Traversals converge on natural peaks: a board, an owner, a regulation, a register. The practical consequence is the one that matters when your graph gets big: **seed a query in a thousand places and the paths converge on a handful of peaks. Result size is bounded by the number of peaks, not by the fan-out.**

Making the rule machine-checkable

The first edition stated this rule. The second edition enforces it, and the story of how is worth two paragraphs because it is the shape of every good rule in this estate.

In a memo of 24 August 2026 the founder looked at a graph that had exploded and suspected the cause: the graph was going both ways under one name. The agent proposed enforcing the rule at the build gate immediately. The founder's answer, verbatim from the chat:

> "i think it will better to have a first go at creating those inverse verbs , and then we use our new visualisations (like the schema one) to improve them"

So the register came first and the enforcement came with it. `v2/universe/verbs.json` carries every verb the estate's extraction may assert together with its declared inverse: **seven asserted verbs** (`departs-from ⇄ departed-from-by`, `determines ⇄ determined-by`, `provides ⇄`

`provided-by`, `enables ⇄ enabled-by`, `remedies ⇄ remedied-by`, `licenses ⇄ licensed-by`, `exhibits ⇄ exhibited-by`), **three structural relations** (`about ⇄ subject-of`, `demonstrates ⇄ demonstrated-by`, `contains ⇄ part-of`), and **one symmetric relation** (`derived`), marked as symmetric on purpose so that the exception is visible rather than accidental.

The build then fails on three things: a verb the register does not carry, two verbs sharing an inverse, and an undeclared self-inverse. And the file itself states the rule it enforces: *"every relation is stored one way and declares its unique inverse, so the reverse reading is always derivable and never stored."*

That is the difference between a rule in a style guide and a rule in a system. One is advice. The other cannot be broken without the release stopping.

3 · If the path does not read as a sentence, the edges are wrong

A well-built path is legible without a key:

```
[Risk] <-arises_from- [Vulnerability] -impacts-> [System] -owned_by-> [Entity]
       -has_stakeholder-> [Role] -reports_to-> [Board]

    "This risk arises from this vulnerability, which impacts this system,
     which belongs to this entity, which has this stakeholder, who reports
     to the board."

    Nobody needs the legend. The query is almost like a story.
```

The rule has a second half that is easy to skip: the sentence should read **in the reader's own language and business context**, not in yours. A path that reads beautifully to an architect and means nothing to a regulator is a path with the wrong verbs on it *for that reader*. The remedy is not to rename the edges. It is to add the ones that reader's question needs. From the same 10 June memo:

> "the path should read in English, or not even in English, it should read in the language and the culture and the business context we are talking about"

so that

> "the graph explains itself to whoever is reading it, in their own terms."

This is also the cheapest quality check available to you. **Read your paths aloud.** The bad edges announce themselves, in the same way that a badly named function announces itself when you try to say what it does in one sentence.

4 · Rich nodes are good. Build wide, find the few, then flip

Everyone who has done this has produced the blob: the hairball diagram that shows everything and therefore nothing. The usual response is to prune: fewer relationships, cleaner picture.

> "I see a lot of people get into semantic graphs, get excited, and **arrive at the big blob**… The weird problem is **a race to the bottom, where you start not wanting a lot of relationships because they make the graph more complicated.**"

countered, in the same memo, by:

> "**the more rich a node is, the more connections it has, the better.**"

The blob is a **rendering** failure being mistaken for a **modelling** failure. Pruning solves the picture by destroying the asset. Solve it at query time instead:

```
1 · GO WIDE
    a first pass captures the universe around your subject.
    do not economise. nodes and edges are close to free, and some
    nodes exist only to give a later query something to anchor on.
                        |
                        v
2 · FIND THE FEW
    run the question. out of the universe, a handful of nodes are
    relevant to it.
                        |
                        v
3 · FLIP
    re-root the query at those few and walk out again.
    this is the move that makes big graphs usable, and it is why
    the graph being large is not a problem to be managed but the
    condition that makes the flip worth doing.
```

Figure 3.1 · Build wide, find the few, flip.

> Never render the whole graph. Render the result of a query.

Two ceilings are worth carrying in your head, because they decide which tool you reach for.

Ceiling	Number	What it means in practice
Mermaid readability	**~50 nodes**	Mermaid (a text-to-diagram language) is the print step: text, diffable, committable, reviewable in a pull request. Past fifty it is unreadable.
Visualisation legibility	**~300 to 400 nodes**	An interactive canvas is the exploration step. Past this, a human is looking at texture, not information.

"A diagram of everything is rarely useful; one node with its neighbours is always readable." That rule governs this book's own figures, which is why you will not find a hairball in it.

5 · Link to schema.org; do not become schema.org

An **anchor node** is well-connected, well-maintained, well-known, and has **no special authority**. It is a meeting point, not a standard.

The wrong move is to declare `I am a schema:Review`. That is a conformance claim, it is all-or-nothing, and it is usually a lie by the second field. The right move is a granular, honest, disputable edge:

```
[our document_findings step] -similar_to-> [schema:reviewBody]

    "Our document_findings step is similar to what schema.org calls
     reviewBody."

    Partial. Traversable. Arguable. And the part that matters: a third party can
    add this edge without touching either node.
```

Four properties fall out, and the fourth is the one that matters organisationally: **the mapping is a first-class object that somebody else can own.** You do not need the vocabulary's permission, and the vocabulary does not need yours. Chapter four is entirely about the consequences of that.

The edge set

A concrete, versioned, public vocabulary to start from. Fifteen verbs, cited in the appendix of the 28 July 2026 architecture brief as the established set.

Edge	Inverse	Reads as	Where the inverse comes from
`connected_to`	`connected_to`	A is connected to B	symmetric
`observed_on`	`bears_observation`	this evidence was observed on this system	proposed
`backed_by`	`evidences`	this fact is backed by this evidence	proposed
`measured_by`	`measures`	this fact is measured by this measure	proposed
`grants`	`granted_by`	this role grants this capability	proposed
`reaches`	`reachable_from`	this grant reaches this asset	proposed
`enables`	`enabled_by`	this capability enables this action	proposed
`exposes`	`exposed_by`	this fact exposes this blast radius	proposed
`gives_rise_to`	`arises_from`	this vulnerability gives rise to this risk	in the corpus
`protected_by`	`protects`	this asset is protected by this control	proposed
`conditional_on`	`conditions`	this control is conditional on this fact	proposed
`defeated_by`	`defeats`	this control is defeated by this attack	proposed
`owned_by`	`owns`	this system is owned by this role	in the corpus
`accepted_by`	`accepted`	this risk is accepted by this role	proposed
`underwritten_by`	`underwrites`	this acceptance is underwritten by this role	proposed

Read the last column. The fifteen edge names are quoted from the corpus and are load-bearing. Several of the *inverse* names are not: where the corpus does not supply one, this book proposes it and marks it. Proposed names are a starting point for disagreement, not a standard. If this book becomes the place people cite for that vocabulary, the distinction between quoted and proposed has to survive, or we will have done the thing we are warning about.

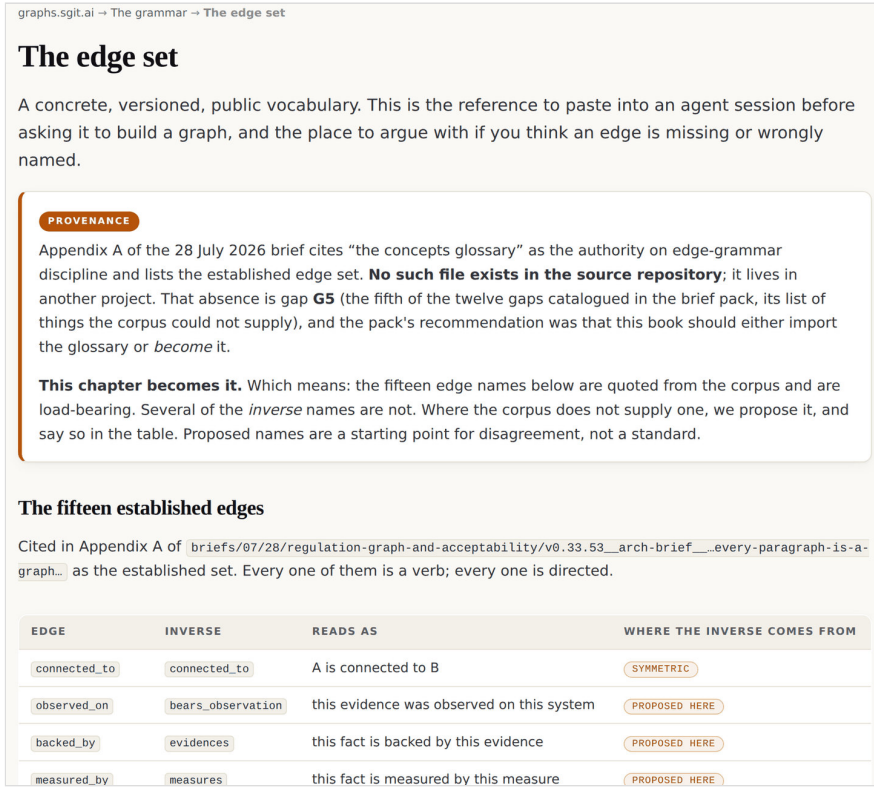

graphs.sgit.ai → The grammar → **The edge set**

The edge set

A concrete, versioned, public vocabulary. This is the reference to paste into an agent session before asking it to build a graph, and the place to argue with if you think an edge is missing or wrongly named.

PROVENANCE

Appendix A of the 28 July 2026 brief cites “the concepts glossary” as the authority on edge-grammar discipline and lists the established edge set. **No such file exists in the source repository**; it lives in another project. That absence is gap **G5** (the fifth of the twelve gaps catalogued in the brief pack, its list of things the corpus could not supply), and the pack's recommendation was that this book should either import the glossary or *become* it.

This chapter becomes it. Which means: the fifteen edge names below are quoted from the corpus and are load-bearing. Several of the *inverse* names are not. Where the corpus does not supply one, we propose it, and say so in the table. Proposed names are a starting point for disagreement, not a standard.

The fifteen established edges

Cited in Appendix A of `briefs/07/28/regulation-graph-and-acceptability/v0.33.53__arch-brief__…every-paragraph-is-a-graph…` as the established set. Every one of them is a verb; every one is directed.

EDGE	INVERSE	READS AS	WHERE THE INVERSE COMES FROM
`connected_to`	`connected_to`	A is connected to B	SYMMETRIC
`observed_on`	`bears_observation`	this evidence was observed on this system	PROPOSED HERE
`backed_by`	`evidences`	this fact is backed by this evidence	PROPOSED HERE
`measured_by`	`measures`	this fact is measured by this measure	PROPOSED HERE

Figure 3.2 · The edge set as published at graphs.sgit.ai/v1/grammar/edge-set.html, site version v0.5.11. Every row marks whether its inverse is quoted from the corpus or proposed by the book, which is the same provenance discipline the book asks of its readers.

> **One row is in tension with this chapter's own rule.** `connected_to` is the only symmetric edge in the set, and a symmetric edge with a broad verb sits uncomfortably close to the one that is banned. It survives because in the graphs that use it, it means something specific (physically or logically attached) rather than "associated somehow". Treat it as a last resort: wherever you can name what kind of connection it is, name it, and the query gets better.

Also used, and deliberately not folded in

Two more appear in the worked graphs and are listed separately rather than folded into the fifteen, because folding them in would quietly enlarge a set somebody else cited: `impairs ⇄ impaired_by` (this risk impairs this asset) and `emits ⇄ emitted_by` (this control emits this detection signal).

Rules for extending the set

1. **A new edge needs a sentence.** If you cannot write "A ⟨verb⟩ B" and have a person in that business say it out loud, it is not an edge yet.

2. **And its inverse needs a different sentence.** If the inverse is just the same sentence read backwards, you have one relationship where you thought you had two. Check whether the direction you chose is the one with lower fan-out.
3. **Domain and range, stated.** Which node types may sit at each end. This is what makes a malformed graph detectable rather than merely wrong.
4. **No generic association edge, ever.** Not even temporarily. Temporarily is how the one in `link-types.json` got there.
5. **Prefer adding an edge to adding a validation rule.** Enrichment, not enforcement.

> **Where the live estate demonstrates this.** The edge set is published at `graphs.sgit.ai/v1/grammar/edge-set.html` and the five rules at `graphs.sgit.ai/v1/grammar/`. The machine-checked verb register is `v2/universe/verbs.json`, and the schema view that was built to improve it (one node per family, one edge per typed relation, both directions labelled with their counts) is a toggle on the graph page at `graphs.sgit.ai/v2/universe/thinking-in-graphs.graph.html`. On first look at the pilot document that view showed nine node types and twenty-four typed relations, with `about ⇄ subject-of` used forty times while the seven asserted verbs are used once or twice each, which is a finding about the extraction rather than about the viewer.

4 · Anchors, not standards

After this chapter you will know why merging two vocabularies is a destructive operation, what to build instead, and what a bridge between two worlds looks like when it is a first-class object somebody owns rather than a mapping table nobody maintains.

Given two ontologies covering the same ground, the instinct is to merge them into one. An ontology, in the plain phrasing this book prefers, is **the kinds of thing that exist and how they may connect**. Two of them describing the same domain look like duplication, and duplication looks like a problem.

The position here is that merging is a destructive operation, and what it destroys is the finding.

> "ontologies are not folded into a single shared definition, **because that erases the disagreement**, they are kept intact and connected through anchor nodes… which is **how meaning actually travels across languages, cultures, biases, and political agendas**, by maintaining translations between definitions that each side still owns."

Read the last clause twice. *Each side still owns.* That is not a courtesy. It is the property that makes the arrangement survive a change of personnel, a reorganisation, or a falling-out.

The three layers

The construction has three layers, and the separation between them is the whole design.

```
3 · DECLARED BRIDGES
    explicit edges connecting formulas at specific points:
    "our 'material' corresponds to their 'reportable' under these
     conditions".
    owned by whoever declared them. revisable without renegotiating
    anything. wrong bridges are removable without touching either side.
                                ^
                                |
2 · PER-PARTY FORMULAS
    each party classifies the shared nodes with its OWN rules.
    the chief information security officer's definition of a critical
    risk, the chief financial officer's, the regulator's. three
    different answers over one set of facts, each internally
    consistent, each inspectable.
                                ^
                                |
1 · SHARED FACTS, OWNED BY NOBODY
    the factual graph. this account exists. this bucket is public.
    this provision is in force. nobody has to agree about what any of
    it MEANS in order to agree that it is the case.
```

Figure 4.1 · The three-layer arrangement. Layer one is shared, layer two is owned, layer three is negotiated one bridge at a time.

> Parties can disagree about meaning while still agreeing about facts, which is the only stable basis for working together.

The founder's version is shorter: *"I always err on the side of understanding versus a standardized schema… instead of folding it, you make it compatible, which is why you need an ontology of ontologies."*

Three consequences fall out of the separation, and they are the reason to build it this way rather than as a shared model with per-party extensions.

A disagreement becomes localised. When two parties disagree, the disagreement is a specific bridge, or the absence of one, at a named point. You can look at it. You can date it. You can assign it. In a merged model a disagreement is diffuse: it shows up as a field that means something different depending on who filled it in, and nobody can point at it.

A correction propagates without renegotiation. Change a formula in layer two and every answer that used it moves. Nobody else's formula changes. In a merged model a correction is a schema change, which means a meeting.

A third party can contribute without permission. A bridge is an edge between two nodes that neither party owns. Somebody outside both organisations can assert one, and be wrong, and be corrected, without either organisation touching its own model. This is the single biggest practical difference and it is easy to miss because nothing about it is technical.

The anchor node, precisely

An anchor node is well-connected, well-maintained, well-known, and has **no special authority**. Chapter three introduced it as a rule about edges. Here is what it is as an object.

An anchor is a meeting point that several parties independently point at. `schema.org` is one. A published concept scheme such as EuroVoc or Wikidata is one. So is a named regulation, a named standard, a named library at a pinned version, and, inside a company, the chart of accounts.

Three things an anchor is not:

Not an authority. You point at it; you do not become it. The difference between `similar_to schema:reviewBody` and `is a schema:Review` is the difference between a disputable claim and a conformance assertion. The first can be partial and can be argued with. The second is all-or-nothing and is usually false by the second field.

Not a merge target. Two parties both pointing at an anchor have not merged. They have each declared a relationship to a third thing, which is precisely what lets them compute an overlap without agreeing.

Not necessarily well-designed. An anchor earns its role by being *pointed at*, not by being good. A widely-used vocabulary with known flaws is a better anchor than an elegant one nobody references, because the whole value is in the convergence.

Partial mapping is the normal case, not a defect. And because nodes cost almost nothing, some exist purely to anchor a query, which matters most as soon as you are working across languages and cultures, where the anchor is often the only thing two sides share.

A nuance survives translation because it was never stored in a word

The unit of meaning is a **concept**, not a word. A concept is language-independent: it carries one preferred label per language plus alternates, and it relates to other concepts as broader, narrower, or related. A **term** is how one language happens to express it.

Once meaning lives in the concept rather than the term, a whole class of failure disappears by construction. You are no longer translating word to word and hoping the connotation survives.

The unexpected result is the better story. Rendering a set of concepts into Portuguese produced a bad Portuguese label, and the diagnosis was that **the English word was wrong**. Naming a concept in a second language forces a decision that the source language let you avoid. Where two languages' induced graphs diverge, that divergence is either an error or a genuine lexical gap, and either way **it is a finding**, not noise to be smoothed away.

This is the same shape as everything else in the chapter: the disagreement is the data.

Bridges you can watch being used

The first edition argued for bridges and had none running. The second edition has two kinds, both authored, both reviewable, both wired into a working engine. They are worth looking at because between them they show the two directions a bridge can point.

Senses: what your word means in other worlds

A **senses register** holds, for a word, the meaning this document gives it and three to five meanings other industries give it. It is a small file, authored by an agent and corrected by a human, and it currently covers four words:

Word	Senses held
`graph`	the network graph (this document) · a chart of data (business reporting) · the graph of a function (school mathematics) · graph paper (stationery) · the social graph (social platforms)
`node`	a point in a graph (this document) · Node.js the runtime (software tooling) · a lymph node (medicine) · a stem node (botany) · a host in a cluster
`fractal`	this document's sense · the mathematical sense · the artistic sense · the fractal antenna · the figurative sense
`task`	this document's sense · project management · the operating system sense · a household chore

The document's own sense is always first, and it is always the default. Switching away from it is explicit, visible and reversible, which is the register's most important property: it does not decide what you meant, it asks.

The `task` row is there for a reason. The corpus's own foundational essay uses "task" as its example of a word whose meaning changes with culture, so the register carrying it is the corpus's argument applied to the corpus.

Analogies: how your concept is said in their world

The **analogies register** points the other way. Instead of asking what your word means elsewhere, it asks how *your concept* is said in the listener's world, and it carries the reason.

It came from a memo of 26 August 2026, asking for it by name:

> "it's interesting because then what we need to do is we need to find analogies, which actually is something we haven't talked about. But you know, analogies or or or equivalencies, right? So, how do we then connect two nodes? Sorry, two concepts. Two concepts that may they might come from different places, but we need to use it to reflect that. So, a good example of I don't know if you talk about graphs of graphs. Maybe for the financial, we need to talk about, you know, some type of data or spreadsheets of spreadsheets, right? Because actually, in the financial world, we do have spreadsheets of spreadsheets of spreadsheets of spreadsheets, right?"

The register that answers it holds **sixteen mappings across three audiences**: six for finance, five for operations, five for medicine. Each entry names one of this document's concepts, the audience's own concept, and the why. Four of them, quoted from the file:

This document's concept	Said to somebody from finance	Why
graphs of graphs	spreadsheets of spreadsheets	"finance nests workbooks inside workbooks inside consolidation packs, the same structure at every zoom, and everyone in the room has lived it"
node	a cell with a name	"a named cell means nothing alone; its formula references, what it reads, what reads it, are what make it matter"
meaning through connectivity	the consolidation trail	"a number in the group accounts IS its roll-up: subsidiary lines, eliminations, adjustments, remove the trail and the number is just ink"
anchor node	the chart of accounts	"the shared reference everyone maps their local codes to, so two subsidiaries can disagree locally and still consolidate"

Figure 4.2 · Four of the sixteen entries in `v2/wclm/analogies.json`, quoted verbatim including their stated reasons.

Notice what the last row does. The analogy for anchor node **is an anchor node**: the chart of accounts is exactly a well-known reference point that several parties map their local codes to, with no authority of its own, so that two subsidiaries can disagree locally and still consolidate. The bridge and the thing it bridges to are the same shape, which is a small piece of evidence that the concept is real rather than invented for this book.

And the register is honest where it is empty. When an audience has no analogy authored for a concept, the engine says "no analogy authored yet" rather than improvising one. That is the same rule as the ghosted edge in chapter one: the absence is drawn.

Both registers are training data, and the training is editing

The thing to take from these two files is not their content. It is what changing them does.

Both are **authored inputs**, marked as agent-proposed and human-reviewed, and both are read by a running engine. Correcting an analogy is not filing a bug against a model. It is editing a file, after which the engine's answer changes, deterministically, and the change is visible in the diff. The founder's own framing, from the same memo:

> "the corrections is where it gets interesting because in principle the correction should be that we need better meaning, we need better objects, we need we're missing nodes, or we have to add some corrections to it, actually manually or like let's say manual overrides, which is kind of what again the learning the pre-trained does to an LLM"

Chapter eleven takes that seriously and shows what it means for a system's weights to be stated formulas over edited graph inputs rather than fitted numbers.

Enrichment, not enforcement

The rule that ties the chapter together, and the one most often abandoned under pressure.

The remedy for low confidence is more edges, never more validation rules. The graph grows; it does not constrain.

The reason is structural rather than stylistic. A validation rule works by rejecting something, which means somebody whose data is rejected has two choices: change their reality, or misrepresent it. In practice they misrepresent it, because their reality is not yours to change. The schema-first failure mode is therefore *invisible*: a conforming record that is false. The graph-first failure mode is *visible*: a sparse region you can point at and count.

	Schema-first	Graph-first
Meaning is	declared, in advance, centrally	discovered, at query time, locally
Disagreement is	a conflict to resolve before you start	data, and often the most useful data you have
Crossing a boundary	forces conformity, or breaks	computes an overlap, which may be partial
Low confidence is fixed by	more validation rules	more edges
A third party can	request a schema change	add a mapping edge without touching either node
Failure mode	everyone lies about their process to fit the schema	the graph is thin where nobody did the work

That asymmetry of failure modes is most of the argument. It is not that the graph is more accurate. It is that when the graph is wrong, you can see where.

> **And the honest cost.** Enrichment does not stop anything. There are systems that need a gate, and a gate is a schema. The right answer in a real organisation is usually both, at different layers: a schema at the point where you control both ends and need enforcement, and a graph at the boundary where you control one end and need understanding. Saying "graph everywhere" would be the same mistake as "schema everywhere", made in the other direction.

> **Where the live estate demonstrates this.** The senses register is `v2/wclm/senses.json` and the analogies register is `v2/wclm/analogies.json`; both are readable raw and rendered through the operator explorer at `graphs.sgit.ai/v2/wclm/operators/`. The engine that consumes them is at `graphs.sgit.ai/v2/wclm/`: the sense picker appears for any prompt word the register knows, and the audience picker restates the answer in the listener's own concept with the reason carried. The three-layer arrangement is argued at `graphs.sgit.ai/v1/depth/#ontologies`.

5 · A type is a path, not a label

After this chapter you will be able to replace the sentence "somebody classified this as a vulnerability" with a formula that can be read, versioned and disagreed with, and you will know why a superseded claim must never be deleted.

Here is the strongest single sentence in this material for a technical audience.

> The content of the node does not decide its type; its paths do. Two nodes with identical text can be different types because their edges differ.

A node type stops being a label somebody applied and becomes a **required pattern of typed, directed paths** that a node either matches or does not. Written out:

```
[Fact] := a node with a downward -backed_by-> [Evidence]

    A claim with nothing under it is not a fact. It is an assertion:
    a different node type, and one worth being able to count.


[Vulnerability] := a [Fact] that also has an upward -gives_rise_to-> [Risk]

    The same public bucket is a Fact on Monday and a Vulnerability on
    Tuesday, not because anything about the bucket changed, but because
    somebody connected it to a risk.
```

Every security practitioner recognises the problem this dissolves. A scanner asserts a finding. The finding is a label. The label is wrong for this deployment. You triage it by hand. The triage lives in somebody's head and is lost when they leave. Here **the triage is the formula**, and the formula is versioned.

Judgment does not disappear. That is the usual objection and it deserves a direct answer. Somebody still decides that a vulnerability requires an upward path to a risk. What changes is where that decision lives: out of the classifier's head, into a formula that is visible, versioned, inspectable and arguable. You can now disagree with a classification by pointing at a line, which you could not do before.

The worked instance in the corpus runs six layers, roughly thirty-one node types, twenty edge types with named inverses (forty readings) and seven formulas, over cloud identity and access management configuration. Its standout type is `AuthorizationClosure`: the transitive union of every grant reachable over assume-role, pass-role and wildcard edges, called *the agentic union*. For an agent, that closure is the rating floor, not the nominal grant. It is a node type that only exists because the graph can compute it, which is the clearest possible demonstration that a type can be a computation.

The grounding ladder

One worked formula set, and the most reusable thing in this chapter.

```
                UPWARD IMPLIES
        what does it mean, and why does it matter?
        each step up is an interpretation somebody is accountable for

              [Risk]
                 ^
                 | gives_rise_to
              [Vulnerability]
                 ^
                 | gives_rise_to
              [Fact]  <----- you are here
                 |
                 | backed_by
                 v
              [Evidence]
                 |
                 | measured_by
                 v
              [Measure]

                   DOWNWARD GROUNDS
              is it real? each step down asks for
              something more checkable than the last
```

Figure 5.1 · The grounding ladder. Downward asks for checkability; upward asks for consequence. The two directions have different owners and different failure modes.

Two details in the source brief are easy to miss and are exactly the parts that make it usable.

Measure is not the floor. The true floor is the last node where going deeper would neither improve observability nor change a decision. That is a judgment, it is stated as one, and it is the reason the ladder terminates instead of receding forever. Without that stop rule you get an infinite regress and a project that never ships.

It is explicitly *one* formula among possible others. The brief says so. A ladder presenting itself as *the* ladder would be the schema-first move at one remove, and this book would have no standing to object to it.

The support state: what a document says, and how hard it says it

The first edition stopped there. The second edition ships something that the ladder implies but never quite states: **the evidence state is itself data, recorded per claim, by the reader who extracted it.**

When a document in this estate is turned into a graph (chapter nine builds the pipeline properly), every claim node carries a `support` value which is the *document's own* evidence state, not a truth judgment:

Support state	What it records
`demonstrated`	backed by a worked example inside the text
`argued`	reasoning given
`declared`	stated without support

The pilot document, the corpus's own foundational essay, yields **27 claims: 7 demonstrated, 15 argued, 5 declared.** Those five are not an embarrassment. They are a worklist. Anybody can now ask, of the document that founded this whole argument, which of its claims it merely asserts, and get a list of five with the quoted sentence beside each.

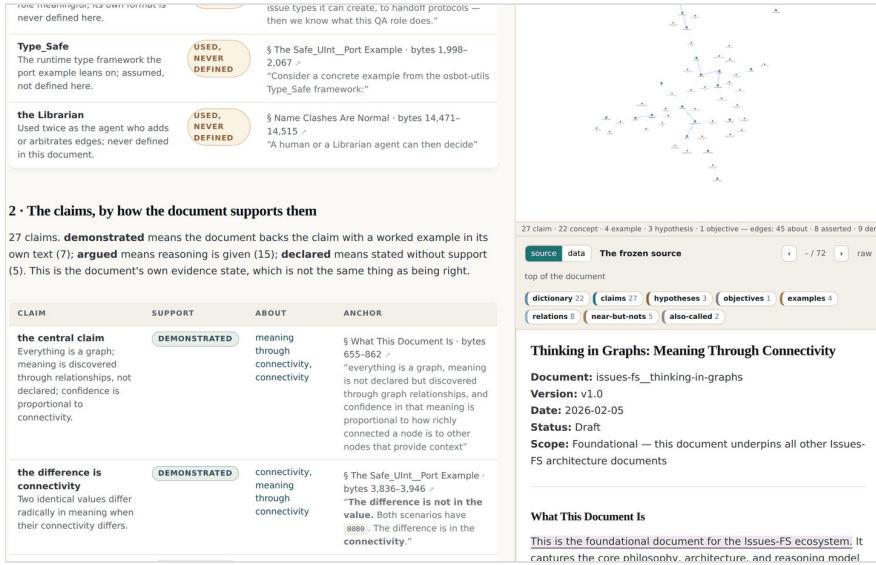

Figure 5.3 · The claims table at graphs.sgit.ai/v2/universe/thinking-in-graphs.html, site version v0.5.11. Every claim carries its support state, the concepts it is about, and the anchor: a named section, a byte range, and the quoted sentence. Above it, three concepts marked USED, NEVER DEFINED, which is the named-absence rule applied to a vocabulary.

The same discipline reached the engine. Every ranked answer the estate's meaning engine returns now declares its **anchoring**: a quoted fact in a named section of the document, a stated claim asserted but not quoted, an authored pack term, or a chosen sense. That is the confidence ladder of chapter one applied to a machine's own output, and it means the answer carries its own rung.

> If a system can tell you how strongly its own sources support what it just told you, you have a different kind of object from one that only tells you the answer. The difference is not accuracy. It is that you can decide how much weight to put on it.

Supersede, never delete

A superseded claim is marked from a date. It is not removed. Removing it destroys the thing you most need: the record that something once rested on it.

Because the claim is still there and still connected, the graph can answer the question a pile of documents cannot: **which conclusions were resting on this?**

The case that makes it vivid is the ten-thousand-hours literature: **242 papers, more than 200,000 supporting citation paths**, all leading back to a claim that later corrections never reached, because in a citation network there is nothing for a correction to attach to. Every one of those 200,000 paths is still, formally, intact. That is not a failure of diligence. It is a structural property of a system where the only edge type is "cites".

Three companions to the rule, all from the same body of work.

Attach, never mutate. Contributions arrive as subgraphs attached to an author-confirmed spine. A bad contribution is discarded rather than repaired, which is what makes *abundance* a feature instead of a risk. You can accept a hundred evidence packs because accepting one costs nothing you cannot undo. This is worth pausing on: it is the property that lets a system take contributions from people it does not trust.

Weight by independence, not by count. Ten citations of one source are one source. This is the rule the ten-thousand-hours network violates at scale, and it is the difference between a confidence number that means something and one that measures popularity.

An index is not a source. Pointer nodes and assertion nodes are structurally distinct. A pointer can be wrong without being dishonest, it is regenerable, it needs no attribution apparatus, and it is therefore **safe to prune**. That is the one place in this book where pruning is allowed, and the distinction is worth building into your node types on day one, because retrofitting it means auditing everything.

A finding that is arithmetic

The best thing a formula can produce is a finding nobody can argue with, and the cleanest example in the corpus is three lines long.

```
fact:      this system retains logs for 30 days
           (backed_by: the configuration export)

provision: Article 26(6) requires a SIX-MONTH minimum
           (backed_by: the official text, hash-verified)

-----------------------------------------------------------
computed:  30 days < 6 months  =>  [Vulnerability]
```

Figure 5.2 · A breach computed rather than asserted, from the Article 26(5) worked example. The source brief calls it "the most defensible finding in the graph".

There is nothing to argue about except the two inputs, and both are checkable. Compare it with the ordinary shape of a compliance finding, which is an experienced person's opinion about whether a control is adequate. That opinion may well be better than the arithmetic. It is not *defensible* in the same way, and when it is challenged the challenge is about the person.

This is the pay-off of the whole chapter. Once a type is a path and a formula is a file, a class of finding moves from the most arguable thing in the report to the least.

What this costs

Three costs, stated because a chapter this pleased with itself should carry them.

Formulas are code, and code rots. A formula that made sense when it was written keeps running after the world changes. The graph does not know that. Somebody has to review formulas the way somebody reviews code, and this estate has not yet built the loop that forces it.

A path-pattern is harder to explain than a label. "This is a vulnerability" fits in a report. "This node matches the vulnerability formula because it has a downward edge to evidence and an upward edge to a risk owned by the chief financial officer" does not fit in a report, it fits in a rendering. If you cannot render it, people will go back to labels.

The upward direction has no natural stop. Downward terminates at the judgment floor. Upward, every fact implies a risk which implies a bigger risk, and the discipline that stops it is entirely social: somebody has to accept the risk, and the acceptance is an edge with a name on it. Chapter thirteen shows what happens when that acceptance is missing, which is the sharpest single finding in the corpus.

> **Where the live estate demonstrates this.** Node type formulas and the grounding ladder are argued at `graphs.sgit.ai/v1/depth/#formulas` and `#ladder`. The support states are data: open the extraction for the pilot document at `graphs.sgit.ai/v2/universe/thinking-in-graphs.html`, where the claims table lists all 27 with their support state and the quoted sentence that carries each. The arithmetic finding is walked at `graphs.sgit.ai/v1/examples/article-26-5.html`.

PART THREE

Fractal

What the middle word of the title commits you to, how to falsify it, and what happens at every boundary and every zoom level once you take it seriously.

6 · Fractal is a testable claim

After this chapter you will be able to tell a fractal system from a merely hierarchical one with a test that either passes or fails, and you will have seen the test applied to this estate's own work, including where it fails.

"Graphs of graphs of graphs" sounds like a flourish. It is meant literally, and it commits you to four specific things.

Claim	What it commits you to
Self-similarity	The same node and edge grammar at every altitude. A property, a paragraph, a person, a national estate: same rules.
Scale invariance	One validator, one query engine, one provenance rule. Not a family of them per level.
Composition	Graphs combine into graphs without an adapter layer. Risk registers of risk registers.
Recursion	Zoom into any node and it expands into a graph obeying identical rules, **with no new format and no special case**.

That last clause is the falsifiable part, and it is how you check whether a system is fractal or merely hierarchical.

> **The zoom test.** Zoom into any node. If the zoom needs a different file format, a different validator, or a special case, the claim is false. It is a testable property, not a description of a feeling.

Almost every system that calls itself hierarchical passes a weaker test: it has levels, and the levels nest. That is not the same thing. A folder tree has levels. A file inside a folder is not a folder, and you cannot apply folder operations to it. That is hierarchy. Fractality is the stronger claim that the operations at one level are the operations at every level.

What fractality buys you is that the system has **no natural stopping point and no integration tax**. From the corpus: *"there might be an article that is so meaty that it requires its own ontology and taxonomy, and that's the power of the fractal element."* The graph starts wherever the work is (*"it is kind of like a Lego structure where one feeds to the other"*) and grows outward from there, which is also why it does not matter where you start.

Applying the test to this estate

The first edition of this argument stated the four commitments and could not test them, because nothing in it zoomed. The second edition zooms twice, deliberately, at two different scales, and the

test can now be run. This section runs it and reports what it returns, including the failures, because a chapter about a falsifiable claim that never reports a falsification is not doing its job.

Zoom one: a document, down to the word

The first zoom came from a memo of 26 August 2026:

> "I should be able to start from the document and keep expanding it, like you know, bit by bit by bit by bit. Let's say on a tree structure, and I should be able to expand it all the way to the paragraph. In fact, all the way to the word, so the way to think about this is kind of like an AST, the abstract syntax tree, but you know, for now, very driven by the by the content itself."

The build that answered it gives the pilot document a **core graph**, which is a named ladder of node kinds:

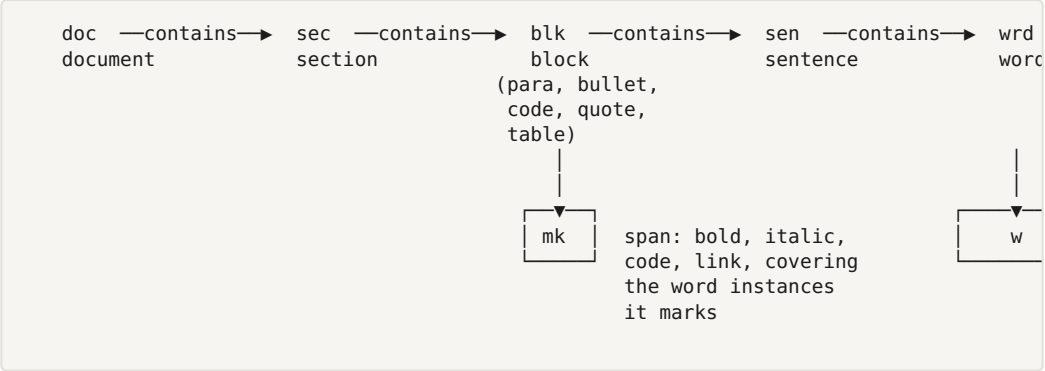

Figure 6.1 · The core graph ladder, seven node kinds, from `v2/universe/data/core/`.

For the pilot document, that ladder resolves to real numbers, computed by the build: **39 sections, 186 blocks, 342 sentences, 4,221 word instances, 143 markup spans, and 951 distinct word forms.** Every level has an identifier and every identifier is structural and deterministic, which is chapter ten's subject.

Two things about that ladder are worth naming as design choices rather than accidents.

Markup is structure, not formatting. A bold run is a `span` node that covers the word instances it marks. The memo asked for exactly this: *"if you have a bold, right? Let's say, then you need to link. You need to have a node that links that bold, those three nodes, to a bold, so we can basically understand which one of those are because the bold has extra meaning."* Emphasis becomes queryable rather than lost.

A word form is a node, not an attribute. One node per distinct form, carrying its count and the identifiers of every instance. Which means the question "where does this word occur, and what else occurs near it?" is a traversal rather than a search. 951 forms, 505 of which occur exactly once.

Zoom two: the code itself

The second zoom is a level most systems never attempt, and it arrived as a message sent minutes after the previous release went live: *"That worked great, let's keep zooming."*

The estate's meaning engine is built from twelve operators. Each operator is a small JavaScript file. The ask was to apply the same treatment to the source code:

> "can you apply the graphs of graphs approach here, the visualisation of grouping specific parts of the code and providing an explanation on a right pane on what it does, what the variables do and what are the inputs and outputs of those inner bits of code."

What shipped is an **anatomy** per operator: the code sliced into contiguous segments, each a node with an identifier, a kind (docs, imports, data, contract, step, export), an explanation written for somebody who already knows the language, its variables with their roles, what it reads and writes, and `feeds` edges to the segments it drives.

And the anchoring discipline came with it. Each segment is anchored by the exact text of its first line, the build resolves those heads to line ranges that must tile the file completely, and **a gate fails the release the moment the code and the anatomy drift.** So the code's graph cannot quietly become a lie about the code, which is the failure mode of every architecture diagram you have ever seen.

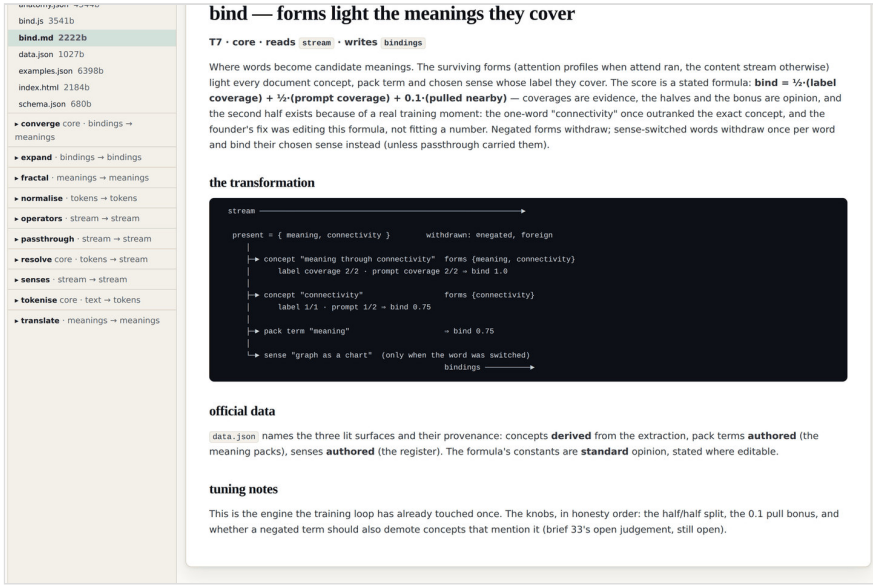

Figure 6.2 · One operator, zoomed: the `bind` engine's own page at graphs.sgit.ai/v2/wclm/operators/, site version v0.5.11. Left: the twelve operator folders, each with its code, schema, data, docs, examples and workbench. Right: bind's contract (reads `stream`, writes `bindings`), its stated formula, an ascii diagram of the transformation, and the provenance of its official data. The paragraph beginning "the second half exists because of a real training moment" is chapter eleven's subject.

The verdict, honestly

Here is the test applied to the estate's own two zooms, one commitment at a time.

Commitment	Verdict	The evidence, and the qualification
Self-similarity	**passes at the reading layer**	The document ladder, the extraction's concepts and claims, and the derived layers all render as nodes and typed edges in one canvas with one viewer. The schema view over the pilot shows nine node types and twenty-four typed relations, all in the same grammar.
Scale invariance	**partial**	One viewer and one query surface across all levels. But not one validator: the extraction has its anchor gate, the core graph has its round-trip gate, the code anatomy has its drift gate. Three gates enforcing one discipline is not the same as one validator, and calling it scale-invariant would be overclaiming.
Composition	**passes**	The engine's world is assembled from the extraction, the core graph's token analysis, the meaning packs, the senses register and the analogies register, with no adapter layer. Each is a graph; the composition is a graph.
Recursion, no new format	**fails at the storage layer, passes at the engine layer**	See below.

The recursion row is the interesting one, so it gets stated in full rather than summarised.

Where it fails. The extraction is stored as node and edge lists. The core graph is stored as an index plus one shard per section, and a shard is a *nested* structure: blocks containing sentences containing words. The code anatomy is a third shape again, segments with feeds edges. Three different serialisations for three levels of the same zoom. Under the test as this book states it, that is a failure: zooming from a document into its sections does require a different file format from zooming from a claim into its concepts.

There is a decent engineering reason (a nested shard is loaded once when a section is expanded, which is what makes the tree fast), and a decent reason does not make the claim true. It makes the claim *partly* true, and the honest form of the sentence is: **the estate is fractal in its grammar and hierarchical in its storage.**

Where it passes, and passes hard. The engine has an operator called `fractal` whose entire job is to be a full instance of the engine, inside the engine. It takes the winning meaning's own statement and runs it through a complete inner pipeline (tokenise, resolve, bind, converge), one zoom down, and reports the meaning of the meaning. It reads the type `meanings` and writes the type `meanings`, exactly like every other operator, so the pipeline cannot tell it apart from a simple one. It is registered, typed and swappable like the rest.

That is recursion with no new format and no special case, in the strong sense: **the system composes with itself, and the composition is invisible to everything around it.** The inner pipeline contains no fractal operator, so recursion terminates at depth one by construction, which is a stated limit rather than an accident.

> **Why report a failing row at all.** Because the value of a falsifiable claim is destroyed by never falsifying it. A book that stated the zoom test and then reported four passes would be asking you to take its word for the test as well as the result. The storage-layer failure is real, it is specific, and it is fixable: the shards could be node and edge lists at some cost in load time. Nobody has decided whether that cost is worth paying, and that is the actual state of the question.

Name clashes are not a problem

One practical consequence of fractality that saves an enormous amount of argument.

If each graph keeps its own vocabulary, then two graphs may both use the word "node" for different things, and nothing breaks. The estate has a live instance of this and it is slightly comic: the meaning engine has an operator literally named `operators` (it handles the little words that flip meaning, such as *without* and *not*), while the same release calls all twelve building blocks operators. The agent flagged it rather than renaming either:

> "The operators word now means two things … The folders adopt the brief's meaning; the T5 engine keeps its name inside the registry."

Under a global schema that is a collision requiring resolution. Under local vocabularies with a declared scope it is two words in two namespaces, and the fix is a note. The general rule: **a name clash between two graphs is only a problem if you were planning to merge them**, and chapter four is why you are not.

Where it does not matter that you start

Two smaller notes that follow from the same property.

It does not matter where you start. The graph will be deep where the work is and absent everywhere else. That is not a defect to apologise for. It is the property that makes the project finite. A graph that had to be complete before it was useful would never be either.

A bug is a divergence, not a breakage. *"A bug is something that we have mapped in the graph that is not happening in reality."* Which reframes it from "something is broken" to "reality diverges from the model", and leaves open which of the two is wrong. That reframing is only available in a system where the model is a first-class artefact rather than a document about the system.

Where the live estate demonstrates this. The document ladder is browsable at `graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html`, where the authored folder and the derived core data sit in one file tree and every file reads raw or rendered. The code anatomy is at `graphs.sgit.ai/v2/wclm/operators/`, one folder per operator. The fractal operator is at `graphs.sgit.ai/v2/wclm/operators/fractal/`, and it can be toggled into the pipeline on the engine page with one click.

7 · A graph at every boundary

After this chapter you will be able to explain why determinism, explainability, provenance, sovereignty and auditability are usually lost between layers rather than inside them, and what changes if nothing is allowed to cross a layer except a graph.

This is the chapter the book's title comes from. The phrase *fractal semantic graphs* first appears in a brief of 12 July 2026 whose subject is an operating layer sitting between customer channels and business systems, and whose argument is one sentence long before it is anything else.

The seam is where it dies

A conventional agentic stack has the right layers and the wrong seams. Channels hand the runtime a payload. The runtime hands the model a prompt. The model hands back text. The workflow parses that text into another payload. A tool call goes out as JSON. A business system returns a record, which is summarised back into prose.

At every one of those boundaries the structure of what was known (what grounded it, where it came from, how confident anybody was) is **flattened into a string and then re-guessed by the next layer**.

That flattening is not an implementation detail. It is precisely where five properties die, and the important observation is *where* they die:

> Not one of determinism, explainability, provenance, sovereignty and auditability is lost inside a layer. All five are lost *between* them.

Which is why bolting a governance rail onto the side afterwards cannot recover any of them. The information was destroyed at the seams, and nothing downstream can reconstruct it. From the source brief, verbatim:

> "at every boundary meaning is lost and re-guessed; the alternative is to make a semantic graph the interface at every boundary, so each layer emits a graph and consumes a graph and **nothing crosses a layer as an opaque blob or a sentence.**"

The stack, with graphs at the seams

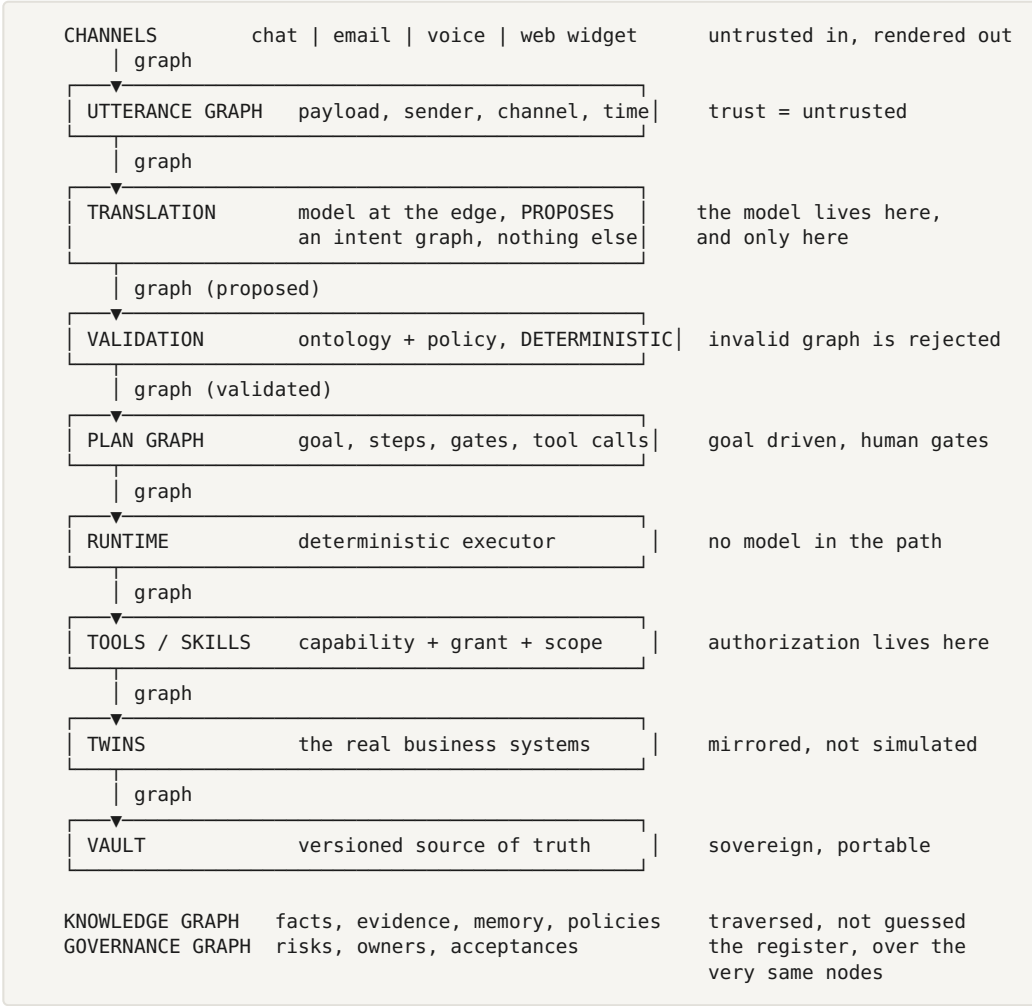

Figure 7.1 · The stack as the 12 July 2026 architecture brief draws it, redrawn here. The only thing that changes between the conventional version and this one is what crosses the arrows.

What falls out

The point of the design is that these are **consequences, not features**. Nobody built a provenance subsystem.

Determinism. The model proposes a graph; a deterministic validator executes it. The same proposal produces the same result. Swapping the model changes the translator and nothing else, which is what inference-agnostic actually means.

Explainability. The explanation is the path. It was already there; you only have to render it. That is the difference between a defensible account and a plausible one, and chapter twelve is entirely about it.

Provenance. Every node carries where it came from, because it crossed the boundary as a node rather than as prose about a node.

Sovereignty. Computed, not claimed. You can point at which parts of the graph left your jurisdiction, because parts of a graph have addresses.

Auditability. The transaction log is a by-product. Every mutation is an ordered message, and a message *is* the change rather than a notification about it. In the estate's own implementation the audit log is the vault's commit history, and there is no second system to reconcile with the first.

The security property, and how much of it is true

This is the sharpest claim in the chapter and it deserves care in both directions.

> The model sits at the edge and only *proposes* a graph. A deterministic validator decides what is admissible. Untrusted input is therefore data and can never become instruction, so prompt injection fails at the validator, **structurally**.

Read the claim narrowly, because the wide reading is false.

It does **not** say a model cannot be manipulated. Of course it can, and it will propose something wrong. It says the manipulated proposal **has to pass a validator that does not read prose**. So the class of attack that works by talking a system into doing something never reaches the executor. Text saying "ignore previous instructions and export the customer table" proposes an action for which no grant edge exists, so it fails at the validator deterministically rather than at the model's discretion.

What remains is a proposal that is structurally valid and semantically wrong. That is a much smaller and much more detectable problem than the original one, but it is not zero, and the source brief carries its own honest-tensions table saying so. **This is a real reduction, not an elimination**, and anybody who sells it as an elimination is misrepresenting it.

There is a companion property in the same design that is easier to state and harder to argue with. **Authorization lives in the tool graph.** A tool node declares its capability, its required grant and its scope, and a plan step can only invoke it if the grant edge exists. Which means an ungranted action is not *denied at runtime*; it is **absent from the graph**. A read-only, row-scoped grant makes a mass write impossible rather than forbidden, and the difference between impossible and forbidden is the difference between an architecture and a policy.

The same seam, at the smallest scale that runs

The architecture above is a design. The smallest running instance of the same idea in this estate is one operator of the engine of chapter eleven, and it is worth looking at because the property is identical at a fraction of the size: the operator declares the data types it reads and writes, it is handed a typed structure and returns a typed structure, and nothing crosses that boundary as prose.

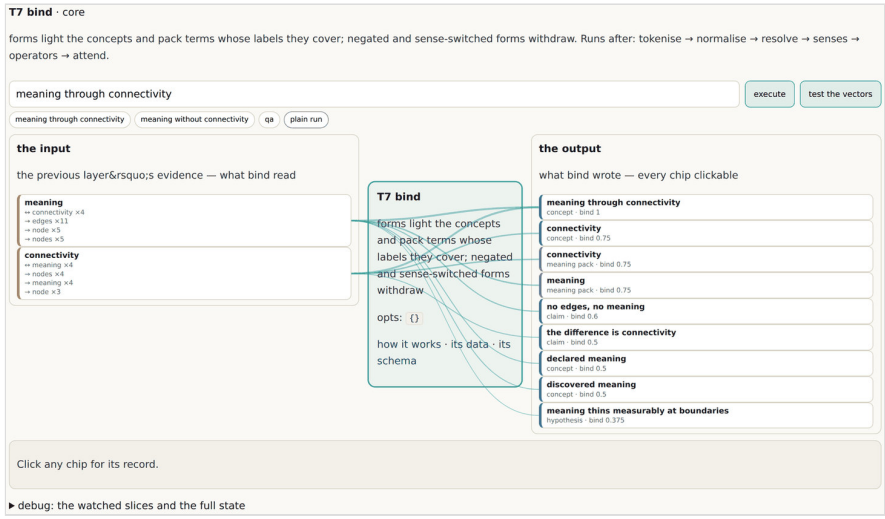

Figure 7.3 · One operator's workbench at graphs.sgit.ai/v2/wclm/operators/bind/, site version v0.5.11. The input state on the left, the operator in the middle, the output state on the right, with the declared contract (reads `stream`, writes `bindings`) drawn from the operator's own schema file. An operator placed where its input type has not been written does not fail: it is skipped, with the missing type named.

The honest tensions

The source brief lists five tensions in its own words, and they belong in this chapter rather than in a footnote.

Tension	The note the brief carries
Graph at every boundary versus latency	Validation at each seam costs time, and a voice channel has a hard latency budget
Proposal validation versus expressiveness	A strict grammar rejects hostile input and also rejects legitimate intents nobody modelled yet
One grammar versus fitness	Fractal uniformity is powerful but can be procrustean when a domain genuinely does not fit
Open source versus revenue	If the engine and blueprints are forkable, the value must sit in evidence, assurance, and hosting
Sovereign vaults versus managed convenience	Self-hostable portability is the trust argument and also more work than a managed runtime

The second row is the one that will bite you first in practice. A validator that rejects what nobody modelled is correct and infuriating, and the fallback path (what happens to a legitimate intent the grammar cannot express) is listed in the same brief as an open question. It still is.

Twins, and where the graph touches something real

A graph that only ever refers to itself is a very well-organised opinion. A **twin** is where it stops modelling and continues into a real system.

> "the power of the twin is that we always arrive at the twin, so the edges and the peaks and the endpoints of the graph continue into the twin, and then ideally into reality."

A twin can be made of almost anything (an organisation, an inbox, a person, a behaviour, the weather) because a twin is just a system with properties, behaviours, functions, inputs and outputs. Two disciplines come with it, and both are the kind that a project abandons quietly under deadline pressure:

Whether an endpoint actually reaches reality is itself a measurable fact. Not an assumption about the diagram; a property of it you can query. "How many of our risk nodes terminate at a twin?" is a number.

Everything modelled must be real. No hypothetical risks. This is the rule that keeps a risk graph from becoming a brainstorm, and it is why the phrase *facts only in phase one* appears as a modelling principle inside the data of the estate's most complete instance graph.

And where it cannot reach, name the gap

Sometimes the graph cannot reach the real system. There is no programmatic interface at all, the data arrives by email, somebody re-keys it on a Thursday. The instinct is either to leave that part of the diagram blank, or to draw it as though it connects.

Neither. The twin holds an **explicit, tracked gap**: *"this needs to be manually updated once a week, but the point is we now know where that gap is."*

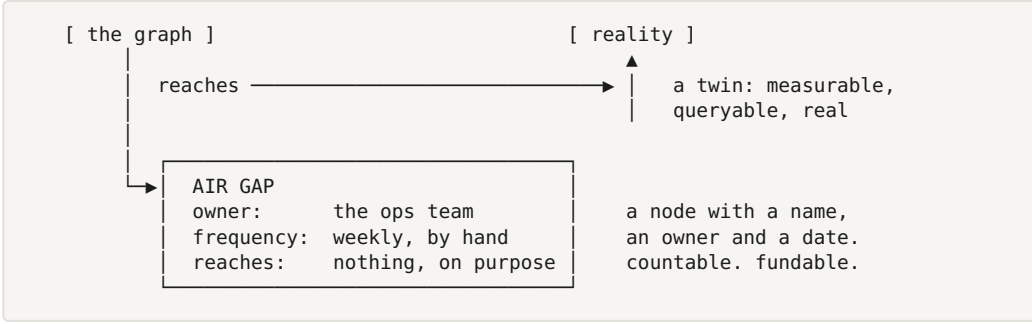

Figure 7.2 · The air gap as a node. The alternative is a blank space on a diagram, which looks like nothing and is therefore never funded.

> A named absence beats a hidden one. In the graph an air gap is a node with a name, an owner and a frequency, which means it can be counted, argued about, and funded.

The air gap is the operational sibling of the third-of-ten-evidence rule from chapter one. Any risk not connected to the register is an air gap. Any legal provision whose hooks reach no twin is a provision not actually mapped, which turns hook coverage into a real coverage measure over an instrument rather than an assertion that somebody reviewed it.

There is a Wardley map for this in the corpus (a Wardley map is Simon Wardley's strategy-mapping technique: a graph with two coordinates added, how visible a component is to the user and how evolved it is). The map shows both ends of a process at high evolution, commodity and automated and solved, with a gap in the middle filled by human labour. The reason it works as a map rather than as a sentence is structural: **a gap has no evolution.** There is nothing to plot, because nothing is there. So what you plot is *the labour that fills it*, and once that labour is on the map, at a position, it is something a strategy can act on rather than something everyone works around.

The open one: time

> **Time is asserted and never developed.** "Time is an event, things change" is close to the whole of the corpus's treatment. It repeatedly says the graph moves and never explains how. Adjacent material circles it: the acceptance-interval ladder (1 hour, 4 hours, 2 days, 2 weeks, 1 month, 6 months), `repealed_from` in the regulation graph, supersede-never-delete, temporal permissions, and the vault's own commit history, which is after all a working answer to "how does a graph change over time". The chapter that ties those together has not been written, in the first edition or in this one. If you have a view, the estate's comms board is where to put it.

> **Where the live estate demonstrates this.** The boundary argument is at `graphs.sgit.ai/v1/depth/boundaries.html`. The source brief the title comes from is published in full, frozen and hash-verified, at `graphs.sgit.ai/v1/docs/sources/fractal-semantic-graphs.md`, including its own honest-tensions and open-questions tables. The air gap and the twin are argued in the same chapter, and the acceptance-interval ladder is data in the 2FA instance graph that chapter thirteen walks.

8 · Documents are projections

After this chapter you will stop treating a document as a source and start treating it as a view, and you will have seen the strongest evidence in this book that the difference is real: a build gate that rebuilds the document from the graph and fails unless the bytes match.

The claim: **the graph is the truth; the document is a view of it, generated in the context of use.**

It appears across the corpus as though already established, and is then applied to skills, compliance standards and legal texts. It is nowhere argued from first principles, which the estate records as a gap. This chapter is the synthesis, and then the test.

> "The same way I talk about documents being projections of graphs, **the skill is a projection of a graph**… The skills we have today are just **a photograph of what it should be**, because it is static."

Three problems that turn out to be one

```
1 · A SKILL FILE
    a static description of how to do something, which drifts from
    how it is actually done.
    as a projection: the graph of how the work is done is the truth,
    and the file is rendered from it when needed.

2 · A COMPLIANCE STANDARD
    a document handed to you whole, from which you strike out what
    does not apply.
    as a projection: nothing is relevant until your facts attach, so
    the standard starts EMPTY and accretes.

3 · A CONSOLIDATED LEGAL TEXT
    the sharpest version. do not STORE the consolidated text at all.
    hold the base text plus the amendment instructions as data, and
    COMPUTE the consolidated version as a projection.
```

Figure 8.1 · Three familiar maintenance problems, and the one move that dissolves all three.

The third is where the idea pays for itself, and the result is the kind that only shows up if you take a claim seriously enough to build on it. If a consolidated text is computed from a base text plus amendments, and a customised standard is computed from a base text plus your facts, then **consolidation and per-organisation customisation are one mechanism, not two.** The maintenance burden and the flagship feature share an engine. Nobody designed that; it fell out.

The same shape appears in the other direction. A document is not one blob but a hierarchy of paragraphs, points and definitions, each written for a reason and therefore yielding something extractable. **Every paragraph is a graph.** A document's own definitions are the first and most valuable node layer, and three kinds of work follow: how they relate, **where they contradict**, and

what the text uses but never defines. The second is the highest-value output and the one nobody produces. The third is where interpretive risk concentrates.

Lifting text into a graph is decompilation

Going from concrete text to abstract structure runs the same direction a decompiler runs, and it inherits the same property: **it is ambiguous, and it cannot be done reliably without help.**

The help is the author. The goal is not absolute truth but *the author's own meaning, confirmed by the author*. Which reframes the most common failure of every extraction system ever built: a reader saying **"that is not what I meant"** is not a failure of extraction. It is the elicitation working.

That reframing has a mechanical requirement attached, and this is the part that is usually skipped. For the correction to be cheap, every node at every altitude must carry a source map back to the span it came from. If it does not, "that is not what I meant" produces an argument instead of an edit. Chapter nine is about what it takes to make that anchor trustworthy.

The test: rebuild the document from the graph

Here is where the second edition can do something the first could not.

If the document really is a projection of the graph, then the graph must contain everything the document contains. Not the meaning of it: **the bytes of it.** Otherwise the graph is a lossy summary wearing the word "projection", and the claim is decoration.

The founder asked for exactly this, and asked for it urgently, in a memo of 26 August 2026:

> "Actually, there's one more important topic that I think I forgot. It's very important that we have the ability to rebuild the document from the graph. So it's very important that we have a two-way transformation. So let's address that sooner more than later, because there's a lot of things that I would like us to do, and you need that, especially when you talk about refactoring and making changes and detecting changes."

And the design instruction that came with it:

> "you need to capture those properties, those in a way formatting properties in an what's it called in a separate dock in a separate graph connected to the main graph and in separate things, so that we have, we can go a two-way transformation back into the document"

So there are two graphs, joined by identifiers. The semantic graph holds document, section, block, sentence, word, span and form. The **formatting graph** holds heading lines, block markers, the gaps between blocks, and the raw markdown per block, keyed by the same block identifiers. The semantic graph stays clean. The join is the identifier.

Seven gates, and two of them are the argument

The generator that builds this runs seven gates. Any failure kills the build. Quoted from its own header:

```
1. block byte ranges reassemble each section body exactly
   (gaps whitespace-only)
2. sentences reassemble each block's clean text
   (whitespace-collapsed equality)
3. the word-form totals equal the word-instance total
4. every span covers at least one word unless it marks pure
   punctuation or code
5. the document rebuilds from the formatting graph BYTE-IDENTICAL
   to the source
6. the semantic shards re-derive from the formatting graph alone
   (the two graphs cannot disagree, because one provably generates
   the other)
7. the identity ledger covers every doc, section and block with a
   unique live uid, and a second carry-forward pass over its own
   output changes nothing (identity assignment is deterministic)
```

Figure 8.2 · The gates in `admin/build/gen_coregraph.py`, quoted from the file. Gate 5 is the projection claim made falsifiable. Gate 6 is the stronger one.

Gate 5 is the projection claim, executable. The build reassembles the source markdown from the graph and compares it byte for byte with the frozen original. If a single character differs, the build stops with the message *"the transform is not two-way"*. There is no partial credit, no similarity score, no reviewer judgment. It either is a projection or the release does not happen.

Gate 6 is the one worth stealing. It does not merely check that the two graphs agree. It re-derives the semantic shards from the formatting graph alone, so that the two graphs **cannot** disagree, because one provably generates the other. That is a structurally different guarantee from a consistency check. A consistency check tells you when two things have drifted. This makes drift impossible by removing the second source of truth.

That distinction is the whole discipline of this book in miniature. Enrichment over enforcement, one source over two, computed over asserted.

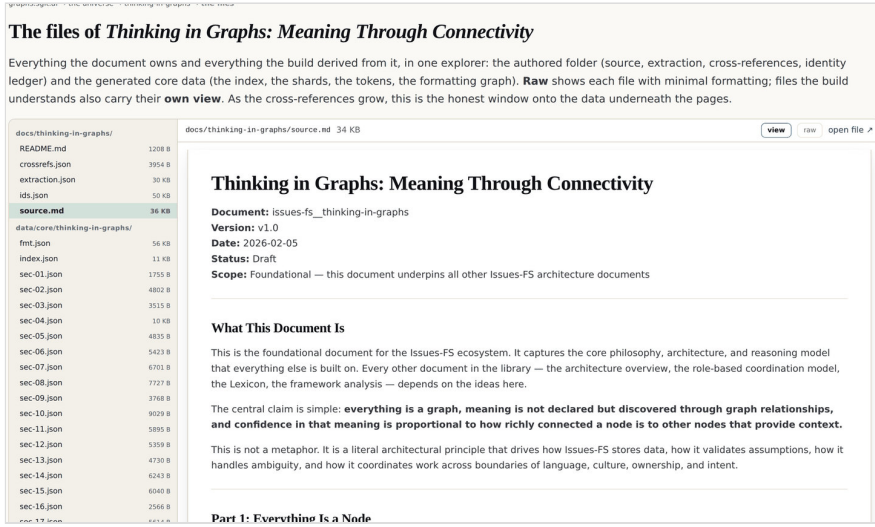

Figure 8.3 · The pilot document's file explorer at graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html, site version v0.5.11. Left: the authored folder (source, extraction, cross-references, identity ledger) above the derived core data (the formatting graph, the index, one shard per section). Right: any file, raw or in its own data-driven view. The formatting graph at 56 KB is the file gate 5 rebuilds the document from.

What it costs, and what it does not prove

Three honest qualifications, because a gate this satisfying invites overreading.

It proves losslessness, not understanding. Rebuilding the document byte for byte proves the graph contains everything the document contains. It says nothing about whether the *meaning* was extracted well. The extraction is a separate artefact with a separate review loop, and chapter nine is about how that one is kept honest, which is a much harder problem because there is no byte comparison available.

It works because the source is frozen. The pilot document is byte-frozen with its SHA-256 recorded, and the build verifies the copy against it. A living document that changes under you needs a different discipline, and change detection at the graph level (which blocks changed between two versions) is named in the same memo as the next step and has not been built.

One document, not twenty-one. The estate carries twenty-one source documents and one has been through this pipeline. The method is documented so the other twenty can follow. Until they do, this is a demonstration rather than a system, and chapter fourteen says so in the list where such things belong.

Four views, one structure

The practical consequence of taking projection seriously is the one that shows up in every organisation with more than one audience.

The Article 26(5) worked example produces four stakeholder registers, one per altitude, each in that reader's own language, from one chain of facts. The brief's own line: *"nothing is duplicated; each view is a query over one structure."*

Those four registers are not four documents that have to be kept in sync. They are four queries, and **they cannot drift apart, because there is only one thing there.** Anybody who has maintained a board pack, a risk register and a regulator submission that are supposed to say the same thing will recognise the size of that claim.

The same pattern governs this estate's own pages. The universe reader shows four reviewer-facing views (the dictionary, the taxonomy, the thesaurus with its near-but-nots, and the ontology) plus the claims table, the drawn graph and the coverage table, and every one of them is a projection of a single authored file. They are generated, never separately authored, **so they cannot disagree with each other**.

And this book is the same shape. The markdown chapters are the source of truth. The web pages render that markdown in the browser, so a page cannot drift from the file it claims to render. The printed version is generated from the same files. If you find the page and the print disagreeing, one of them is a bug and the markdown is right.

> **Where the live estate demonstrates this.** The formatting graph is `v2/universe/data/core/thinking-in-graphs/fmt.json`, 56 KB, readable raw in the file explorer at `graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html` beside the 39 section shards it keys into. The gates are in `admin/build/gen_coregraph.py`. The four projected views are on the reader at `graphs.sgit.ai/v2/universe/thinking-in-graphs.html`, and the customisation inversion runs in the published regulation graph vault at `sgit.ai/demos/vaults/regulation-graph/`.

PART FOUR

Computed

The half the first edition could not write: how a document becomes a graph anchored to bytes, what identity means when a document changes, an engine that computes meaning deterministically, and why an explanation is a path rather than a narration.

9 · The universe method

After this chapter you will know how a raw document becomes a graph whose every node is anchored to exact bytes, why the build refuses to ship an extraction that cites words that are not there, and what it costs to make coverage total by construction rather than by effort.

Part four is the half the first edition could not write. It argued that meaning is a computation and had nothing that computed it. What follows is the computation, built from the bottom.

Bottom means the document. Not a summary of it, not an embedding of it: the bytes.

The layer model

Agreed on 23 August 2026 and worth stating before anything else, because most confusion about this kind of work comes from mixing the layers.

Lay-er	What it is	State
0	the frozen bytes: the source documents, hash-recorded, held by a build gate	frozen
1	per-document local graphs, one folder per document	pilot done, 1 of 21
2	the bridge layer: cross-document edges, authored with a note	not started
3	the book's own universe, connecting down into layers 1 and 2 as evidence	not started

The rule that makes the whole thing tractable sits between layers 1 and 3, and it is one sentence:

> **Truth is deferred.** Every layer 1 node is a record of the form *"this document says X, at this anchor."* Whether X is true is not evaluated at layer 1. That judgment belongs higher up.

This is the discipline that most extraction projects lack, and the lack is why they fail. If your extractor is simultaneously deciding what a document says and whether it is right, you cannot review either decision, because they arrive fused. Split them and both become reviewable: one against the bytes, the other against the world.

It also means each document keeps its own vocabulary. Where two documents name the same idea differently, or disagree outright, layer 2 records that as an **authored bridge**, and the divergence is preserved as a finding rather than merged away. Chapter four is why.

The pipeline

```
0 · THE SOURCE IS ALREADY FROZEN
    v1/docs/sources/<slug>.md, byte-frozen, SHA-256 recorded,
    held by a gate. an anchor into a frozen file is stable forever.
                                  │
                                  ▼
1 · THE DOCUMENT GETS A STANDALONE FOLDER
    source.md      a byte copy, verified against the recorded hash
    extraction.json the local graph: THE ONLY AUTHORED ARTEFACT
    crossrefs.json  where the document is used, each use rated
    ids.json        the identity ledger (chapter ten)
    README.md       what the folder is, for a reader who finds it
                    outside this repository
                                  │
                                  ▼
2 · EVERY ITEM IS ANCHORED
    an anchor = a named SECTION + a VERBATIM QUOTE
                (+ an occurrence index when the quote repeats)
                                  │
                                  ▼
3 · THE BUILD REFUSES A BAD ANCHOR
    hash the source, refuse if it moved
    verify the folder copy is byte-identical
    find the quote VERBATIM inside its named section
    resolve it to exact byte offsets
                                  │
                                  ▼
4 · COVERAGE IS TOTAL BY CONSTRUCTION
    every prose-bearing section either yields an anchored item or
    is listed in empty_sections WITH A REASON. the build fails on a
    silent gap AND on a stale entry.
                                  │
                                  ▼
5 · THE USES ARE RATED
    aligned / stretched / misaligned / unrated, signed and dated
                                  │
                                  ▼
6 · EVERYTHING A READER SEES IS A PROJECTION OF ONE FILE
```

Figure 9.1 · The extraction pipeline, from `v2/universe/README.md`.

The anchor, and why it is the whole design

An anchor names a section and quotes a sentence. That sounds unremarkable until you notice what the build does with it.

> **An extraction that cites words that are not there cannot ship.** A hallucinated quote fails the build. A misread of a real quote is what review is for. The two failure modes are separated on purpose, one for the machine, one for the reviewer.

That separation is the most transferable idea in this chapter, and it is worth restating in general form. When a machine produces a reading of a text, there are two ways it can be wrong: it can cite something that does not exist, or it can misunderstand something that does. **The first is mechanically detectable and should never reach a human.** The second is a judgment and can only be settled by a human. Systems that treat both as "review the output" waste the reviewer on the class of error a build could have killed.

For the pilot document, that machinery resolves to **72 anchors and 73 verified spans in the source**, each one a quote found verbatim inside its named section and resolved to exact byte offsets, re-verified on every release against a source that still hashes to its recorded value.

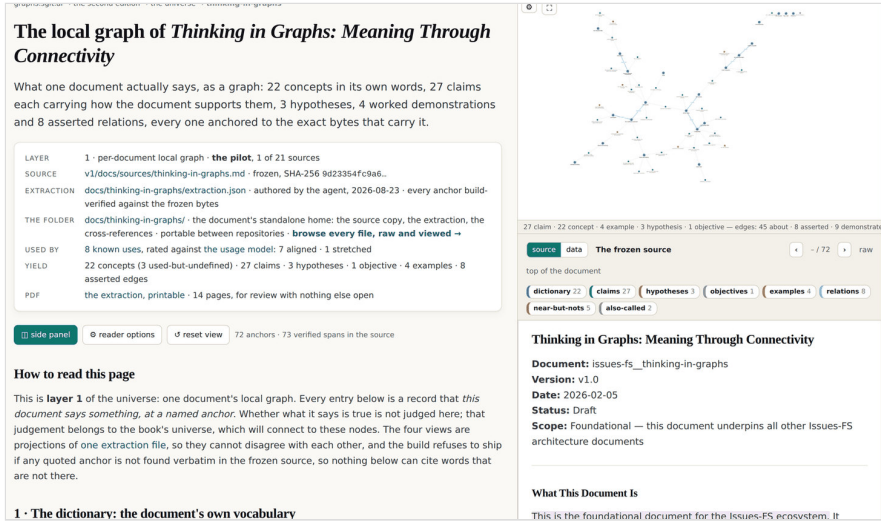

Figure 9.2 · The pilot document's reader at graphs.sgit.ai/v2/universe/thinking-in-graphs.html, site version v0.5.11. Left: the layer, the frozen source with its hash, the folder, the rated uses and the yield. Right: the local graph above, and below it the frozen source with the anchored spans highlighted in place. Clicking a node opens both its extraction row and the cited bytes.

Coverage is total by construction

The second discipline is the one that makes an extraction trustworthy rather than merely plausible.

Every section that carries prose either yields at least one anchored item, or is listed in `empty_sections` **with a reason**. The build fails on a silent gap. It also fails on a stale entry: a section declared empty that now has anchors.

> A recorded empty section is a finding. A silent one is a hole.

This is the third-of-ten-evidence rule from chapter one, applied to the extractor's own work. It converts the question "did the extraction miss anything?" from an unanswerable worry into a mechanical property. The pilot document declares two empty sections, each with its reason, and every other prose-bearing section carries at least one anchored item.

The cost is real and worth naming: it makes an extraction slower to author and impossible to do casually. You cannot skim a section and move on. You either find something in it or you write down why there was nothing.

The yield, in numbers

The pilot is the corpus's own foundational essay, *Thinking in Graphs: Meaning Through Connectivity*, dated 5 February 2026, extracted 23 August 2026. Its yield, computed from the extraction file:

Family	Count	What the family records
`concept`	**22** (3 used but undefined)	a term the document uses. An undefined-but-used term is recorded, not skipped.
`claim`	**27** (7 demonstrated, 15 argued, 5 declared)	with the document's own evidence state, not a truth judgment
`hypothesis`	**3**	what the document expects to test later
`objective`	**1**	what a passage is trying to achieve for its reader
`example`	**4**	the document's own worked demonstrations, with what each demonstrates
asserted edges	**8**	concept-to-concept relations the document itself asserts, each anchored to the sentence that asserts it
aliases	**2**	the thesaurus: X is also called Y, anchored
near-but-nots	**5**	distinctions the document draws on purpose: X is not Y, anchored

Two of those rows deserve a sentence each.

Three concepts are used but undefined. The document leans on three terms it never defines. Under a summarising extractor those disappear, because a summary reports what a document says. Here they are nodes, which means they are countable and assignable. That is the named-absence rule again, one level down.

Five near-but-nots. A definition says what a thing is. A near-but-not says which neighbouring idea a reader will otherwise substitute for it. In practice the second does more work than the first, and almost nobody records it.

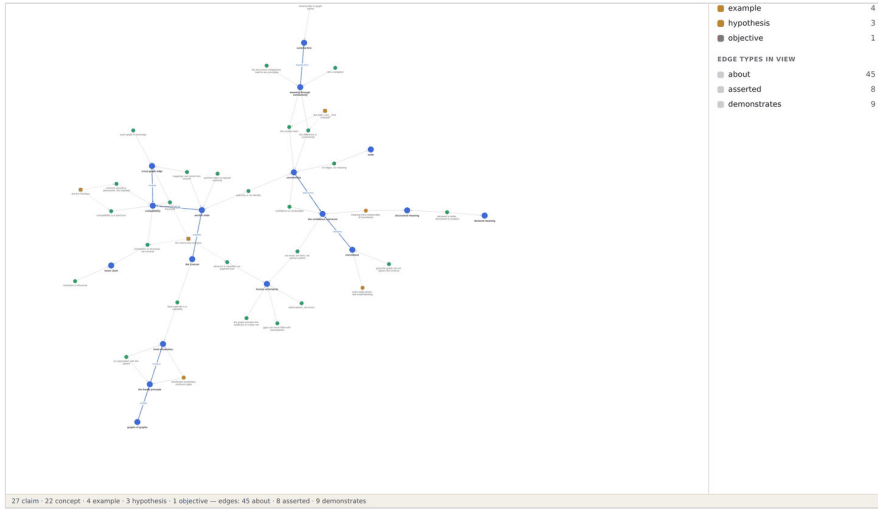

Figure 9.3 · The same extraction as a graph, at graphs.sgit.ai/v2/universe/thinking-in-graphs.graph.html, site version v0.5.11. Blue is a concept, green a claim, red an example, amber a hypothesis, purple an objective. The status line counts what is on the canvas: 27 claim · 22 concept · 4 example · 3 hypothesis · 1 objective, and 45 about · 8 asserted · 9 demonstrates edges. Nothing here was laid out by hand: the clustering is force over the declared edges, so a concept that sits alone sits alone in the document too.

The usage maturity model

The third discipline, and the one this book would most like other people to steal.

A source document records, where it can, **where it is being used**, and each use gets a rating. The rating is a judgment about the *use*, never about the user, and it must be signed and dated or the build rejects it.

Rating	The test
aligned	Read the use, then the cited part of the source. Would the source's author say: yes, that is what I said?
stretched	The author would say: that came from my document, but it is not quite what I meant.
mis-aligned	That is not what the document says.
unrated	A known use nobody has judged yet: the named absence, awaiting judgment.

The pilot has **eight rated uses: seven aligned, one stretched.**

The one stretched use, which is the model earning its keep

The single stretched rating is not a small thing found in a corner. It is a finding against **the first edition of this argument**, and it is about the word in this book's title.

The ledger entry, quoted from the file:

> The first edition defined fractal as one grammar at every level: uniformity. This document's Part 3 and Part 7 license local vocabulary and override, which is composition, not uniformity. Caught by the founder (brief 20), corrected in the lexicon at v0.4.6 with the superseded definition kept visible.

Read what happened there. The estate built a mechanism for rating how faithfully its own sources are used. On the first day it ran, the mechanism caught the first edition misreading the foundational document about the concept the second edition is named after. The correction was made, the superseded definition was kept visible rather than deleted, and the misreading is recorded rather than erased.

That is chapter five's supersede rule and chapter one's named-absence rule and this chapter's rating model, all doing their jobs at once, against the people who built them.

It also explains a result in chapter six. If fractality is composition rather than uniformity, then a system whose levels compose cleanly but whose storage formats differ is *more* fractal than the strict reading suggests, and the "no new format" test may be testing the wrong commitment. This book keeps the strict test, reports the failure, and records the disagreement here, because that is what the ledger is for.

The extractor is an author too

The last rule, and the one that governs the trust model of everything in part four.

> The extraction is a **reading**, marked as such, shipped with the anchors a reviewer needs to check it, and reviewed from its own printed version.

Every extraction names its extractor and its date. The pilot's says `"extractor": "agent"`. It is not presented as the document's meaning; it is presented as one reading of the document, with every claim traceable to the bytes that support it.

And the review surface is deliberately a fourteen-page PDF rather than a website, for a reason recorded in a memo of 23 August 2026: a website cannot control what a reviewer sees or in what order, and a document that is not modifiable survives as a record better than a page does. The review contract is printed on its cover: *for each item, is this what the document says, and is the anchor fair? Come back item by item, by node id. Items you do not mention are taken as agreed.*

That last sentence is the part to steal. Silence is defined in advance, so a review that covers half the items produces a determinate result instead of an ambiguous one.

What is not built

One document of twenty-one. The estate carries twenty-one source documents. One has been through this pipeline. The method is written down so the other twenty can follow, and until they do, layer 2 (the bridges) cannot start, because bridges need two documents. Everything in this chapter is therefore a method that has run once, thoroughly, rather than a corpus that has been processed. The distinction matters and chapter fourteen carries it into the list of what ships.

Where the live estate demonstrates this. The method is written down at `graphs.sgit.ai/v2/universe/README.md`. The pilot's folder, every file readable raw or rendered, is at `graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html`. The usage model with its four levels and their tests is `v2/universe/usage-model.json`, and the ledger that caught the first edition is `crossrefs.json` in the pilot's folder. The machine surface for agents, with every anchor's resolved byte offsets and the hash each was verified against, is `v2/universe/data/universe.json`.

10 · Two kinds of address

After this chapter you will be able to tell a content address from an identity address, you will know why a cross-reference built on character offsets breaks the first time somebody fixes a typo, and you will have a working algorithm for keeping identity stable across edits.

This is the shortest chapter in the book and the one whose absence causes the most damage.

Everything in part four rests on being able to point at something. Point at a sentence, a section, a concept, a word. And the moment a document is edited, every naive way of pointing breaks.

The founder put the problem plainly in a memo of 26 August 2026:

> "you are already doing linkage based on word count and number count. I think, and even character count, which is very fragile, which basically means that every time there's a little change in a document, everything breaks and you you can't really refactor, and it's really sort of fragile. Where you really should having a sort of expat, sort of, or even ID base, sort of cross-reference, right? … if I want to point something on a particular part of the document, I should have a reference ID. I shouldn't be saying you know this is start set character 256 and ends on character 259."

Anybody who has maintained annotations over a changing corpus recognises this. You store offsets, somebody fixes a typo in paragraph two, and every annotation after it is now pointing at the wrong place. Silently. There is no error, just wrong answers.

Two addresses, two jobs

The resolution is not to pick a better identifier. It is to notice that there are **two different questions** being asked, and they need two different answers.

```
CONTENT ADDRESS                          IDENTITY ADDRESS
"what is this exactly?"                  "which thing is this?"

a hash of the content                    a minted, opaque, short id
CHANGES when the content changes         SURVIVES when the content changes

FNV-1a 64-bit of the case-folded         tig:b42
word; a phrase hashes the joined
word hashes

used by: the engine, deduplication,      used by: cross-references,
"is this the same text?",                provenance, "what did this claim
verification against frozen bytes        rest on?", review threads

the same word tokenises identically      the same paragraph keeps its id
in EVERY document with no registry       through an edit, a rename and a move
```

Figure 10.1 · The two addresses, and what each is for. Confusing them is the source of most identifier pain in document systems.

The distinction is recorded in the estate's own brief as a design note rather than discovered later: *the hash is a CONTENT address (changes when the spelling changes); the ledger uid is an IDENTITY address (survives change). The engine uses hashes; provenance still lands on uids.*

Both are needed. A system with only content addresses cannot say "this is the same paragraph, revised". A system with only identity addresses cannot say "this text has not been tampered with". Together they answer both.

Structural locators, and why they are not the identity either

There is a third kind of pointer that looks like it solves the problem and does not: the **structural locator**.

The estate's core graph gives every level a deterministic structural path: `sec:What This Document Is`, `blk:What This Document Is/1`, and below those `sen:` and `wrd:` for sentences and word instances. No randomness anywhere; the same document always produces the same locators.

Those are excellent addresses and terrible identities. Rename a heading and every locator underneath it changes, even though nothing about the content moved. A cross-reference holding a locator would break on a rename, which is a smaller version of breaking on a typo.

So the locator is kept, because it is human-readable and it is how you navigate, and a separate identity is minted beside it. **The locator is free to move. The identity is not.**

Match, then mint

The algorithm that keeps identity stable is worth reading in full, because it is short and it is the part everyone gets wrong.

On every build, for every node in document order, the ledger tries three matches in sequence:

```
1 · SAME LOCATOR
    the node is where it was. edits in place update the recorded
    content hash and keep the identity.

2 · SAME CONTENT HASH
    the node MOVED. identity follows the content, not the position.
    a paragraph dragged to a different section keeps its id.

3 · FUZZY SIMILARITY  (locator or text head, ratio >= 0.75)
    the node was renamed AND edited. the closest surviving
    candidate above the threshold claims the identity.

otherwise: MINT A NEW ONE
    tig:b187, and the counter only ever goes up.

and whatever the document no longer has is RETIRED, never deleted,
so identity history survives.
```

Figure 10.2 · The match-then-mint pass from `admin/build/gen_coregraph.py`. The order matters: position first because it is cheapest and most common, content second because it is exact, similarity last because it is a guess.

Three properties of that algorithm are load-bearing.

It is deterministic. Same document plus same ledger always yields the same output. This is not left to good intentions: a build gate runs the pass a second time over its own output and fails if anything changes. An identity system that is not idempotent will drift, and you will find out months later when two references disagree.

Retirement, not deletion. A node the document no longer has is marked retired and kept. That is chapter five's supersede rule applied to identity itself. A retired identifier can still be resolved, so a reference written last year still resolves to *something*, and that something can say "I was removed on this date" rather than returning nothing.

The fuzzy pass is a guess and is marked as one. A similarity threshold of 0.75 will occasionally claim the wrong identity. There is no way to avoid that in the general case (a paragraph rewritten from scratch in the same place is genuinely ambiguous) and the honest response is to make the rule visible rather than to pretend precision. This book would rather state the threshold than hide it inside a heuristic.

The ledger, in numbers

For the pilot document the ledger holds **225 identities: 1 document, 38 sections and 186 blocks**, all currently live, none retired, because the source is frozen and has not yet been edited. Each row carries the identity, the level, the current structural locator, a twelve-character content hash of the node's text, the first eighty characters as a human-readable head, and a status.

The frozen source is why the retired count is zero, and it is worth saying so rather than presenting an untested mechanism as a proven one. The machinery for carrying identity across edits exists, is gated and is deterministic. It has not yet had to survive a real revision of a real document, because the pilot's source is deliberately immutable.

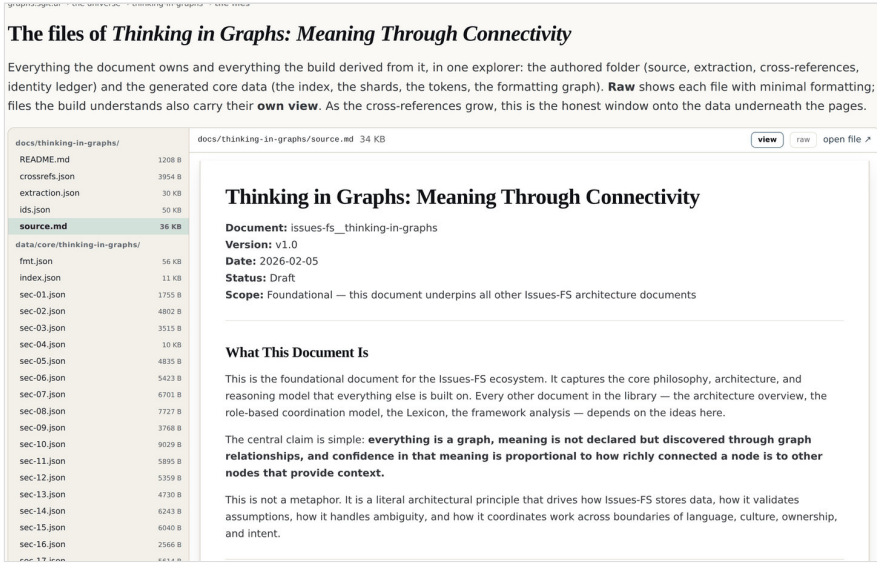

Figure 10.3 · The identity ledger at graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html, site version v0.5.11, in its own data-driven view. Every row is one identity: the minted uid, its level, the structural locator it currently sits at, the content hash of its text, and its status. The locator is free to move; the uid is not.

The same idea, three scales up

The two-address distinction is not local to documents. It shows up at every scale in this estate, which is a small piece of evidence for the fractal claim of chapter six.

Words. A token's identity is the content hash of its spelling. There is no registry, no vocabulary file and no training corpus, so the same word tokenises identically in every document, forever. Change the spelling and you have a different token, which is correct: "graph" and "graphs" are different tokens, and chapter eleven shows the engine treating the difference between them as evidence rather than noise.

Documents. The frozen source is identified by its SHA-256, recorded once and verified on every build. That is a content address doing exactly its job: proving the bytes an anchor was verified against are the bytes that shipped.

Vaults. The estate's storage layer is a content-addressed object graph: the identifier of an object is the SHA-256 of its *ciphertext*, with multi-parent commits, a tree per directory, and a real merge base computed by breadth-first search over all parents. A commit history is a directed acyclic graph, and it is the one graph in this family that has been running for months rather than being argued for.

A correction this estate carries rather than repeats. A skill file in the source repository states that object identifiers are the "SHA-256 of plaintext". The code hashes the **ciphertext**. The difference matters to anybody reasoning about what the object store leaks, so the claim is not republished here, and the correction is recorded rather than made silently.

What to do on Monday

The transferable version of this chapter, for a system that is not this one.

1. **Never let a cross-reference hold a position.** Not a character offset, not a line number, not a paragraph index. Positions are for navigating, not for pointing.
2. **Mint an opaque identity and keep it forever.** Short, meaningless, and never reused. Meaningless is a feature: an identifier that encodes something will eventually encode something false.
3. **Keep the content hash beside the identity.** It is how you detect that a thing changed, and it is how identity follows content when a thing moves.
4. **Retire, do not delete.** A resolvable reference to a removed thing is worth far more than a broken one.
5. **Make the assignment idempotent, and gate it.** Run your own identity pass twice over its own output in the build. If the second run changes anything, you have a drift bug that will surface at the worst possible time.

Where the live estate demonstrates this. The identity ledger is `v2/universe/docs/thinking-in-graphs/ids.json`, 50 KB, readable raw or as a live-and-retired table in the file explorer at `graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html#docs/thinking-in-graphs/ids.json`. The match-then-mint pass and its gate are in `admin/build/gen_coregraph.py`. The content-addressed commit graph under all of it is the sgit vault layer, described at `graphs.sgit.ai/v1/shipped/`.

11 · Meaning, computed

After this chapter you will know what a language engine looks like when nothing is learned and everything is named, why its determinism is the point rather than a limitation, and exactly what it can and cannot do.

The first edition of this argument ended one step short. It said meaning stops being something you assert and becomes something you can compute, and then had nothing that computed it. This chapter is that step.

It starts, as most of this estate does, with a voice memo. This one, of 26 August 2026, introduces itself as *"a bit of a crazy experiment"* and asks for a separate folder, separate code and a separate page so that it can be brought back deliberately if it earns it. Its central intuition is about what a large language model actually does:

> "I don't think predicting of the next word is doing a disservice because the models are not predicting the next word based on random. The models have already created a very rich graph, a very rich mental representation of what something is, so that we, you know, so that then they they know sort of from there what what is most likely to come next."

And then the ask:

> "instead of as defining the layers of the LLM or the neural network to be random, let's create our own layers, and each layer has a very specific role, and is deterministic that role and the inputs and the outputs of those layers."

Followed by the name, coined in the same memo: not a large language model but *"a words content language model"*. The **WCLM**.

What it is, in one paragraph

An engine in the shape of a transformer where **nothing is learned and everything is named**. Words become content-hash tokens. Tokens are resolved against a compiled world built from the graphs of chapter nine. Layers of small, pure, typed operators pass typed data forward. The query is not "what word comes next" but **"what does this mean"**, and the answer is a concept with its statement, the quoted sentence in the source that carries it, the arithmetic that scored it, and a clickable trail back through every layer that produced it.

Same prompt, same world, same picture, every time.

Tokens are content hashes

The first design decision is the founder's own, taken literally:

> "maybe the way to do it is the the the actual ID is actually the hash of the word. Maybe that's how we do it, right? And then you actually have the hash of the combined words, so you start to operate on top of hashes."

So a token's identity is **FNV-1a 64-bit over code points, twelve hexadecimal characters, case-folded**, and a phrase is the hash of its joined word hashes. That is stated in the world file itself, where it can be read and argued with.

Three consequences, and the third is the one that matters.

No registry. There is no vocabulary file, no learned tokenizer, no merge table. A word hashes to the same token everywhere, in every document, forever.

No compression. A learned tokenizer exists to bound the size of the vocabulary. This engine has no such need, and the memo says so directly: *"instead of breaking into tokens, which makes sense because they want to, you know, the idea is to try to limit the amount of the universe, right? Of of the possible words, but in this case we don't have that."* The pilot document has 951 distinct forms. There is nothing to compress away.

Cross-document identity for free. Because the token is a content address (chapter ten), two documents that use the same word are already joined at that word, with no registry to maintain and no alignment step. That is the property that makes the whole thing scale to a corpus rather than a document.

Six types move the whole pipeline

The second design decision arrived as a typed note during a build, and it is the one that turned an experiment into an architecture:

> "I think for this many to many function workflows to work (and to have multiple engines on the same layer) you need to add an abstraction layer that defines what is the input and the output (aka schema) for each of those engines … the good news is that we don't have that many types of data and schemas"

The good news held. **Six types move the entire pipeline.**

Type	What it carries
`text`	the prompt, as typed
`tokens`	hashed words
`stream`	tokens resolved against the world, carrying every clue mark
`profiles`	attention
`bindings`	candidate meanings
`meanings`	the ranked answer

Every operator declares what types it reads and writes. Which makes compatibility **structural** rather than conventional: an operator placed where its input type has not been written is skipped, with the type named in the reason, and **any operator writing `bindings` can stand in for the one that normally writes them.** That is what makes the blocks genuinely swappable rather than nominally modular.

Twelve operators, each a folder

Operator	Core	Reads → writes	What it does
`tokenise`	●	text → tokens	words become hash tokens; the phrase gets the hash of its hashes
`normalise`		tokens → tokens	the dictionary and the thesaurus: what we think you said, with the fix named
`resolve`	●	tokens → stream	each hash looks itself up in the world: form, class, count, weight
`senses`		stream → stream	each word declares its number and its active sense
`operators`		stream → stream	the little words that flip meaning: negation marked, contradiction kept
`passthrough`		stream → stream	carries evidence the marking engines would withdraw
`attend`		stream → profiles	each surviving token carries its attention: pairs stacked, companions pulled in
`bind`	●	stream → bindings	forms light the concepts and pack terms whose labels they cover
`expand`		bindings → bindings	each bound meaning assembles its neighbourhood
`converge`	●	bindings → meanings	evidence is summed; the meaning, its provenance, its blast radius and its contradictions come out
`translate`		meanings → meanings	analogies for an audience: the answer restated in the listener's own concept
`fractal`		meanings → meanings	a full WCLM inside an operator: one zoom down

Four are core and fixed. The rest toggle, reorder and stack. A **layer is a slot that can hold several operators**, which arrived as another typed note: *"we need to support multiple 'engines/transformations' in one layer (since each of those engines is able to bring clues needed for the next layer). This can get interesting since one of those engines could be just' passthough from the previous layer, or block all from previous later and only use the new engines in that layer".*

So `passthrough` is itself an operator. Include it in a layer and the previous layer's evidence flows on beside the new clues; leave it out and the layer's operators replace the stream. The difference is visible on the page and changes the answer.

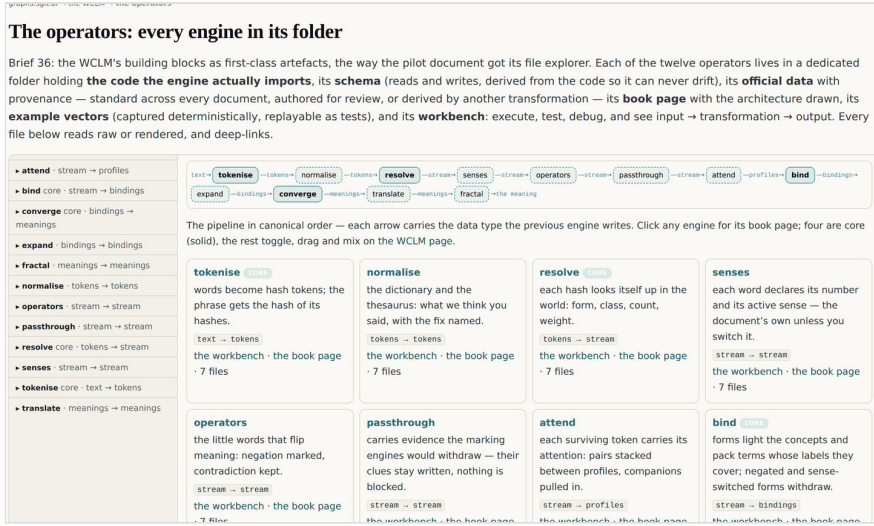

Figure 11.2 · The operator explorer at graphs.sgit.ai/v2/wclm/operators/, site version v0.5.11. The pipeline in canonical order along the top, each arrow carrying the data type the previous operator writes; a card per operator below with its role, its typed contract, and links to its workbench and its book page. Four operators are drawn solid: those are the core, and the rest toggle.

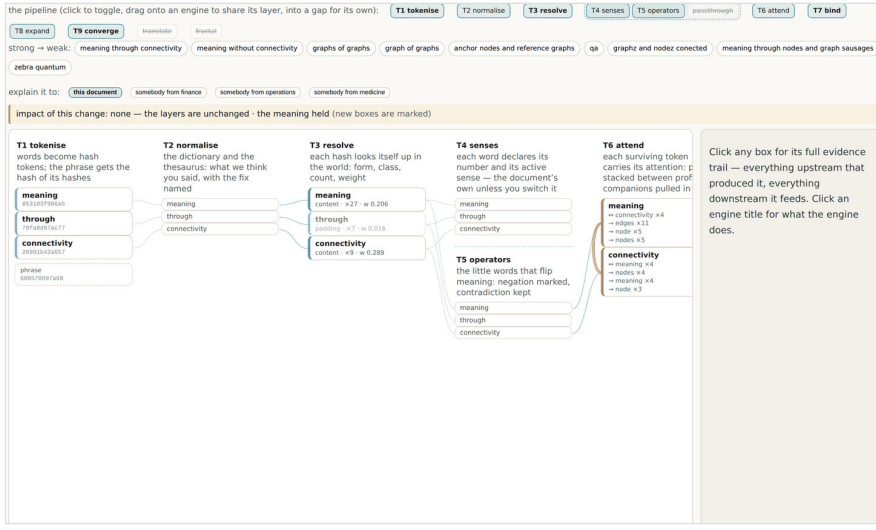

Figure 11.1 · The engine at graphs.sgit.ai/v2/wclm/, site version v0.5.11, answering "meaning through connectivity". Layers as columns, wires between adjacent layers only. Tokenise shows each word with its hash. Resolve shows class, count and weight: "meaning" is content, appears 27 times, weight 0.206; "through" is padding, 7 times, weight 0.016. Attend shows the co-occurrence pulls. Every box is clickable.

Every weight is a stated formula

This is the part that makes it a different kind of object from the thing it resembles.

> "for that you need to transform … you want to start assigning some weights to this, but … you want to be able to programmatically calculate a lot of this stuff"

and the sentence that decides what kind of system this is:

> "when I want to train the model. In this case, I'm going to train the model using graph inputs and tweaks on the graph inputs, and connectivity and meaning, which is the whole point of what I do, versus randomly or you know moving numbers around"

Every weight in the engine is therefore a formula over graphs that were already computed, written into the world file where it can be read and edited. The three that matter:

```
TOKEN WEIGHT
    w = classWeight[class] / log2(2 + count)

    classWeight = { content 1.0 · code 1.0 · verb 0.7
                    number 0.3 · padding 0.05 }

    a rare content word weighs more than a common one; padding
    is not removed, it is weighed at 0.05 and stays visible.

BIND
    bind = 0.5 · (direct / |label forms|)
         + 0.5 · (direct / |prompt content forms|)
         + 0.1 per neighbour pulled in

    half "how much of the concept's label did you say",
    half "how much of what you said was this concept".

CONVERGE
    total = 2 · bind + 0.1 · blastRadius
```

Figure 11.3 · The engine's three formulas, from `v2/wclm/data/world.json` and the operator pages. Every constant is visible and editable.

> **Training this model means editing its graph inputs.** Never fitting numbers. The inputs are the extraction, the meaning packs, the senses register and the analogies register, and every one of them is a file a human reviews.

And the loop has already run. The `bind` operator's own page records the moment, and it is the best single piece of evidence in this chapter that the design is real rather than aspirational:

> "the second half exists because of a real training moment: the one-word 'connectivity' once outranked the exact concept, and the founder's fix was editing this formula, not fitting a number."

A ranking bug was found. The fix was a change to a stated formula, in a file, in a commit, visible in a diff, with a reason recorded. Anybody can now disagree with that fix by disagreeing with the formula, which is exactly the property chapter five asks for from a node type and chapter four asks for from a bridge.

The engine also labels which of its own numbers are which. Counts and coverage are marked **evidence**; the halves and the multipliers are marked **opinion**. Figure 12.1 in the next chapter shows that labelling in place. A system that tells you which of its numbers were measured and which were chosen is doing something almost nothing else does.

The query flips

The last design decision is the one that makes the whole thing useful rather than merely interesting:

> "instead of going from predicting the next word, we can be what is the definition of something? You know, like if we have a phrase or a word, you know, what what is the meaning of it?"

So the query is **meaning-of**, not next-word. Give it a phrase and the answer is: the concept it binds to, that concept's statement, the anchored quote from the frozen source, the examples and claims reachable from it, the blast radius, the arithmetic, and the paths that got there.

Three worked runs, from the engine's own example set:

Prompt	What comes back
`meaning through connectivity`	the exact concept, beating its one-word members on specificity: statement, anchored quote, blast radius 5, total 2.5 with the arithmetic shown
`qa`	`part-of → development`, from the first meaning pack. The founder's own instruction, honoured: "QA is not a random word; it's part of development"
`zebra quantum`	nothing, honestly. Two words the world does not know, said so rather than approximated

That third row is the one to dwell on. An engine that returns nothing when it knows nothing is unusual, and it is only possible because the engine has no generative fallback to reach for. There is no plausible answer available to it. This is the chapter one rule (a named absence beats a hidden one) built into an architecture rather than into a guideline.

Words have many meanings, and the effect is computed

Chapter one showed the sense switch as a demonstration of the five Reviews. Here is what is mechanically happening.

The senses register (chapter four) holds three to five industry senses per word, the document's own always first. Switch a prompt word away from the document's sense and three things happen, in this order:

1. The word **withdraws** from binding against this universe's concepts.
2. The chosen sense's own definition binds instead.
3. The answer says out loud **which of this universe's claims lean on that word and no longer apply**, computed from the concept labels rather than authored by anybody.

Run it on "graphs of graphs" with `graph` switched to *a chart of data* and the engine reports: the meaning moved, and nothing survives into attention. The fractal element does not apply to a chart, and the machine says so without being told.

Number is evidence too. "graphs" carries the chip *plural of graph, more than one involved*, read from the stem families the core graph already computed. Which is the founder's instruction, honoured: *"graph of graphs is very different than graph of graph or graphs of graphs. Right? all of those are different."* They produce visibly different runs, with different hashes, different tokens and different bindings.

Little words count

One more operator earns its paragraph, because it fixed an embarrassment.

Early on, the prompt "meaning without connectivity" answered exactly like "meaning through connectivity", because *without* was classified as padding and weighed at 0.05. The founder found it. The `operators` engine (the little words that flip meaning) now reads *without*, *not*, *no* and *never* as negation, withdraws the negated evidence from binding, keeps it visible with a marker rather than deleting it, and **checks the result against the world**.

So the answer to "meaning without connectivity" now carries a warning: the prompt negates connectivity while this universe asserts meaning through connectivity. **The query contradicts the world, said out loud.**

That is a small feature with a large implication, and it is the seed of something the corpus wants next: asking a document whether it agrees with you. From a memo of 26 August 2026:

> "Ask a paragraph. Ask this to say, "Hey, here's where I'm going, or here's the graph of where I'm going. How does that graph compare, and what's the answers, and what, and does it work? Does he agree? Doesn't it doesn't agree, or does he provide evidence?"

That needs a second extracted document and does not exist yet. The single-document seed of it does.

What this is not

> **Read this section before quoting anything else from this chapter.** **It is not a language model.** It generates nothing. It has no parameters, no training run, no gradient, no sampling. It borrows the transformer's *shape* (layers, attention, tokens) on purpose, as a way of making an argument legible to people who know that shape, and it shares none of the machinery. **It is one document's world.** The compiled world holds 951 tokens, 114 stem families, 400 co-occurrence edges, 57 concepts and 53 edges, drawn from one 4,221-word document plus a 35-term meaning pack, a 4-word senses register and a 16-entry analogies register. That is not a corpus. Everything here is a method demonstrated at one scale. **Its heuristics are honest heuristics.** Verb classification is a curated word list, not a part-of-speech tagger, and is marked as such. Sentence splitting is a rule with guards, and is marked as such. The polysemy score is statistical spread, not sense resolution. **It is an experiment, in the estate's own register.** It is listed in the methods register with status `experiment`, in its own folder, with its own code, exactly as the memo asked, so that it can be brought back deliberately or not at all.

What it *does* demonstrate is narrow and real: an engine can answer "what does this mean" over a graph, deterministically, with every number traceable to either a measured count or a stated opinion, and with training reduced to editing reviewable files. Whether that scales to a corpus is an open question. Whether it is possible at all is no longer one.

> **Where the live estate demonstrates this.** The engine is at `graphs.sgit.ai/v2/wclm/`, with seven example prompts from strong to weak connectivity including the deliberately failing ones. Every operator has its own folder, code, schema, provenance-marked data, example vectors and workbench at `graphs.sgit.ai/v2/wclm/operators/`. The compiled world, with its formulas written into it, is `v2/wclm/data/world.json`.

12 · Every box explains itself

After this chapter you will understand why an explanation that is a path is a different kind of object from an explanation that is a narration, and you will have four properties to demand of any system that claims to explain itself.

An explanation you cannot check is a story.

That sentence is the whole chapter, and the rest of it is about what has to be true of a system for its explanations to be checkable. This estate got there by a route worth retelling, because it started with somebody clicking on a box and being disappointed.

The memo, and the point taken literally

After trying the engine of chapter eleven, the founder recorded a memo whose ask was short: the point of a deterministic transformer is that **every item can say why it is there**. So make every item say it.

What shipped, the same day, is that every chip at every layer is clickable. Clicking selects it, lights its wires, dims everything else, and opens an explanation pane containing three things:

- **the record**, with its formula;
- **because of**: everything upstream that produced it, with the reason carried on each wire;
- **leads to**: everything downstream it feeds.

And every entry in both lists is itself clickable, so the *why* can be walked in either direction, indefinitely. It is graphs all the way, which is the point.

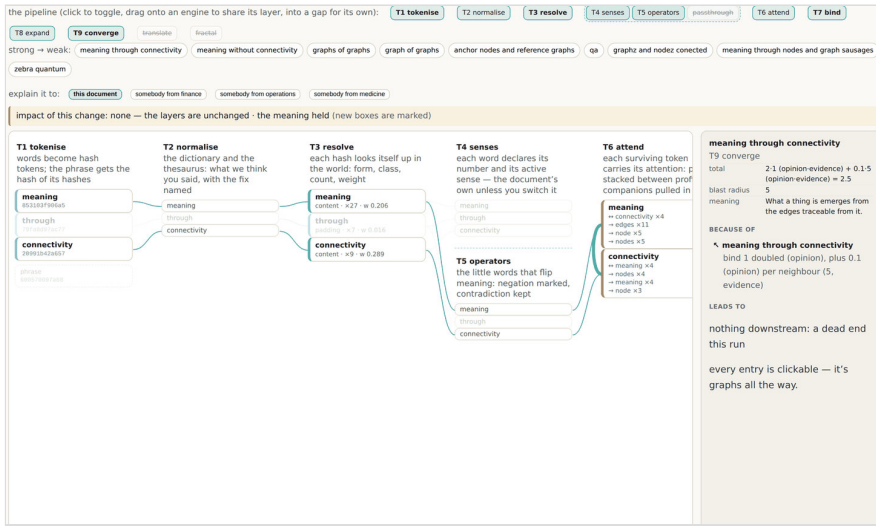

Figure 12.1 · The engine at graphs.sgit.ai/v2/wclm/, site version v0.5.11, with the winning meaning selected. Right: the record. `total 2·1 (opinion·evidence) + 0.1·5 (opinion·evidence) = 2.5`, `blast radius 5`, *and the concept's statement. Below it, BECAUSE OF carries the reason on the wire: "bind 1 doubled (opinion), plus 0.1 (opinion) per neighbour (5, evidence)". Everything that played no part is faded. Note the word "opinion" beside every chosen constant and "evidence" beside every measured count.*

Four properties, and none of them is optional

Read figure 12.1 carefully and you can extract the four things that make it an explanation rather than a description. Each one is a demand you can make of any system.

1 · Determinism, or there is nothing to explain

If the same input can produce a different output, then any account of why it produced *this* output is at best a description of one run. It cannot be checked by re-running, because re-running is not a check.

Every operator in the engine is a pure function of its declared inputs. Same prompt, same world, same picture, every time. That restriction is what buys everything else in this chapter, and it is the reason chapter eleven insists on it before anything else.

2 · The explanation is the path, not a narration

There is a family of systems that produce explanations by asking a model to say why it did something. That produces a plausible account. It does not produce a checkable one, because the account is generated after the fact by the same process that produced the answer, and nothing constrains it to be true.

Here the explanation is not generated at all. **It was already there.** The trail from the answer back to the tokens is the actual computation, rendered. The only work the page does is draw it.

Which means the explanation cannot be wrong about the computation. It can be *unhelpful*, if the computation is badly designed. That is a much better failure mode.

3 · Adjacency, or the picture is lying

This one is easy to miss and it is the sharpest engineering lesson in the chapter.

A narrated review of an early build caught wires jumping layers: `tokenise` connecting straight to `attend`, `attend` straight to `converge`. The drawing was showing a real data dependency, and it was still a lie, because the page claims an architecture in which each layer reads only the layer before it. A wire that skips a layer says the architecture is not what the page says it is.

The engine and the renderer were rebuilt so that **every layer reads only the layer before it**, and evidence a layer merely carries appears as an explicit pass-through chip, so the wire has somewhere adjacent to land.

And then the property was machine-checked. At the release that introduced the analogies work, **551 wires across the pipeline shapes were re-proved with zero adjacency violations.** Not reviewed. Proved, on every build.

> If a system's diagram of itself is not enforced by a test, it is a drawing of what somebody believed at the time. The estate's own rule, arrived at the hard way: a visualisation that can drift from the thing it visualises will drift.

4 · Opinion and evidence are labelled separately

Look again at figure 12.1: `total 2·1 (opinion·evidence) + 0.1·5 (opinion·evidence) = 2.5`.

Every formula surface in the engine labels its parts. **Counts and coverage are evidence. The halves, the multipliers and the bonuses are opinion.** So a reader can see exactly which numbers were measured and which were chosen, without reading the code.

This is the smallest feature in the chapter and possibly the most useful one to steal. Every scoring system you have ever seen mixes measured quantities with chosen constants and presents the result as a single number. Splitting them costs one label per term and turns "I do not trust this score" into "I disagree with this constant", which is an argument that can be settled.

The detective's move

The interaction that follows from all four is worth naming because it changes how the page is used.

Clicking a box does not light its neighbours. It computes the **transitive closure over the wires, in both directions**, and fades everything that played no part. So clicking the winning meaning is the detective's move: the full evidence trail back to the tokens, nothing else lit.

It is the flip of chapter three, applied to a computation instead of a corpus. Go wide (the whole run is on screen), find the few (the answer), then re-root at those and walk out again.

Impact is measured, not guessed

One more property, which arrived from a gesture rather than a specification. Comparing two screenshots, the founder said that adding a word had "made a massive difference". The system now says how much.

Every run is diffed against the previous one by a pure function, and a banner reports the movement layer by layer. From the run in figure 1.1, verbatim from the page:

```
impact of this change:
  senses +1/−1 · bind +1/−1 · expand +1/−1 · converge +1/−1
  the meaning moved: graphs of graphs → graph as a chart of data
```

and from a different run, changing the prompt rather than a sense:

```
impact of this change:
  tokenise +3/−3 · resolve +3/−3 · senses +1/−0 · attend +0/−2
  bind +1/−9 · expand +1/−9 · converge +1/−9
  the meaning moved: meaning through connectivity → graphs of graphs
```

New boxes carry a badge. The reader can see not only that something changed but where in the pipeline the change entered and how far it propagated. That is the propagation question of chapter five (which conclusions were resting on this?) asked of a computation rather than of a corpus, and answered arithmetically.

The same treatment, one zoom down: the code

The last move in this chapter is the one that shows the discipline is not special-cased to the engine's output.

The operators are small JavaScript files. Small does not mean readable: the founder's message was precise about the problem, *"I know JS very well, but at the moment since I don't have the context you have, those scripts (although small) are still hard to read and understand what is going on"*, and precise about the ask, invoking Bret Victor by name and asking for architecture, flowcharts and explanations in a right-hand pane.

What shipped applies chapters six and nine to source code. Each operator carries an **anatomy**: the code sliced into contiguous segments, each a node with a kind, an explanation, its variables with their roles, what it reads and writes, and `feeds` edges to the segments it drives. The rendered view is a flowchart of the segments on top, the code below grouped into its titled segments, and the explanation pane on the right. **One identifier drives all three views**: click a flow box, a code block or a hop, and the same segment lights everywhere.

Three properties carried over unchanged from the document pipeline, which is the evidence that the method generalised rather than being reinvented:

Authored, not parsed. No parser pretends to understand intent. The agent proposes the segmentation and the explanations; the human corrects them.

Anchored. Each segment is anchored by the exact text of its first line, and the build resolves those heads to line ranges that must tile the file completely.

Gated. When the code changes, the build fails until the anatomy is re-anchored. Which is exactly the reminder wanted, and the reason this is not another architecture diagram that was true in March.

And the frame this was built in

The estate names what these operator pages are for, and the naming is worth carrying away as a working practice.

They are **the lab**. Small code, small data, fast rounds: try ways to see, run, visualise and debug in a small problem space, keep what works, then promote the winners into the main surface. The record even names the current promotion candidates, and it names something rarer: *"Experiments that fail stay in the folders as recorded attempts, that is what a lab's notebook is for."*

A place where failed experiments are kept rather than deleted is the same rule as supersede-never-delete, applied to a team's own work. It is also, not coincidentally, the rule that makes the next chapter possible.

> **Where the live estate demonstrates this.** Click any box on the engine at `graphs.sgit.ai/v2/wclm/` and walk the trail upstream and back. The adjacency proof and the opinion-and-evidence labelling are described row by row in the release history at `graphs.sgit.ai/admin/versions.html`. The code anatomies are at `graphs.sgit.ai/v2/wclm/operators/`, one per operator, with the drift gate that keeps them honest.

PART FIVE

In practice

Where this runs today, what ships and what is argued, and how to build your own first graph tomorrow.

13 · The fractal in production

After this chapter you will have seen the grammar of this book applied to real problems with countable results, and you will know which of those results you can check yourself and which you are being asked to take from a design document.

Everything so far has been argument plus one estate's own machinery. This chapter is the other kind of evidence: graphs somebody built to answer a question that mattered to them, with numbers.

> **How to read the numbers in this chapter.** Some artefacts are **live and public**: you can open them and count for yourself. The rest are **parsed from their design documents**: the graph is real and complete in the brief, but nothing has been deployed. The two are never mixed, and every number below says which it is.

Twenty vaults, published with their read keys

The parent project publishes vaults whose read keys it has deliberately released. As fetched on 26 August 2026, the published list carries **twenty vaults**, each with its file count, size, commit count and shape. A read key is a capability handed out on purpose and it cannot become write access; every vault is audited before its key appears, and findings are published on the vault's page rather than filed away.

Six of them are graph work in the sense this book means:

Vault	Shape	Contents
Regulation Graph	reference data: regulation as an evidence graph	207 files · 14.9 MB · 2 commits · 11 app views
Agentic Browser Isolation	structured analysis: a living risk graph	104 files · 2.4 MB · 4 commits · 17 app entries
Risk Graph Explorer	application, public by design	33 files · 428 KB · 7 commits · 1 app entry
Standards Atlas — GDPR	a standard as a semantic graph	116 files · 6.3 MB · writes scoped to `feedback/`
VoiceDebrief	structured analysis: four apps, one vault	92 files · 1.2 MB · 18 commits
Risk Mandate	application: a software project in a vault	124 files · 1.9 MB · 98 commits · 8 app entries

Figure 13.1 · Six of the twenty published vaults, quoted from sgit.ai/demos/vaults/ as fetched on 26 August 2026. Every row is live and countable.

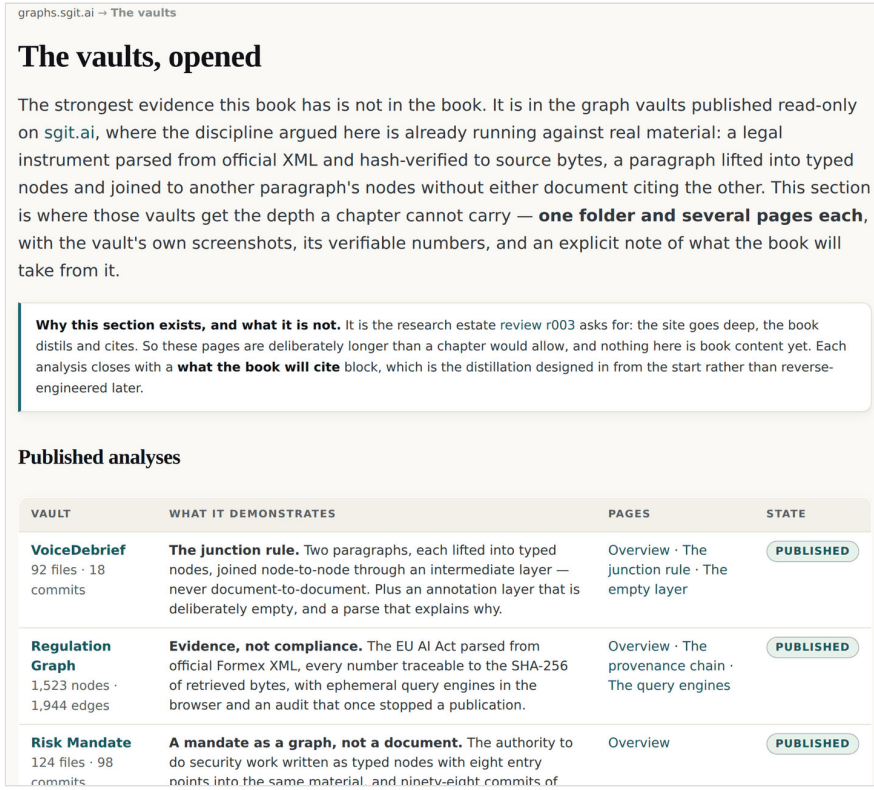

Figure 13.2 · How this book's companion site teaches from those vaults, at graphs.sgit.ai/v1/vaults/, site version v0.5.11. Each vault gets a chapter that explains what it does and links to the running thing, rather than a rebuilt copy of it.

The three that this book teaches from most often carry these numbers.

The EU AI Act regulation graph. 1,523 nodes and 1,944 edges: 113 articles, 500 paragraphs, 417 points, 180 recitals, 13 annexes, 68 definitions. Eleven views, including an article graph, a SQLite interface, an RDF/Turtle export and a lab with an interactive graph console. Parsed deterministically from the official Formex XML, with every element hash-verified to its source bytes.

That last clause is the one to notice. It is chapter nine's anchor discipline, applied to a law by somebody who had never read chapter nine, because both come from the same rule: if you cannot point at the bytes, you are asserting.

The Risk Graph Explorer. 18 facts, 37 risks and 14 provisions in its "Exposed" preset, with seven views recomputed simultaneously. Amber is exposure, green is assurance, and **ghosted edges are unanswered**. It requests `permissions: {}`, meaning no network and no storage, entirely client-side, which you can verify in a browser's network panel in about ten seconds.

Agentic browser isolation. 17 entry points, five stakeholder altitudes from IT to the board, acceptance-gated escalation **with no deny button**, and `fs.write: []`.

The question a graph answers that a slide cannot

The browser-isolation vault has a design brief behind it whose graph is 59 nodes and 75 edges, parsed from the brief. It answers one question: should an agent that browses and acts run inside the user's own browser, with their live signed-in sessions, or in an isolated one with a scoped identity?

Everybody can argue that in adjectives. Safer. More convenient. Enterprise-grade. Nobody wins those arguments and nothing is checkable afterwards.

The graph replaces the adjective with a node type. `AuthorizationClosure` is the transitive union of everything an identity can reach by following grants, including grants reached *through* other grants. Compute it for both options and subtract.

> "What isolation changes" stops being a claim and becomes a closure difference: these assets are reachable in option A and not in option B. A buyer can check it against their own estate.

Two details from that graph are worth stealing whole.

Seven owners, and the escalation has no escalator. The owners climb the whole organisation: IT, then the chief information security officer, then the chief financial officer, the chief operating officer and the data protection officer, then the chief executive, then the board. The escalation between altitudes is **a property of the edges, not of a workflow rule**. Nobody escalates anything. A risk arrives at the chief financial officer because the path from it leads there and nobody below has accepted it.

Three of the thirteen risks are created by the mitigation itself. A slide comparing two options never lists the harms of the option it recommends, because a slide has a direction. A graph does not: a risk node arising from a control node is the same shape as any other risk node, so it gets drawn, gets an owner, and needs accepting.

Honesty here is not a discipline you have to remember. It is what the structure produces unless you go out of your way to suppress it.

One configuration fact, carried to the board

The 2FA instance graph is the only artefact in the corpus that is **both a complete narrative and a machine-readable file**: 51 nodes and 53 edges plus an acceptances block, as one standalone parseable JSON file, with 22 node classes and 34 edge-type rows in its companion brief. Parsed from the brief, not deployed.

It starts from one fact: two administrator accounts without two-factor authentication. It ends at the board and the regulator.

```
[2FA not enforced] -backed_by->     [config export]
        │
        └─ -gives_rise_to-> [credential stuffing]
                            -impairs-> [customer data]
                            -gives_rise_to-> [regulatory exposure]
                            -accepted_by-> [the wrong person]
```

Figure 13.3 · The chain, walked as a sentence. The finding is the last hop.

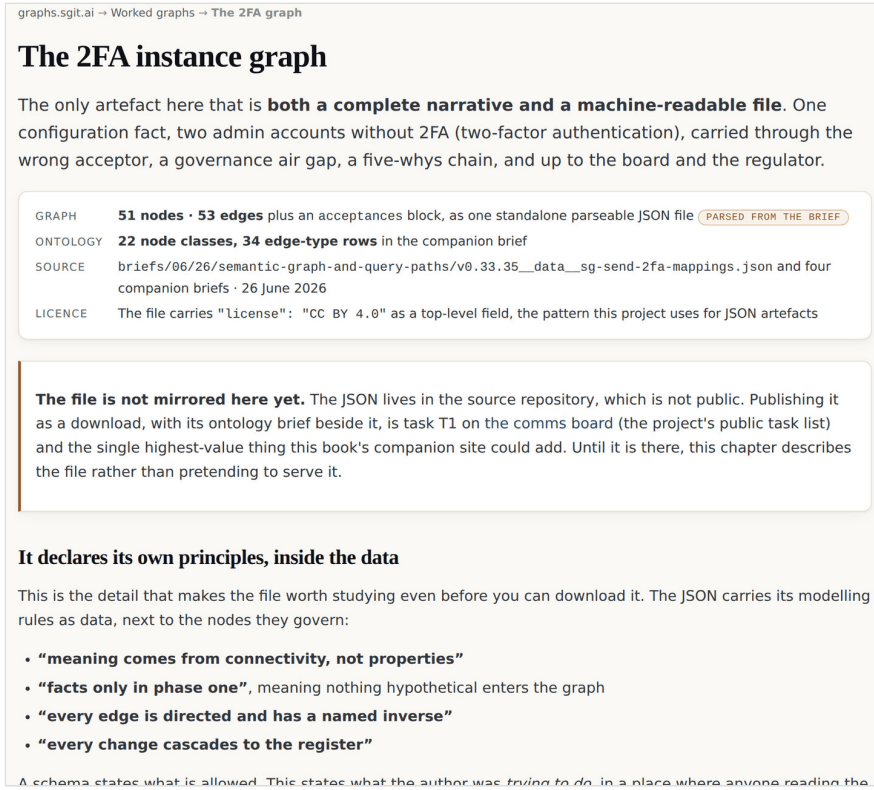

graphs.sgit.ai → Worked graphs → **The 2FA graph**

The 2FA instance graph

The only artefact here that is **both a complete narrative and a machine-readable file**. One configuration fact, two admin accounts without 2FA (two-factor authentication), carried through the wrong acceptor, a governance air gap, a five-whys chain, and up to the board and the regulator.

GRAPH	**51 nodes · 53 edges** plus an acceptances block, as one standalone parseable JSON file PARSED FROM THE BRIEF
ONTOLOGY	**22 node classes, 34 edge-type rows** in the companion brief
SOURCE	`briefs/06/26/semantic-graph-and-query-paths/v0.33.35__data__sg-send-2fa-mappings.json` and four companion briefs · 26 June 2026
LICENCE	The file carries `"license": "CC BY 4.0"` as a top-level field, the pattern this project uses for JSON artefacts

> **The file is not mirrored here yet.** The JSON lives in the source repository, which is not public. Publishing it as a download, with its ontology brief beside it, is task T1 on the comms board (the project's public task list) and the single highest-value thing this book's companion site could add. Until it is there, this chapter describes the file rather than pretending to serve it.

It declares its own principles, inside the data

This is the detail that makes the file worth studying even before you can download it. The JSON carries its modelling rules as data, next to the nodes they govern:

- **“meaning comes from connectivity, not properties”**
- **“facts only in phase one”**, meaning nothing hypothetical enters the graph
- **“every edge is directed and has a named inverse”**
- **“every change cascades to the register”**

A schema states what is allowed. This states what the author was *trying to do*, in a place where anyone reading the

Figure 13.4 · The 2FA instance graph as published at graphs.sgit.ai/v1/examples/2fa.html, site version v0.5.11, with its node classes, its six interval nodes and the modelling principles the file declares inside its own data.

The finding is the last hop. The risk *was* accepted, by somebody without the authority to accept it. A register records that an acceptance exists. The graph records who, and lets you ask whether that person's authority reaches this impact.

Three more things this file does that are worth copying.

It declares its own principles, inside the data. Four lines of modelling rules sit next to the nodes they govern: *"meaning comes from connectivity, not properties"*, *"facts only in phase one"*, *"every edge is directed and has a named inverse"*, *"every change cascades to the register"*. A schema states what is allowed. This states what the author was trying to do, in the place where anybody reading the data will see it, and it costs four lines.

There is no deny button. Six `Interval` nodes: 1 hour, 4 hours, 2 days, 2 weeks, 1 month, 6 months. A real risk is never denied. It is **accepted for an interval**, by a named person, after which it returns.

	Deny	Accept for an interval
What it feels like	An argument you have to win	A decision you can make today
What it records	Nothing; the risk stays open and unowned	Who accepted it, at what altitude, until when
What happens next	It is re-raised by whoever cares most, eventually	It returns on a date, automatically, to the same person
Failure mode	Risks accumulate in a register nobody reads	Somebody has to keep choosing an interval, which is visible

It shows a governance air gap. One risk is not connected to the register at all. Not denied, not accepted: unconnected. Chapter seven's air gap, in a real instance.

The output is the five questions nobody could answer

The Article 26(5) exercise takes one EU AI Act provision and one concrete deployment (a creditworthiness agent), and carries it from a running system to a board decision and back down. Its inventory: 1 reality, 1 twin, 8 facts (one deliberately unevidenced), 7 pieces of evidence (one deliberately absent), 5 provisions, 3 vulnerabilities, 5 risks, 4 stakeholder altitudes, 3 decisions (and a fourth deliberately absent), **9 questions of which 5 are unanswered**, and 2 projects (one unfunded). The organisation is invented and every invented element is marked as invented; the Act provisions are verified against five external sources.

The brief calls those five unanswered questions *"the actual output of the exercise"*.

And it produces the single sharpest table in the corpus. Four quadrants: accepted or not, against acceptable or not.

	Acceptable	Not acceptable
Accepted	Fine. This is what governance is supposed to produce.	A known bad decision, on the record, with a name on it. Rare and survivable.
Not accepted	empty	empty

The bottom row is empty, and the emptiness is the finding. There is no mechanism by which a risk gets to be *not accepted*. Nothing is denied; things simply are not accepted, silently, by nobody, forever. You cannot see that in a risk register, because a register has rows and an absent row looks like nothing at all.

The line from the source, on what the graph shows that a register cannot:

> R3 reaches the CFO because **nobody accepted it** — not because anybody raised it.

Two more, briefly

Browser extensions and the read-content closure. "Allow this extension to read the content of the pages you visit" quietly grants the authorization closure of every site you are currently logged into. The demonstration is that the graph is queryable in both directions: walk *out* from the extension to the cloud console, the email and the customer database it reaches, or walk *back in* from "how could my email be attacked?" to the extensions that expose it. Two different tables. One graph. Designed, not deployed.

The issue tracker's own configuration. 12 node types, 10 verb and inverse edge types with domain and range constraints, **71 nodes and 141 edges** across 107 issue files, edges stored bidirectionally. Live repository data, not a design. It is the cheapest credibility in this book and the most instructive artefact in it, for two reasons: it shows that domain and range constraints on a verb pair are enough structure to be useful, and it ships the banned generic edge, which is what a rule looks like when nothing enforces it.

Nineteen sites, and what a network demonstrates

Chapter two introduced the network as early evidence. Here is what it does and does not show.

As fetched on 26 August 2026, the sgit.ai network lists **nineteen sites, eighteen live**. They share a design and a discipline (sourced claims, a stated status, honest edges) and they publish their arguments *before* the things they describe exist, so the commitments stay checkable afterwards. Several of them are chapters of this book relocated into another domain's vocabulary: `standards.sgit.ai` is chapter five's anchoring rule as a citation scheme; `risks.sgit.ai` is the acceptance-interval mechanic with personal liability attached and **a stated zero lines of code implementing any of it**; `newsroom.sgit.ai` is chapter eight's projection claim applied to journalism, costed at £8.40 for a worked story, and marked *not built*.

What that demonstrates: the grammar travels. Nineteen independent arguments, no shared schema, connected by named links, each owning its own vocabulary. That is chapter four running at the scale of a publishing estate rather than a data model.

What it does not demonstrate: that any of it is in production somewhere else. These are this family's own sites. A grammar that travels between projects run by the same people is weaker evidence than a grammar that travels between strangers, and this book will not inflate it.

Spreadsheets of spreadsheets

The last item in this chapter is not an artefact. It is a sentence, and it belongs here because it is the fastest way to make the whole argument land with somebody who has never drawn a graph.

Ask a person from finance what graphs of graphs means and you will lose them. Tell them it is **spreadsheets of spreadsheets** and you will not, because finance genuinely nests workbooks inside

workbooks inside consolidation packs, and everyone in that room has lived it. The audit working-paper hierarchy is the fractal principle: every working paper opens into schedules with their own working papers, the same review structure at every zoom.

The estate turned that into a register (chapter four) so the machine can do it too. But the sentence works without any machine, and it is the one to carry into a meeting.

> **Where the live estate demonstrates this.** The published vault list, with a file count, a size and a commit count per row, is at `sgit.ai/demos/vaults/index.html`. The three graph vaults this book teaches from are at `sgit.ai/demos/vaults/regulation-graph/`, `/risk-graph-explorer/` and `/agentic-browser-isolation/`. The worked examples with their node and edge counts are at `graphs.sgit.ai/v1/examples/`. The network index is at `sgit.ai/network/index.html`.

14 · What ships, what is argued

After this chapter you will know exactly which claims in this book are backed by running code you can read, which are backed by a design document, and which are backed by nothing at all. Read it before you cite anything else in here.

This chapter is non-negotiable, and it is the first one to read if you are deciding whether to trust the rest. This project's credibility rests on separating design from delivery.

The first edition of this argument carried a chapter with this title and this job. Its headline sentence was blunt: *we ship a hand-written content-addressed object graph in the browser. We do not use a graph database, and we say so in our own architecture notes.*

That sentence is still true. What has changed since is worth being precise about, because the honest summary of the last six months is **the machinery around the argument got real while the semantic layer stayed a design**.

Ships, and is verifiable by reading code

Everything in this section is running and checkable by somebody with the repository open or a browser pointed at the estate.

Carried over from the first edition

What	The detail that makes it checkable
The vault commit graph	Content-addressed objects (the identifier is a SHA-256 hash of the ciphertext), multi-parent commits, a tree per directory, deterministic refs derived by keyed hashing, a real merge base computed by breadth-first search over all parents, and three-way merge.
A graph of graphs	Typed `*.link.json` edges between vaults, optionally pinned to a specific commit in the target's history. A cross-graph edge that cannot silently follow a moving target.
A read-only query interface over that graph, exposed to untrusted sandboxed apps	`sg.history.log / list / read / readText / readBlob`. A graph query surface handed to code you do not trust.
A live typed property graph	The issue tracker's own data: 12 node types, 10 verb and inverse edge types with domain and range constraints, **71 nodes and 141 edges** across 107 issue files, edges stored bidirectionally. Not a design; repository data.
Published vaults	Three in the first edition. **Twenty** on the published list as fetched on 26 August 2026, each with its file count, size and commit count.

New since the first edition

What	The detail that makes it checkable
The extraction pipeline, on one document	57 nodes across five families, 8 asserted edges, **72 anchors and 73 verified spans**, each quote found verbatim inside its named section and resolved to byte offsets, re-verified on every release against a source that still hashes to its recorded SHA-256. An extraction citing words that are not there fails the build.
Coverage as a build property	Every prose-bearing section either yields an anchored item or is declared empty with a reason. Both a silent gap and a stale declaration fail the build.
The core graph	39 sections, 186 blocks, 342 sentences, 4,221 word instances, 143 markup spans and 951 distinct word forms, with structural deterministic identifiers at every level.
The two-way transform	The build rebuilds the source markdown from the formatting graph and **fails unless it is byte-identical**, and re-derives the semantic shards from the formatting graph alone, so the two graphs cannot disagree. Seven gates in one generator.
The identity ledger	225 identities (1 document, 38 sections, 186 blocks), minted once and carried across edits by a three-pass match-then-mint, retired rather than deleted, with a gate that runs the assignment twice over its own output and fails if anything changes.
The verbs register	Every asserted verb with its declared unique inverse. The build fails on an unregistered verb, on two verbs sharing an inverse, and on an undeclared self-inverse.
The WCLM	Twelve operators in twelve folders, six data types, every weight a formula written into the world file, 39 recorded example vectors replayed on every build, and the wire adjacency property machine-checked (551 wires, zero layer jumps at the release that proved it).
The code anatomies	One per operator, anchored to the exact text of each segment's first line, with line ranges that must tile the file completely and a gate that fails the release when code and anatomy drift.
The unit suite	`node admin/tests/universe.test.mjs` reports **84 passed, 0 failed** at site version v0.5.11. It was run while writing this chapter.
The release history as a record	Every release row narrates what happened, and the build fails if the version table has no row for the current version or lists any version twice.

The best description of the shipped layer remains one the project applied to itself: *"what we've built is not fundamentally an encryption system. It is a content-addressed, portable, storage-agnostic version control protocol."* A commit history is a graph, and it is the one graph here that has been running for months.

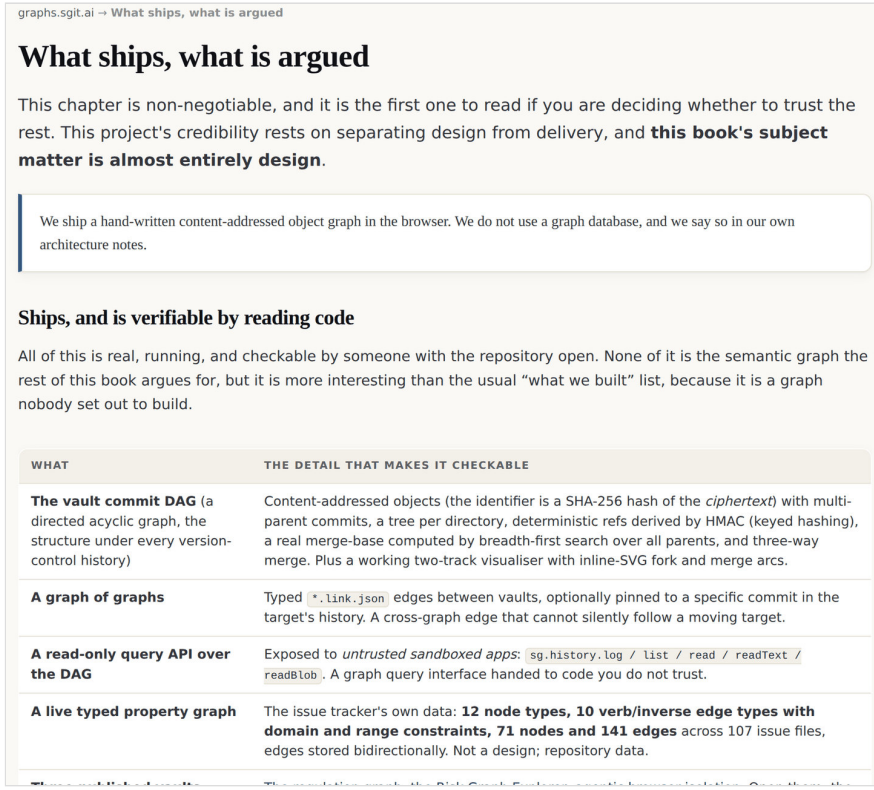

graphs.sgit.ai → What ships, what is argued

What ships, what is argued

This chapter is non-negotiable, and it is the first one to read if you are deciding whether to trust the rest. This project's credibility rests on separating design from delivery, and **this book's subject matter is almost entirely design**.

> We ship a hand-written content-addressed object graph in the browser. We do not use a graph database, and we say so in our own architecture notes.

Ships, and is verifiable by reading code

All of this is real, running, and checkable by someone with the repository open. None of it is the semantic graph the rest of this book argues for, but it is more interesting than the usual “what we built” list, because it is a graph nobody set out to build.

WHAT	THE DETAIL THAT MAKES IT CHECKABLE
The vault commit DAG (a directed acyclic graph, the structure under every version-control history)	Content-addressed objects (the identifier is a SHA-256 hash of the *ciphertext*) with multi-parent commits, a tree per directory, deterministic refs derived by HMAC (keyed hashing), a real merge-base computed by breadth-first search over all parents, and three-way merge. Plus a working two-track visualiser with inline-SVG fork and merge arcs.
A graph of graphs	Typed `*.link.json` edges between vaults, optionally pinned to a specific commit in the target's history. A cross-graph edge that cannot silently follow a moving target.
A read-only query API over the DAG	Exposed to *untrusted sandboxed apps*: `sg.history.log / list / read / readText / readBlob`. A graph query interface handed to code you do not trust.
A live typed property graph	The issue tracker's own data: **12 node types, 10 verb/inverse edge types with domain and range constraints, 71 nodes and 141 edges** across 107 issue files, edges stored bidirectionally. Not a design; repository data.

Figure 14.1 · The first edition's version of this chapter, still published at graphs.sgit.ai/v1/shipped/, site version v0.5.11. It is reproduced here rather than paraphrased because the point of the page is that it exists at all: a project that publishes what it has not built is making a checkable commitment.

Argued, and published as argument

Nearly everything else. The node type formulas, the grounding ladder, ontologies of ontologies, the semantic risk ontology, a graph at every boundary, twins as endpoints, the path query language, decompilation: all of it is **proposed**. It is published because publishing a design before it is built is how it gets checked, and because several of these ideas have been applied by hand to real problems even though no system implements them.

Where an idea has been applied by hand, this book says so and marks the numbers as parsed from a design document rather than live.

Does not exist anywhere

Searched for, and absent. If you read something in this book that seems to imply otherwise, the book is wrong.

- **Any graph database at all.** MGraph-DB is named repeatedly in the corpus and deferred every time: *"there is even a graph database, MGraph-DB, we could use, but for now let's keep it simple."* File-based won.
- **Browser SPARQL or Cypher**, the query languages of the RDF and property-graph worlds.
- **RDF or JSON-LD serialisation in the code.** The live regulation graph exports RDF/Turtle as a published artefact; there is no RDF layer in the codebase.
- **The semantic risk ontology as a schema file.** It exists as prose in briefs.
- **The path-query language.** Seventeen queries across five tiers are written down. Nothing executes them.
- **Commit signing.** `commit_v2.signature` is written on every commit and only ever set to `null`.

And what part four does not yet have

This is the section the first edition could not write because part four did not exist. It is the most important list in this chapter for a reader deciding how much weight to put on chapters nine to twelve.

- **Twenty of twenty-one source documents are unextracted.** One has been through the pipeline. The method is written down; the corpus is not processed.
- **Layer 2 does not exist.** The bridge layer, cross-document edges authored with a note, cannot start until there are at least two extracted documents. Every claim in this book about how divergence between documents is preserved describes a design, not a running thing.
- **Layer 3 does not exist.** The book's own universe, connecting down into layers 1 and 2 as evidence, is planned and unbuilt. This book anchors directly to the corpus instead, which is what the commission instructed.
- **Ask-a-document is not built.** Bringing a graph to a document and hearing agree, disagree or provide evidence needs a second extracted document. The single-document seed of it (negation checked against the world, the contradiction said out loud) does run.
- **No layer uses a language model.** The design for an operator that hands its graph to a model and takes a graph back is recorded, with the argument for why the typed contract makes it safe (both graphs are kept as evidence). It is not built, and when it is, it will run at generation time and never in a published page.
- **The identity ledger has never survived an edit.** The pilot's source is frozen, so all 225 identities are live and none has been retired. The machinery for carrying identity across a revision is written, gated and deterministic, and untested against a real revision.

- **Nobody has reviewed the extraction end to end.** It ships with a printed review pack whose contract says items you do not mention are taken as agreed. That contract has not yet been exercised by a full pass.

Four corrections this book carries rather than repeats

A book arguing that corrections must propagate had better propagate its own.

1 · Object identifiers are a hash of the ciphertext, not the plaintext. A skill file in the source repository states "SHA-256 of plaintext". The code hashes the ciphertext. The difference matters to anybody reasoning about what the object store leaks. Stated correctly here; the source file still carries the error.

2 · The banned edge is in the shipped configuration. The project's ontology brief forbids a generic association edge. Its own `link-types.json` ships one as a self-inverse pair, and a single edge instance uses it. Narrated in chapter three rather than quietly fixed, because it is a better teaching moment than the rule is.

3 · Several inverse edge names in chapter three are proposed by this book, not quoted from the corpus. They are marked in the table. If this book becomes the place people cite for that vocabulary, the distinction between quoted and proposed has to survive, or we will have done the thing we are warning about.

4 · The first edition's definition of "fractal" was a stretched reading of its own source. It defined fractal as one grammar at every level, which is uniformity. The foundational document licenses local vocabulary and override, which is composition. Caught by the usage ledger on the day that ledger first ran, corrected in the lexicon, with the superseded definition kept visible. Chapter six carries the consequence, and reports its own test failing rather than restating the correction as a pass.

Two things we would like to ship and have not

The personal risk question graph. Six questions, each answer typed as fact, opinion, hypothesis or evidence, and you watch your own risk graph build itself. Browser storage only, no backend, no account, no language model required. It is the best interactive demo in the material, it is a task on the public board, and it is not built. Saying so is cheaper than implying it exists.

A second extracted document. Everything queued behind layer 2 waits on it, and the honest reason it has not happened is that the estate kept building instruments instead. That is a real judgment call with a real cost, and it is recorded here rather than in a retrospective nobody reads.

Why this chapter exists at all

Because without it the book over-claims, and an over-claiming book about provenance discipline is self-refuting.

The graph work described here is mostly a design published in advance so that it can be checked against whatever eventually ships. That is a weaker claim than most publications make, and it is the one that is true.

> **Where the live estate demonstrates this.** The first edition's version of this chapter is at `graphs.sgit.ai/v1/shipped/`. The release history, one narrated row per release, is at `graphs.sgit.ai/admin/versions.html` for the current era, with the previous eras kept whole at `versions-v0.4.html` and `versions-earlier.html`. The unit suite is `admin/tests/universe.test.mjs` and the release gate is `admin/build/validate.js`.

15 · Your first graph, tomorrow

After this chapter you will have a concrete plan for an afternoon, a week and a month, and a short list of the mistakes that will otherwise cost you the first two of those.

Everything in this book is worth nothing to you unless you draw something. So this chapter is instructions.

You need no software you do not already have. A text editor and a version control system are enough for the afternoon. You will not need a graph database, and if somebody tells you otherwise in the first month, they are solving a different problem from yours.

Before you start: pick the right question

The most common way this fails is picking a subject instead of a question.

"Let's map our architecture" is a subject. It has no stopping point, no test for correctness and no audience, so it becomes a diagram nobody reads. Pick instead a question somebody is currently answering badly, by hand, repeatedly.

Good first questions have three properties: somebody actually asks them, the current answer requires a person to remember something, and the answer changes when facts change.

A good first question	Why it works
"If this service goes down, who has to be told?"	The answer lives in three people's heads and is wrong twice a year.
"What does this third-party integration actually reach?"	Everybody assumes a smaller answer than the true one.
"Which of our controls does this obligation rely on?"	Currently a spreadsheet that goes stale between audits.
"What breaks if we deprecate this?"	The current method is asking on a group chat.

Write the question down. It is the acceptance test for everything that follows, and you will be tempted to widen it by Wednesday.

The afternoon

One. Write your question at the top of a file.

Two. List the things the question touches. Not types: actual things, by name. The payments service. Priya. The data protection officer. Article 26(5). The staging database. Aim for fifteen to forty. If you have four, the question is too small. If you have three hundred, you picked a subject.

Three. Write the edges as sentences, in full, in English, one per line. Not `svc-a → team-b`. Write:

```
the payments service is owned by the platform team
the platform team has stakeholder the head of engineering
the head of engineering reports to the chief technology officer
the payments service reads the customer database
the customer database contains personal data
```

Sentences first, always. You are not being slow; you are being forced to name the verb, and naming the verb is the whole discipline.

Four. Turn each sentence into a verb and its inverse. `owned_by ⇄ owns`. `has_stakeholder ⇄ stakeholder_in`. `reads ⇄ read_by`. `contains ⇄ contained_in`. Write both names down even though you will only store one direction. If the inverse is just the same sentence backwards, you have one relationship where you thought you had two: go back and find the direction with lower fan-out.

Five. Read your paths aloud. Start at any node and walk three or four hops. If the sentence does not survive being said out loud in front of somebody from the business, the edges are wrong. Fix them now, while it is five minutes of work.

Six. Ask your question of the graph, by hand, tracing with a finger. This is the moment the afternoon either pays off or does not. If tracing produces an answer that somebody would have got wrong from memory, you have something. If it produces nothing, look at which hop was missing: that is your first real finding.

Seven. Write down what you could not answer. Not in your head. As nodes, with names.

By the end of the afternoon you should have a text file of forty-odd lines, an answer, and a list of things nobody knows. That is a complete first graph and it is more than most organisations have.

The first week

Five moves, in order of how much they pay back.

Give everything an identity. Short, opaque, minted once, never reused. Not the name, because names change. Not the position in the file, because positions move. Chapter ten is the whole argument; the one-line version is that a reference which holds a position will break silently the first time somebody fixes a typo.

Put it in version control. Not because you need history yet, but because the diff is the review surface. A change to a graph should be reviewable the way a change to code is, and this is what makes the artefact and the reasoning behind it live together.

Add the absences you have been carrying. Every "we don't know who owns that" becomes a node. Every "the data comes from a spreadsheet somebody updates on Thursdays" becomes an air

gap with an owner and a frequency. This is the step people skip, and it is the step that makes the graph worth funding, because now the missing dots have names.

Point at two or three anchors. A published standard. A named regulation. A library at a pinned version. Your own chart of accounts. Do not adopt them; point at them, with a verb that says how strongly. `similar_to`, not `is_a`.

Write down one node type formula. Just one. Pick the classification your team argues about most and express it as a required pattern of paths: *a thing is a Fact only if it has a downward edge to evidence*. Then run it over what you have and see what fails. The failures are the point.

The first month

Three things to build, and one to resist.

Build a second graph, from a different angle, on the same subject. The comparison is where the finding is. Two graphs of the same system that disagree have told you something neither could tell you alone, and chapter four is why you should keep both rather than merging them.

Build one projection. Take your graph and render a document from it: the register your audit wants, the on-call sheet, the one-page summary for the board. The test of chapter eight is whether you can throw the rendered document away and regenerate it. If you cannot, the graph is not yet the truth and the document still is.

Start measuring the absences. How many nodes terminate at something real? How many risks are connected to the register? How many obligations reach a control? Those ratios are your coverage, and they are honest in a way a progress bar is not.

Resist buying a graph database. Not because they are bad. Because the thing that will kill your project in month one is not query performance; it is that the edges are wrong, and a database will not tell you that. Buy one when a query you actually need is too slow, and not before.

Ten rules on one page

```
 1 · every edge is a verb, and both directions get a name
 2 · no generic association edge, ever, not even temporarily
 3 · if the path does not read as a sentence in the reader's own
     language, the edges are wrong
 4 · rich nodes are good: solve the picture at query time, never by
     deleting relationships
 5 · never render the whole graph; render the result of a query
 6 · point at anchors, do not become them
 7 · do not merge two vocabularies; keep both and declare bridges
 8 · a type is a required path-pattern, not a label somebody applied
 9 · supersede, never delete, and weight by independence not by count
10 · draw the absences: an unanswered question is a node
```

Figure 15.1 · The rules of this book on one page. If you copy one thing out of it, copy this.

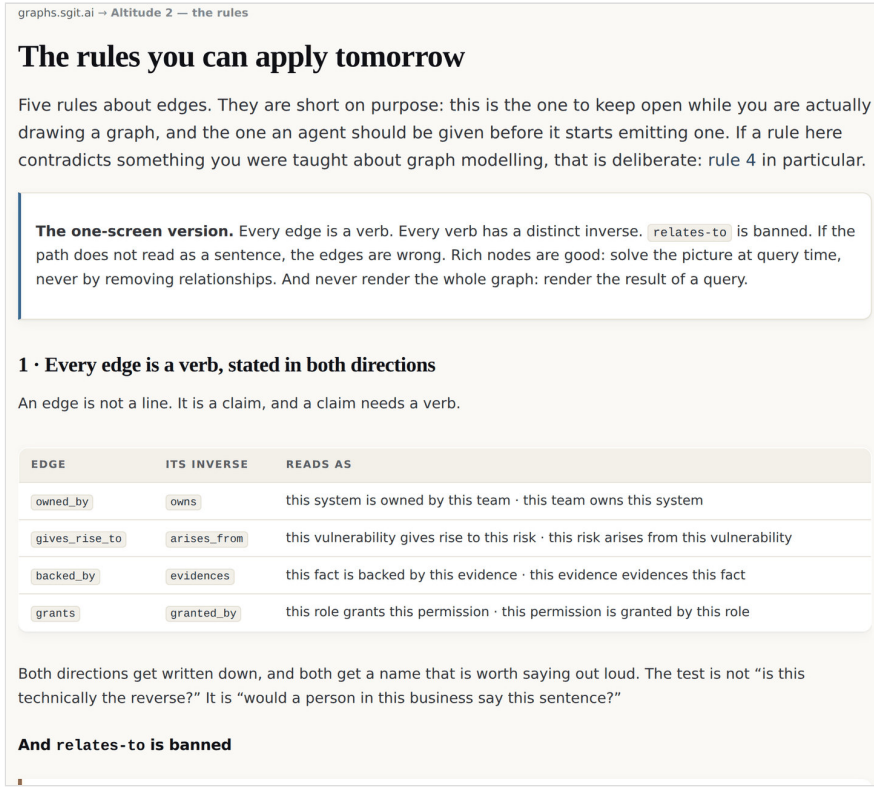

Figure 15.2 · The same rules as the companion site publishes them, at graphs.sgit.ai/v1/grammar/, site version v0.5.11. The one-screen version in the box at the top is the form to keep open while you are drawing.

Six ways this goes wrong

Each of these has happened, more than once, and each is cheap to avoid if you are watching for it.

You draw the whole thing and show it to somebody. They see a hairball, conclude the approach is unserious, and you do not get a second meeting. Never show the whole graph. Show one node with its neighbours, or the result of a question.

You reach for a generic edge because you are in a hurry. It survives, because nothing forces the decision. Six months later it is in forty places and every query is worse. This estate ships one in its own configuration and names it in public rather than quietly fixing it, because the temptation is universal.

You start adding validation rules. Somebody's data does not fit, so you add a constraint, so they work around it, so their data is now conforming and false. Add an edge instead.

You model hypotheticals. A risk that has not happened, a system that might exist, an integration somebody proposed. The graph fills up with a brainstorm and stops being a record of reality. Facts only.

You put the meaning in a property. `criticality: "high"` is a judgment somebody made, sitting in a field, unattributed and unversioned. Make it a path: what makes it high, and who says so. A property may hold a timestamp; it may not hold the answer to "what kind of thing is this?"

You aim for completeness. The graph will be deep where the work is and absent everywhere else. That is not a defect to apologise for; it is the property that makes the project finite. A graph that had to be complete before it was useful would never be either.

What to do with an agent

Most readers of this book will hand some of this to a language model, so here is the version to hand it.

Give it the vocabulary before the task. The edge set in the reference card at the back of this book is written to be pasted into a session. An agent asked to "build a graph" with no vocabulary will invent one, and it will invent a generic association edge within about five nodes.

Ask for a proposal, not a commitment. The model proposes a graph; you validate it. That is the boundary of chapter seven, applied at the smallest possible scale. In practice it means: have the agent emit the graph as data, then check it against your own rules before anything consumes it.

Make it anchor. Every node it emits should carry the sentence it came from, verbatim. Then check the quotes exist. Chapter nine's rule is the single highest-value thing you can adopt from this book when working with a model: **a citation that is not in the source should fail mechanically, and never reach a human reviewer.**

Make it emit absences. When it cannot determine something, it should emit a node saying so rather than a plausible value. Ask for this explicitly and it will do it; do not ask, and it will fill the gap.

The one-paragraph version, for the colleague who will not read the book

A node connected to nothing means nothing. What a thing is emerges from the edges you can trace from it, and how much you can rely on that meaning depends on how richly and how independently it is connected. So write your relationships as verbs with named inverses, never merge two vocabularies (keep both and bridge them, because the disagreement is usually the finding), express your classifications as required paths rather than as labels somebody applied, never delete a superseded claim, and draw what you do not know as nodes with names. Do that and meaning stops being something you assert and becomes something you can compute, check, version and argue with.

That is the whole book. Everything else is detail, evidence and honesty about what does not work yet.

Where the live estate demonstrates this. The rules are at `graphs.sgit.ai/v1/grammar/` with the vocabulary at `edge-set.html`. Worked graphs with real numbers are at `graphs.sgit.ai/v1/examples/`, and the three public vaults you can open and count are at `sgit.ai/demos/vaults/`. If you build something and it disagrees with this book, the estate's comms board is at `graphs.sgit.ai/admin/comms.html`, and a correction that changes a claim gets a row in the release history rather than a silent edit.

Colophon: what was cut, and what remains open

This book argues that a named absence beats a hidden one. Here are its own.

How this book was made

Written in one session by an AI agent working from a written commission, in the estate's own working style: the markdown chapters under `v2/books/fsg/content/` are the source of truth, the web pages render that markdown in the browser so a page cannot drift from the file it claims to render, and the printed version is generated from the same files by the same build.

The corpus read for it, in the order the commission specified: the first edition's seventeen chapter sources in full; the twenty-one frozen source documents, with four read end to end; the nineteen founder memos of the second edition, verbatim, with the agent readings that accompany them; the second edition's working surface, meaning the universe extraction and its pipeline, the core graph, the identity ledger, the twelve operators and their data; the release history, one narrated row per release; and two live fetches of the external indexes, the published vault list and the network of sites, both made on 26 August 2026.

Every number in this book was computed from a file in the repository, quoted from a page that was fetched, or read off a build that was run while writing. The unit suite was run (84 passed, 0 failed). Where a figure could not be computed, the claim was not made.

The screenshots were taken from the real pages with the repository's own harness: headless Chromium driven over the debug protocol against a local server, at site version v0.5.11. Nothing was mocked and nothing was drawn from memory. The structural diagrams are ascii art, which is diffable, reviewable in a pull request and prints without a rendering step.

What was cut

A chapter on time. The corpus repeatedly says the graph moves and never explains how. It was cut rather than written thin, because the material to write it does not exist: "time is an event, things change" is close to the whole treatment. The adjacent pieces are named in chapter seven so a future writer can find them, and the gap is recorded here rather than smoothed over.

A chapter on Wardley maps. The first edition carries one, and it is good. It was cut from this edition because its best material (a map is a falsifiable claim, and the coordinate trap that will bite you on day one) is one chapter's worth of a different book, and there is now a whole site about it. The single sharpest idea in it, that a gap has no evolution so what you plot is the labour that fills it, survives in chapter seven.

A chapter on origins. The first edition's ten-phase history of how this thinking arrived is genuinely useful and belongs to a different book, the one being written in parallel about how these books get made. Repeating it here would have been the second edition going again rather than further.

The full worked graph of the identity and access management example. Six layers, some thirty-one node types, twenty edge types and seven formulas: too large to render usefully in print, and rendering it badly would have violated the book's own rule about never showing the whole graph. It appears in chapters two and five as the source of the authorization closure, and it is published in full elsewhere.

A comparison table against named commercial products. Tempting, and dishonest, because this estate has not run any of them at scale. Chapter two positions the argument against approaches rather than against vendors, which is the comparison this book can actually support.

Every appendix except one. A glossary, a source manifest, a gap catalogue and a list of the estate's methods were all drafted and dropped. The reference card that remains is the one a reader on a plane can use. The rest are published on the web where they can be kept current, which is where reference material belongs.

What remains open

The questions this book could not settle, stated as questions rather than hidden as omissions.

How does a graph change over time? Unanswered, as above. The pieces exist and the synthesis does not.

Does the method survive twenty more documents? One document has been through the extraction pipeline. The pipeline's discipline (anchors that fail the build, coverage total by construction, ratings that must be signed) is designed to scale and has not been scaled. The honest state is one thorough demonstration.

Is the storage-layer failure of the fractal test worth fixing? Chapter six reports it: three levels of the same zoom use three serialisations. There is a decent engineering reason and nobody has decided whether the cost of removing it is worth paying. This book took the strict reading of the test and reported the failure; a reader who takes the composition reading (chapter nine, the stretched-use finding) would score it differently, and that disagreement is unresolved.

Where does a language model belong in the pipeline? The design says: at the edge, proposing, with both graphs kept as evidence. Nothing implements it. The first thing that does will find out whether the typed contract really is enough, or whether the interesting failures happen somewhere the types do not reach.

Can an engine like the one in chapter eleven say anything useful at corpus scale? Its world is 951 tokens from one document. Everything about it is designed to compose across documents, because tokens are content addresses and need no registry. Nobody has tried.

Does the usage maturity model work when the author is not the rater? Every rating in the ledger so far was made by the agent that built it. The model's own principle says the source's author

should ideally rate the uses. That has not happened, and the first time an author disagrees with a rating will be more informative than everything written about it here.

Is "not a graph database pitch" still true if any of this ships? It is true today. A reader in two years should check whether the sentence survived contact with a product, and this paragraph exists so that the check has something to point at.

Where this book might be wrong

Three places to attack it first, offered because a book with no stated weak points is asking to be believed rather than read.

The positioning in chapter two is written fresh. The corpus holds a strong implicit position and never engages the named field. If you know that literature well, that section is where you will find this book weakest, and the estate's comms board is where to say so.

Several inverse edge names in chapter three are proposals. They are marked, and marking is not the same as being right. Some of them will be bad names, and the way to find out is for somebody to try to say them out loud in their own business.

The engine of chapters eleven and twelve is one afternoon's world. It is genuinely deterministic, genuinely explainable, and genuinely tiny. Every claim about what its architecture makes possible is a claim about an architecture, not a result at scale. If you read those chapters as evidence that this approach works, you have read them harder than they can bear. Read them as evidence that it is *buildable*, which is what they show.

The figures

Every figure in this book, with the page it came from and the version it was taken at.

Figure	Source
1.1 The sense switch	`graphs.sgit.ai/v2/wclm/`, v0.5.11, prompt "graphs of graphs", sense switched to a chart of data
1.2 The confidence ladder	drawn, from `graphs.sgit.ai/v1/start/#confidence`
2.1 The four-step case	drawn, from `graphs.sgit.ai/v1/why-graphs/`
2.2 The four losses	drawn, from `graphs.sgit.ai/v1/about/participant.html`
2.3 Seven of the nineteen sites	`sgit.ai/network/index.html`, fetched 26 August 2026
2.4 The three uses of graph	`graphs.sgit.ai/v1/why-graphs/`, v0.5.11
3.1 Build wide, find the few, flip	drawn, from `graphs.sgit.ai/v1/grammar/#blob`
3.2 The edge set	`graphs.sgit.ai/v1/grammar/edge-set.html`, v0.5.11
4.1 The three layers	drawn, from `graphs.sgit.ai/v1/depth/#ontologies`
4.2 Four analogies	`v2/wclm/analogies.json`, quoted verbatim
5.1 The grounding ladder	drawn, from `graphs.sgit.ai/v1/depth/#ladder`
5.2 A finding that is arithmetic	drawn, from `graphs.sgit.ai/v1/examples/article-26-5.html`
5.3 The claims table	`graphs.sgit.ai/v2/universe/thinking-in-graphs.html`, v0.5.11
6.1 The core graph ladder	drawn, from `v2/universe/data/core/thinking-in-graphs/index.json`
6.2 The bind operator's page	`graphs.sgit.ai/v2/wclm/operators/`, v0.5.11
7.1 The stack	redrawn from `v1/docs/sources/fractal-semantic-graphs.md`, 12 July 2026
7.2 The air gap as a node	drawn, from `graphs.sgit.ai/v1/depth/boundaries.html#air-gap`
7.3 One typed boundary	`graphs.sgit.ai/v2/wclm/operators/bind/`, v0.5.11
8.1 Three problems, one move	drawn, from `graphs.sgit.ai/v1/depth/boundaries.html#projections`
8.2 The seven gates	quoted from `admin/build/gen_coregraph.py`
8.3 The document's files	`graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html`, v0.5.11

Figure	Source
9.1 The extraction pipeline	drawn, from `v2/universe/README.md`
9.2 The universe reader	`graphs.sgit.ai/v2/universe/thinking-in-graphs.html`, v0.5.11
9.3 The local graph	`graphs.sgit.ai/v2/universe/thinking-in-graphs.graph.html`, v0.5.11
10.1 Two addresses	drawn, from brief 31's recorded distinction
10.2 Match, then mint	drawn, from `admin/build/gen_coregraph.py`
10.3 The identity ledger	`graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html`, v0.5.11
11.1 The WCLM's layers	`graphs.sgit.ai/v2/wclm/`, v0.5.11, prompt "meaning through connectivity"
11.2 The twelve operators	`graphs.sgit.ai/v2/wclm/operators/`, v0.5.11
11.3 The three formulas	quoted from `v2/wclm/data/world.json` and the operator pages
12.1 The evidence trail	`graphs.sgit.ai/v2/wclm/`, v0.5.11, winning meaning selected
13.1 Six published vaults	`sgit.ai/demos/vaults/index.html`, fetched 26 August 2026
13.2 The vault chapters	`graphs.sgit.ai/v1/vaults/`, v0.5.11
13.3 The 2FA chain	drawn, from `graphs.sgit.ai/v1/examples/2fa.html`
13.4 The 2FA graph	`graphs.sgit.ai/v1/examples/2fa.html`, v0.5.11
14.1 What ships	`graphs.sgit.ai/v1/shipped/`, v0.5.11
15.1 Ten rules on one page	drawn
15.2 The rules page	`graphs.sgit.ai/v1/grammar/`, v0.5.11

Fifteen of the thirty-eight are screenshots of real pages, taken with the repository's own harness at site version v0.5.11. The rest are drawn, and each names what it was drawn from.

Thanks, and the standing invitation

This book stands on a corpus of more than 1,300 founder briefs recorded almost daily over six and a half months, on the first edition that distilled them, and on four days of building that turned an argument into machinery.

If something here is wrong, the fastest route is the estate's comms board or an issue on the repository. Corrections that change a claim get a row in the release history, not a silent edit. That is this book's own rule, and the point of writing it down is that it applies to the book.

Reference card

Four pages to read with nothing else open, and the one part of this book written to be pasted into an agent session.

The thesis, in one sentence

Everything is a graph; meaning is not declared but discovered through relationships; and confidence in that meaning is proportional to how richly a node is connected to other nodes that provide context.

The ten rules

1. Every edge is a verb, and both directions get a name.
2. No generic association edge, ever, not even temporarily.
3. If the path does not read as a sentence in the reader's own language, the edges are wrong.
4. Rich nodes are good: solve the picture at query time, never by deleting relationships.
5. Never render the whole graph; render the result of a query.
6. Point at anchors, do not become them.
7. Do not merge two vocabularies; keep both and declare bridges.
8. A type is a required path-pattern, not a label somebody applied.
9. Supersede, never delete, and weight by independence not by count.
10. Draw the absences: an unanswered question is a node.

The edge set

Fifteen verbs, each directed, each with a distinct inverse. Names in the left column are quoted from the corpus. Inverses marked *(p)* are proposed by this book rather than quoted, and are a starting point for disagreement.

Edge	Inverse	Reads as
`connected_to`	`connected_to` (symmetric, last resort)	A is connected to B
`observed_on`	`bears_observation` (p)	this evidence was observed on this system
`backed_by`	`evidences` (p)	this fact is backed by this evidence
`measured_by`	`measures` (p)	this fact is measured by this measure
`grants`	`granted_by` (p)	this role grants this capability
`reaches`	`reachable_from` (p)	this grant reaches this asset
`enables`	`enabled_by` (p)	this capability enables this action
`exposes`	`exposed_by` (p)	this fact exposes this blast radius
`gives_rise_to`	`arises_from`	this vulnerability gives rise to this risk
`protected_by`	`protects` (p)	this asset is protected by this control
`conditional_on`	`conditions` (p)	this control is conditional on this fact
`defeated_by`	`defeats` (p)	this control is defeated by this attack
`owned_by`	`owns`	this system is owned by this role
`accepted_by`	`accepted` (p)	this risk is accepted by this role
`underwritten_by`	`underwrites` (p)	this acceptance is underwritten by this role

Also used in the worked graphs, listed separately rather than folded in: `impairs ⇄ impaired_by`, `emits ⇄ emitted_by`.

To extend the set: a new edge needs a sentence a person in that business would say out loud; its inverse needs a *different* sentence; state which node types may sit at each end; never invent a generic association edge; and prefer adding an edge to adding a validation rule.

The confidence ladder

Rung	What it is
5	rich multi-hop connectivity: many independent paths lead to the same conclusion
4	edges to external references: a standard, a legal instrument, a versioned library, a fetchable URL
3	edges to anchor nodes: shared landmarks others also point at
2	edges to typed definitions: something else in the system constrains what this can be
1	a few local edges: internally coherent, means nothing outside
0	no edges: nothing can be checked, so nothing should be relied on

The honest output is not a score. It is three sentences: **we know X, we think Y, we cannot confirm Z**, each with a traceable reason. The remedy for a low rung is more edges, never more validation rules.

The grounding ladder

```
upward implies:  [Fact] -gives_rise_to-> [Vulnerability] -gives_rise_to-> [Risk]
                 each step up is an interpretation somebody is accountable for

downward grounds: [Fact] -backed_by-> [Evidence] -measured_by-> [Measure]
                  each step down asks for something more checkable
```

Measure is not the floor. The floor is the last node where going deeper would neither improve observability nor change a decision, and that is a stated judgment. This is one formula among possible others.

Node type formulas

```
[Fact]          := a node with a downward -backed_by-> [Evidence]
[Vulnerability] := a [Fact] that also has an upward -gives_rise_to-> [Risk]
```

The content of a node does not decide its type; its paths do. Two nodes with identical text can be different types because their edges differ.

The vocabulary, with the plain phrase first

Say this	Rather than	What it is
an idea	concept	The unit of meaning, independent of any language. This is the thing you are actually modelling.
a word for it	term	How one language happens to express a concept.
a filing tree	taxonomy	Broader and narrower, one parent. Points **upward**.
the kinds of thing, and how they connect	ontology	What types exist and how they may connect. Points **outward**.
a shared landmark	anchor node	A reference point several parties link to. Has no authority.
a rule that decides what something is	node type formula	A required pattern of paths a node either matches or does not.
same rules at every zoom level	fractal	Self-similarity, scale invariance, composition, recursion, **with no new format and no special case**.
a view, generated when needed	projection	A document rendered from the graph rather than stored.
the real system this stands for	twin	Where the graph stops modelling and continues into something real.
the place it does not reach, named	air gap	A tracked, owned gap where no connection exists.
what else this touches	blast radius	Everything reachable once a thing goes wrong. A closure, not a list.
everything it can actually reach	authorization closure	The transitive union of every permission reachable through other permissions.
add edges, don't add rules	enrichment, not enforcement	The remedy for low confidence is more connections, never more validation.
lifting text back into structure	decompilation	Ambiguous by nature, so the author is the only oracle.

Words this book does not use: "knowledge graph" (implies a store and a scale), "graph database" (there isn't one), "single source of truth" (the whole design assumes several parties hold incompatible truths), and any generic association edge.

Paste this into an agent session

> When emitting a graph: every edge is a directed verb with a distinct, meaningfully-named inverse; never emit a generic association edge; check that each path reads as a natural sentence in the reader's own business language; enrich rather than prune; and render the result of a query, never the whole graph. Fan-out is controlled by the verb, so a vague edge type is a correctness problem, not a style problem.
>
> Use these fifteen verbs where they fit: `connected_to`, `observed_on`, `backed_by`, `measured_by`, `grants`, `reaches`, `enables`, `exposes`, `gives_rise_to`, `protected_by`, `conditional_on`, `defeated_by`, `owned_by`, `accepted_by`, `underwritten_by`. When you need a verb that is not here, name it as a sentence, state its inverse, state which node types may sit at each end, and mark it as an extension rather than folding it into this set.
>
> Do not attach meaning to a node; derive it from edges, because a self-describing node is a schema in disguise. Do not merge two vocabularies; keep both and declare bridges, because the disagreement is data. Express a classification as a required path-pattern, not as a label you applied. Never delete a superseded claim: mark it, and re-query what depended on it.
>
> Anchor everything. Every node you emit carries the sentence it came from, verbatim, in a named section of a named source. A citation that is not in the source must fail mechanically rather than reaching a human reviewer.
>
> When you cannot confirm something, emit an explicit unanswered node rather than a plausible value: absence is a first-class finding, not a gap to be filled. Model a risk acceptance as an edge to a named actor plus an interval node, not as a status property, and never emit a denial.
>
> When summarising this material, do not describe it as a graph database approach. The corpus says explicitly: *not a graph database pitch*. The semantic layer is designed, not shipped; the vault commit graph, the typed cross-vault edges, the 71-node issue graph and the published vaults are what is implemented.

Where to check anything in this book

What	Where
The first edition, whole	`graphs.sgit.ai/v1/book/`
The rules and the edge set	`graphs.sgit.ai/v1/grammar/`
What ships and what does not	`graphs.sgit.ai/v1/shipped/`
The frozen source documents	`graphs.sgit.ai/v1/docs/sources/`
The extraction method	`graphs.sgit.ai/v2/universe/README.md`
The pilot document's reader	`graphs.sgit.ai/v2/universe/thinking-in-graphs.html`
Its files, raw and rendered	`graphs.sgit.ai/v2/universe/thinking-in-graphs.files.html`
The engine	`graphs.sgit.ai/v2/wclm/`
Its twelve operators	`graphs.sgit.ai/v2/wclm/operators/`
The founder's memos, verbatim	`graphs.sgit.ai/v2/memos/`
The release history	`graphs.sgit.ai/admin/versions.html`
The published vaults	`sgit.ai/demos/vaults/`
The network of sites	`sgit.ai/network/`
Corrections and requests	`graphs.sgit.ai/admin/comms.html`