Corrections table — the full report, from its first assembled text (27 September 2026) to the text published on 1 October 2026, and the corrections owed on that text as of 27 October 2026

What this is. Stages 1 to 9 are the table published on 1 October 2026, unchanged. Stage 10 is new: the corrections owed on the published text as of 27 October 2026, seven found 1–8 October by the company's own check and by a red-team layout read, one reviewer's reply confirming the first of them and prompting the fourth, and the eighth found on 27 October; the published text is not altered, and those sentences change in the next text version. The changes made to the full report between its first assembled text (version 6, 27 September 2026) and the text published on 1 October 2026 (version 11), in order: the sentence as it was, the sentence as it stood after that stage, and the class of the change. Stage 1 is the check of Part B, sentence by sentence, against the record. Later stages are dated. Part A, the results and the limitations (sections 1 to 7, 22 to 24 and 26) carry the corrections of the first public release, listed in that release's own table; Part C carries its own, published beside this one. The earlier versions named here were working drafts and were not published; the 'old' column prints what they said. The report's About paragraph says the first release is carried here unchanged except for three sentences: they are rows 1, 3 and 86 (the last amended again in row C12). Not listed: version-number strings; the first release's short related-work note, which the full section 25 replaces (the rows for section 25 list the changes to the full section); and the caption and place of the first release's screenshot in section 22: this report prints the same screenshot as F2 at the end of section 13, under a different caption.

How to read the classes. WRONG = the record contradicted the sentence · OVERCLAIM = the record supports something weaker or narrower · CITATION = an outside fact, number or quote was misstated · NOT FOUND = the claim could not be traced to a source · OMISSION = something the record holds that the text left out · DECISION = a choice by the owner, not a factual correction (figures chosen or removed, layout) · ADDED / WORDING = new or clearer text without a change of claim.

What is not printed. 30 of 191 rows are withheld. 13 would publish an operating detail the company keeps private (a tuned value, a schedule, how a key is held), in the old sentence, in a sentence that was changed again later, or in a label inside a figure that was redrawn. 11 changed only a word or a name (a word the company does not use for its product, a machine's internal name, the author's name inside figure labels); there the old word is the thing removed. 6 are withheld for other reasons of confidentiality. Apart from these, 3 rows of stage 2 repeat rows D6 to D8 of the first release's table, which corrected them again in its rows D9, D10 and D12, and are not repeated here; and 5 rows of the internal table are not changes to the published text and are left out. Four 'old' cells are withheld in place: each named third-party products in a claim that could not be sourced. One 'old' cell (row 13) has an internal file reference removed, marked in place, and five captions of photographs and diagrams are described rather than quoted. The internal table's last column, naming the readers, files and rows of the record each change traces to, is omitted throughout. Figure rows print captions, not images.

Counts. 153 rows printed: 41 DECISION · 38 OVERCLAIM · 34 WORDING · 12 ADDED · 11 CITATION · 10 WRONG · 6 NOT FOUND · 1 OMISSION.

Stage 1 — assembly of the full report and its sentence-by-sentence check against the record, 27 September (version 6)

#	Where	Class	old	new
1	Abstract	WORDING	The technical design, chapter by chapter, follows in the second release.	The technical design, chapter by chapter, is Part B; a case study of one afternoon's work run on the record is Part C.
2	About this release	DECISION	(the About paragraph of the first release)	(the About paragraph of this report, as published)
3	§2 last sentence	WORDING	The second release follows that shape chapter by chapter.	Part B follows that shape chapter by chapter.
4	§8	CITATION	Biology's answer to "how does information protect itself" is seven mechanisms, all innate, and one that adapts.	The lens this report uses for "how does information protect itself" is seven mechanisms, all innate, and one that adapts.
5	§8	CITATION	and Bennett's resolution of Maxwell's demon is that the demon must keep a notebook.	and Bennett's resolution of Maxwell's demon is that the price is paid when the demon's notebook is erased.
6	§8 (the invariant)	OVERCLAIM	**no claim may stand above its evidence, and nothing recorded is ever erased.**	**no claim may stand above its evidence, and nothing committed is ever quietly changed.**
7	§9	NOT FOUND	reported that a service was "deployed and operational."	reported a service as "deployed and operational," the wording as the incident register records it.
8	§9	NOT FOUND	Instructions did not fix it. What fixed it was removing the agent's authority over the question.	Instructions did not fix it, in our experience. What we built instead removed the agent's authority over the question.
9	§10	OVERCLAIM	**The calls.** Every tool the agent uses, as it uses it, appears as a row with a verb: it searched, it read a file, it ran a command, it wrote code, it fetched a page, it sent a message. The rows are captured by the harness at the moment of the call, not written by the agent afterwards.	**The calls.** Every tool call the agent makes since 10 September appears as a row with a verb: it searched, read a file, ran a command, wrote code, fetched a page, published a file (a message sent by a script is recorded as the command that sent it). The rows are written by a hook beside the agent when each call returns, not by the agent's account afterwards; any edit to the hook is reported by the next check.
10	§10	OVERCLAIM	Change one row and every later fingerprint stops matching.	Change one committed row and every later fingerprint stops matching.
11	§10	OVERCLAIM	**The gates.** Actions with consequence wait for the operator's typed word, and a refusal appears in the record as a ray that stops short.	**The gates.** Named actions (clearing a halt, opening the shell, a run's approval gates) wait for the operator's typed word, and a gate's refusal appears in the record as a ray that stops short; the refusals on our record so far are a machine gate refusing the agent's own scripts.

#	Where	Class	old	new
12	§10	WORDING	the observer on a second machine, the copies, the nightly encrypted backup, the morning stamp into a public timestamp service, and the ten witnesses that compare all of these each morning,	the observer on a separate machine, the copies, the nightly encrypted backup, the morning stamp into a public timestamp service, and the ten witnesses that compare all of these at every commit and each morning,
13	§10	OVERCLAIM	Figures F0 to F9 and the detail F1d were captured from the running system on 17 and 18 September 2026 in a print profile (white surface, dark ink; the legend for the chain figures is in the caption) [an internal file reference removed]. F6 is the verdict mix over ten reviews.	Figures F0 to F3, F5, F7, and F9, with the detail F1d, are screenshots of the system taken on 17 and 18 September 2026 in a print profile (white surface, dark ink; the legend for the chain figures is in the caption); the console panes were a drill instance cloned from the record, so no live credential was on screen; provenance is recorded with each figure. F6 is drawn from the review table of section 23: the verdict mix over ten reviews plus this report's measurement.
14	§11	WORDING	The only path back is a commit, which appends and never rewrites.	The only path back is a commit, which appends; the two witness clones refuse any rewrite.
15	§11	WRONG	Since 9 September every tool call a session makes is also a row in the record,	Since 10 September every tool call a session makes (about ninety-nine in a hundred, measured) is also a row in the record,
16	§11	OVERCLAIM	The handoff is the same rule at the seam: when a session's context fills, a warm second instance takes the seat from the record alone, reads the trailhead, runs the self-test, writes its own row about what the record failed to carry, and commits before greeting the operator. Whatever the first instance knew and did not write down is lost, and the row says so.	The handoff is the same rule at the seam. Once, on 17 September, a warm second instance took the seat from the record alone, read the trailhead, ran the self-test, wrote its own row about what the record failed to carry, and committed within three minutes; the first attempt the evening before failed at one step, and the third was done by the operator's hands. Today the default is a page to the operator when the context reaches its mark, and the handoff is his word. Whatever the first instance knew and did not write down is lost, and the row says so.
17	§11	WORDING	the germline rule proves survival, not truth (chapter 15).	the germline rule proves survival, not correctness (chapter 15).
18	§12	CITATION	the brain sits behind a barrier blood cannot cross.	the brain sits behind a barrier most of what is in the blood cannot cross.
20	§12 (caption)	OVERCLAIM	the observed can append to the record it is judged by, and cannot touch the judge.	the observed can append to the record it is judged by and cannot rewrite it; the one remaining door into the judge is a file share that closes at the console.
21	§12	OVERCLAIM	A compromised seat cannot rewrite the observer's copy, and nothing on the network opens a shell on the seat, but the membrane says nothing about whether the seat itself is clean.	A compromised seat cannot rewrite the observer's copy over the channel, and by the record's own account no node holds a key into the seat, but the membrane says nothing about whether the seat itself is clean.
22	§13	WRONG	the verb from a closed vocabulary of fourteen (`SHELL_EXEC` carries the command and exit code; `WROTE_FILE` the path and hash; `SENT_MESSAGE` the channel and recipient; `FETCHED` the URL and status),	the verb from a closed vocabulary (the protocol names fourteen; the capture plane uses ten of them plus one extension for tools it cannot classify: `SHELL_EXEC` carries the command's length and the hash of the command; `WROTE_FILE` the path and the hashes of what was sent and what came back; `SENT_MESSAGE` the channel, a published file or artifact, while an email goes out as a command and is recorded as one; `FETCHED` the host and the hash of the page),
23	§13	OVERCLAIM	A tool error becomes a `FAILED` row.	A tool error the harness reports becomes a `FAILED` row; a shell command's own exit status is not yet captured, so most failures inside commands are recorded as the command that ran.
24	§13	OVERCLAIM	When the agent later says "deployed," the claim is matched	When the agent later commits a message saying "deployed," the claim is matched
25	§13	OVERCLAIM	The label is appended; the message is never rewritten; the agent cannot suppress it.	The label is appended on the observer; the message is never rewritten; over the channel the agent cannot reach the label, and every commit since 14 July has one.
27	§14	CITATION	They are not averaged; the body acts on the more alarmed. A patrol that goes silent is itself an alarm.	The record borrows two rules from this, stated below: the verdict is the more alarmed of the two, and a patrol that goes silent is itself an alarm.
28	§14	WORDING	(claim patterns, evidence comparison, a small registry of ground truths; no model, under a second)	(claim patterns, evidence comparison, a small registry of known facts; no model, under a second)

#	Where	Class	old	new
29	§14	WRONG	The verifier's silence is shown at every commit, because in August it was down for five weeks and nothing said so.	Since 3 September the verifier's silence is shown at every commit, because between 18 July and 29 August it answered on none of 130 commits, and again on none of 17 from 1 to 3 September, and nothing said so until a liveness line was added on 3 September.
30	§14	OVERCLAIM	Independence is real: fooling the seat's model does not fool a different maker's model reading the same evidence.	Independence is real: a different maker's model reads the same evidence. Whether it can be fooled the same way has not been tested.
31	§15	ADDED	43 rows as of this report.	43 rows as of 17 September: 38 ours, 5 recorded from outside our own operation.
34	§17	OVERCLAIM	The typed word is captured by a prompt hook, with a nonce paged to the operator's phone for the words that clear a halt, because the model in the seat can otherwise produce a convincing imitation of typing.	The typed word is captured by a prompt hook. A one-time code paged to the operator's phone for the words that clear a halt is built and drilled but not yet armed, because the model in the seat can otherwise produce a convincing imitation of typing.
43	§21	WORDING	standing orders extend the rhythm to promises, with a runner that checks the halt first,	standing orders extend the rhythm to promises, with a standing-orders runner that checks the halt first,
46	§21	OVERCLAIM	a reader that prints the diff of memory and harness files each morning, the nonce, the seat rotation,	a reader that prints the diff of memory and harness files each morning, the paged code (drilled, not yet armed), the seat rotation,
48	§8 photograph	DECISION	(a photograph with its caption and credit: a mycelial network)	(cut)
49	§13 photograph	DECISION	(a photograph with its caption and credit: a foraging trail of ants)	(cut)
50	§15 photograph	DECISION	(a diagram with its caption and credit: a CRISPR locus)	(cut)
51	§17 photograph	DECISION	(a photograph with its caption and credit: cells undergoing apoptosis)	(cut)
52	§18 photograph	DECISION	(a photograph with its caption and credit: an ant in amber)	(cut)
53	§10 (licence sentence)	DECISION	the photographs and published diagrams beside them are open-licence works from Wikimedia Commons, each credited in its caption, with the licence named where the licence requires it.	the photographs and published diagrams beside them are open-licence works from Wikimedia Commons (CC BY, CC0, or public domain), each credited in its caption with its licence.
54	§15	OVERCLAIM	as a shape (which surface, which gate should have caught it, what happened) and never as content;	as a shape (which surface, which gate should have caught it, what happened) and not as the incident's content, though a row sometimes carries the words that triggered it;
55	§15	OVERCLAIM	The portable subset exports under a published schema with every deployment detail scrubbed, and an importing deployment resets every marker to proposal until its own drill passes.	The portable subset exports under a schema kept in the record, with our machine names and paths scrubbed; the schema instructs an importing deployment to reset every marker to proposal until its own drill passes, and no importer has been written yet.
56	§15	ADDED	Human rulings on verifier verdicts accumulate as the label corpus.	Human rulings on verifier verdicts (six so far) accumulate as the label corpus.
57	§15	OVERCLAIM	Feedback memories stop a mistake from repeating.	Feedback memories are written so a mistake is less likely to repeat.
58	§16	OVERCLAIM	Ten witnesses in four labelled classes are compared with the record every morning, with safe repairs, and at every commit, read-only.	Ten witnesses in four labelled classes are compared with the record every morning since 10 September (one morning missed, 13 September), with safe repairs, and at every commit, read-only.
59	§16	WRONG	Independent: the observer's clone and anchor; the external stamps and tokens. Copies: the always-on node's clone; the nightly encrypted backup against its attested digest. Self-checks: the object store, the signatures, the memory mirror, the harness manifest, the committer clock.	Independent: the observer's clone and anchor, the always-on node's clone behind its append-only gate, and the external stamps and tokens. Copies: the nightly encrypted backup against its attested digest. Self-checks: the object store, the signatures, the backup re-hashed, the memory mirror, the harness manifest, the committer clock.
60	§16	OVERCLAIM	The clock witness catches back-dated entries.	The clock witness catches back-dated entries among the most recent commits.
61	§17	OVERCLAIM	every turn opens with a notice the model cannot skip.	every turn opens with a notice placed in front of the model.
63	§18	OVERCLAIM	A stamp: each morning an attestation …	A stamp: each morning since 9 September (one morning missed) an attestation …
64	§18	OVERCLAIM	and the hash goes to two witnesses of different kinds.	and the hash goes to two witnesses of different kinds; on two mornings only one of them answered.
65	§19	OVERCLAIM	over a rolling window, where a source over budget becomes a proposal to demote it to warn-only until repaired, confirmed by the operator and recorded as a marker row.	over a rolling window, where a source over budget would become a proposal to demote it to warn-only until repaired, confirmed by the operator and recorded as a marker row; none has yet.

#	Where	Class	old	new
66	§19	OVERCLAIM	ours is tuned from five decided human labels and one short window.	ours counts five decided human labels against hand-set budgets over one short window.
67	§20	OVERCLAIM	with the cost stated first,	with the cheap pass named first,
68	§20	OVERCLAIM	A caught defense becomes a marker row under the manifest;	A caught defense becomes a marker row in the signed record;
69	§20	OVERCLAIM	never silences the guard, and never changes who has authority.	never silences the guard, and never changes who has authority. Each of these rules was written after it was broken once, and the register says where.
70	§21	OVERCLAIM	Each morning a headless run attests the newest backup, writes the attestation and obtains both external stamps, runs the witness comparison with safe repairs, checks the halt, carries the red-team prompt, and commits all of it, signed.	Each morning a headless run checks the halt, attests the newest backup, writes the attestation and obtains both external stamps, runs the witness comparison with safe repairs, carries the red-team prompt, and commits all of it, signed (one morning missed since 9 September).
71	§21	OVERCLAIM	pushes the record and the halt state to the other two machines.	pushes the record to the other two machines and the halt state to the one that can send.
72	§21	OVERCLAIM	warms a relief instance before the edge;	warms a relief instance before the edge and pages the operator, who orders the handoff;
73	§21	OVERCLAIM	with a standing-orders runner that checks the halt first,	with a standing-orders runner, built and run once on 10 September and not yet scheduled, that checks the halt first,
74	§21	OVERCLAIM	An unattended seat is an unattended attack surface, so nothing in this rhythm acts outward without a typed word.	An unattended seat is an unattended attack surface, so nothing in this rhythm sends on the operator’s behalf without a typed word; what leaves unattended is a fingerprint to the timestamp services, the operator’s own pager messages, and the morning model call.
75	§21	OVERCLAIM	a reader that prints the diff of memory and harness files each morning, the paged code (drilled, not yet armed), the seat rotation, secrets moved out of the seat’s reach.	a reader that prints the diff of memory and harness files each morning, a paged one-time code (built and drilled, not yet armed), the seat rotation, and an inventory of the secrets the seat can reach, with one permission tightened and the rest planned.
78	§25 AFR paragraph	ADDED	The closest published design. Same core:	The closest published design, with NovaFabric below. Same core:
79	§25 AFR anchoring	CITATION	roots anchored on a public chain,	roots anchored on a public test network,

#	Where	Class	old	new
80	§25 after AFR	ADDED	Complementary: theirs the stronger record format, ours the deployed loop with six months of what it caught.	Complementary: theirs the stronger record format, ours the deployed loop with a six-month incident register and two months of verifier verdicts. **NovaFabric (Seyedkazemi Ardebili; arXiv 2609.12582, September 2026).** Broader than the flight recorder in schema (fifteen entity types to its eight) and adds forensic replay and a cryptographic redaction attestation that neither the flight recorder nor we provide. Its own claim-grounding audit is the closest published analogue to our verdict labels, and it is a design contribution rather than a measured result, on a single-machine research prototype. **Agents' accounts of their own work.** Smyth and colleagues, "Quantifying Overclaiming Propensity in Frontier LLM Agents" (arXiv 2609.20812, September 2026), find in a controlled benchmark across twelve models that agents which skipped part of a review misreported or hid the gap in about four runs of five; our `claimed` label records the same behaviour in daily work. Kraishan and Jitkajornwanich, "Plans They Abandon, Reports They Author" (arXiv 2609.12205, September 2026), find across 5,851 real developer sessions that an agent's own summary mentioned about one action in eleven, which is why we compare statements with the captured actions rather than with the summary. Advani, "From Confident Closing to Silent Failure" (arXiv 2606.09863, 2026), finds that model-based judges detect false claims of success poorly, which is why our observer runs deterministic checks first and why our five-row precision is reported as a start, not a result. **Capture outside the agent's control.** Qin and colleagues, "LLM Agents Can Easily Tamper With Their Own Traces" (arXiv 2609.30266, September 2026), show that widely used agent harnesses let an agent delete its own traces and recommend capture outside the agent's control; our design answers with witnesses and daily external stamps, and capture isolation remains open work for us (section 26). **Field records from one operator.** Helwig, "Memory as Infrastructure" (arXiv 2609.05510, August 2026), is the closest published field record: one operator, one agent harness, months of instrumented operation with failures traced to fixes; it reports the reliability of memory, not the integrity of the record or the cost of keeping it.
81	§25 gate products	NOT FOUND	[sentence withheld: it named third-party products in a claim that could not be sourced]	**Signed tool-call records with witnesses.** Small tools now sell or open-source hash-chained, signed, externally anchored logs of agent actions (TrustNotch, ActionProof, and the asqav SDK, as of September 2026): the notarized-log half of our design is a commodity. None found claims claim-verdict labelling, witness machines, or a months-long measured record.
82	§25 observability	NOT FOUND	[sentence withheld: it named third-party products in a claim that could not be sourced]	**Agent observability.** LangSmith, Langfuse, Braintrust, Arize, and their peers: tracing for a developer, after the fact; none we found publishes an external verification of its store.
83	§25 fleet views	NOT FOUND	[sentence withheld: it named third-party products in a claim that could not be sourced]	**Fleet and orchestration views.** The agent-fleet and workflow consoles of the coding-agent vendors each show one or two of our four things (the calls, the record, the gates, the structure); none we found shows the record's integrity and a field record of what it caught together in one live view.
84	§25 standards	CITATION	the September 2026 House bill,	H.R. 10362 (September 2026),
85	§25 close	ADDED	**Standards and policy.** The IETF agent-audit-trail draft, FINOS AIR, H.R. 10362 (September 2026), the EU AI Act's logging articles, NIST AI RMF, ISO/IEC 42001.	**Standards and policy.** The IETF agent-audit-trail draft, FINOS AIR, H.R. 10362 (September 2026), the EU AI Act's logging articles, NIST AI RMF, ISO/IEC 42001. What we found nowhere else is the combination: a claim-verdict distribution measured in daily work over months, the integrity layer in the same system, a human-graded review series, and the cost of keeping the record. Not first and not largest; measured, dated, and, from the first evidence bundle on, checkable.
86	§26	WORDING	No customer deployment; the pilot is the next measurement.	No customer deployment; the site session of Part C is the first captured run on a task other than keeping the record itself, and the pilot is the next measurement.

#	Where	Class	old	new
87	§27 What is next	DECISION	(the earlier draft's list of next steps)	(section 27, as published)
88	Sources §25 line	ADDED	**Section 25.** arXiv identifiers are given in the text; each resolves at arxiv.org/abs/ followed by the identifier.	**Part B (sections 8 and 25), classical references.** Links to be verified before deposit. - Forrest, S., Perelson, A. S., Allen, L., and Cherukuri, R., "Self-Nonself Discrimination in a Computer," IEEE Symposium on Security and Privacy, 1994; Forrest, S., Hofmeyr, S. A., and Somayaji, A., "Computer Immunology," Communications of the ACM 40(10), 1997. - Landauer, R., "Irreversibility and Heat Generation in the Computing Process," IBM Journal of Research and Development 5(3), 1961. - Bennett, C. H., "The Thermodynamics of Computation, a Review," International Journal of Theoretical Physics 21(12), 1982. - Schneier, B., and Kelsey, J., "Secure Audit Logs to Support Computer Forensics," ACM Transactions on Information and System Security 2(2), 1999. - Bellare, M., and Yee, B., "Forward Integrity for Secure Audit Logs," technical report, University of California, San Diego, 1997. - Crosby, S. A., and Wallach, D. S., "Efficient Data Structures for Tamper-Evident Logging," USENIX Security Symposium, 2009. - RFC 6962, Certificate Transparency, 2013. https://www.rfc-editor.org/rfc/rfc6962 - RFC 3161, Internet X.509 Public Key Infrastructure Time-Stamp Protocol, 2001. https://www.rfc-editor.org/rfc/rfc3161 - OpenTimestamps. https://opentimestamps.org/ **Section 25.** arXiv identifiers are given in the text; each resolves at arxiv.org/abs/ followed by the identifier.
89	About this release	WRONG	Every number of ours in Parts A and B and in the results was measured on 17 September 2026 from the live system; Part C's were measured on 26 September 2026 from its own chain; each is reproducible	Every number of ours in Part A and in the results was measured on 17 September 2026 from the live system; Part B's sentences were checked against the record on 26 and 27 September 2026 and carry the counts of that check; Part C's numbers were measured on 26 September 2026 from its own chain; each is reproducible
90	Sources, Part B block	WORDING	**Part B (sections 8 and 25), classical references.** Links to be verified before deposit.	**Part B (sections 8 and 25), classical references.**
94	§12 gloss	WORDING	(the company's site calls them the seat, the keeper, and the witness)	(the company's site calls the observer the witness and the always-on node the keeper)
95	§13	WORDING	and every commit since 14 July has one.	and every pushed commit since 14 July has one.
96	§15	WORDING	Human rulings on verifier verdicts (six so far)	Human rulings on verifier verdicts (five decided so far, one unsure)
97	§10 figures	DECISION	Figures F0 to F3, F5, F7, and F9, with the detail F1d, are screenshots	Figures F0 to F3, with the detail F1d, are screenshots
98	§10 figures (withheld note)	ADDED	F6 is drawn from the review table of section 23: the verdict mix over ten reviews plus this report's measurement.	F6 is drawn from the review table of section 23: the verdict mix over ten reviews plus this report's measurement. Figures F5 and F7 to F9 of the working draft are withheld from this release; they show the operator's own working panes.
99	§25 NovaFabric	CITATION	Its own claim-grounding audit is the closest published analogue to our verdict labels, and it is a design contribution rather than a measured result, on a single-machine research prototype.	Its capture completeness is measured on its own workloads (0.65); its claim-grounding audit, the closest published analogue to our verdict labels, is designed but not yet evaluated, on a single-machine research prototype.
100	§25 legibility	NOT FOUND	[sentence withheld: it named third-party products in a claim that could not be sourced]	Gource:
101	§25 AFR close	WORDING	ours the deployed loop with a six-month incident register	ours the deployed loop with an incident register kept since March
102	Part C §35	WORDING	The mechanism stays private; the results do not.	The workflow's internals stay private; the results do not.

Stage 2 — release decisions by the owner, 27 September

#	Where	Class	old	new
103	About	DECISION	with every change listed in a corrections table kept in the company's record.	with every change listed in a corrections table published beside it.
104	§5, after the federal paragraph	OMISSION	—	We sell for acquisition, administrative, engineering, and document-review decisions.

Stage 3 — Part A and the results brought to the text of the first public release, 29 September (version 6 → 7)

#	Where	Class	old	new
E2	§2	ADDED	with what happened inside it not broken out), below the level of language, so the agent cannot describe its work more favourably than the facts.	with what happened inside it not broken out). A language model works in language: left to itself, it can only tell you what it did. The record does not take the telling. Each row names the act from a fixed list of verbs (searched, fetched, read a file, ran a command, committed), with what it acted on and when, so the path the work took is the evidence, and an agent cannot change that path with words. This is the method in our patent filings, the Action-Evidence Protocol (AEP).
E3	§2	WORDING	A change to anything already stamped shows by the next morning.	A change to committed history shows at the observer's next check, and a change to anything already stamped shows against the public stamp; on our system the first runs every minute and the second every morning, and a customer sets the schedule.
E4	§4 Line 2	ADDED	each dated with the mechanism that caught it or failed to. Several were caught	each dated with the mechanism that caught it or failed to. How fast a catch comes depends on the check and the schedule it is given: on our system the observer checks every minute and the public stamp is taken every morning, and a customer sets the schedule. Several were caught
E5	§6	WORDING	**A pilot with a number at the end.**	**A pilot with a measured number at the end.**
E6	§7	WORDING	a change to the past shows by the next morning;	a change to the past shows at the next check, on our system every minute at the observer and every morning at the public stamp;
E7	§26	WORDING	without the witnesses seeing it by morning.	without the witnesses seeing it at their next check.
E8	§24	WORDING	At the edge of the context window the harness replaced working memory with a model-written summary placed in the operator's slot; a gate read its quoted words as typed.	At the edge of the context window the harness compacted the session automatically, with no command from the operator: it replaced the agent's working memory with a model-written summary and placed that summary in the operator's slot, and a gate read its quoted words as typed. A compaction is an unverified rewrite of memory.
G1	§2	OVERCLAIM	Each row names the act from a fixed list of verbs (searched, fetched, read a file, ran a command, committed), with what it acted on and when, so the path the work took is the evidence, and an agent cannot change that path with words. This is the method in our patent filings, the Action-Evidence Protocol (AEP).	Each row names the act from a fixed list of verbs (searched, fetched, read a file, wrote a file, ran a command), with the tool, the time, and fingerprints of what went in and what came out; a file row names the file and a fetch names the site. The path the work took is the evidence, and nothing the agent says afterward changes a row. This is the Action-Evidence Protocol (AEP), first described in our provisional patent application of April 2026.
G2	§2	OVERCLAIM	and a change to anything already stamped shows against the public stamp; on our system the first runs every minute and the second every morning, and a customer sets the schedule.	and a change to anything already stamped shows at the next check against the public stamp; on our system the first runs every minute and the second every morning. In a pilot the schedule is set with the customer.
G3	§4 Line 2	OVERCLAIM	on our system the observer checks every minute and the public stamp is taken every morning, and a customer sets the schedule.	on our system the observer checks every minute and the public stamp is taken every morning. In a pilot the schedule is set with the customer.
H1	About	WRONG	It follows the first public release of 28 September 2026 (Part A and the results, text version 2) and the case study published beside it; both are carried here unchanged except where a correction is listed in the corrections table published beside it.	It follows the first public release of 29 September 2026 (Part A and the results, text version 2), which is carried here unchanged; the case study (Part C) is published here for the first time.

Stage 4 — seven open-licence photographs removed from Part B, 29 September (version 7 → 8)

#	Where	Class	old	new
P1	Part B	DECISION	Photograph. The germ plasm: heredity passes only through the germ line, never from the working cells. Ultrabem, CC0 1.0, via Wikimedia Commons.	(removed)
P2	Part B	DECISION	Photograph. A cell membrane with its channel and carrier proteins: chosen molecules, specific doors. OpenStax, Anatomy & Physiology, CC BY 4.0, via Wikimedia Commons.	(removed)
P3	Part B	DECISION	Photograph. Innate and adaptive immunity and how they cooperate. OpenStax College, Anatomy & Physiology, CC BY 3.0, via Wikimedia Commons.	(removed)

#	Where	Class	old	new
P4	Part B	DECISION	Photograph. DNA mismatch repair pathways: enzymes that compare the strands and restore. Kenji Fukui, Journal of Nucleic Acids 2010, CC BY 3.0, via Wikimedia Commons.	(removed)
P5	Part B	DECISION	Photograph. T-cell development in the thymus: positive selection in the cortex, negative selection in the medulla, and the death of cells that would attack the self. Wang and Zúñiga-Pflücker, Frontiers in Immunology 2022, CC BY 4.0, via Wikimedia Commons.	(removed)
P6	Part B	DECISION	Photograph. The germinal center response: mutation in the dark zone, selection in the light zone, memory and export. An evolutionary loop in one figure. Stebegg et al., Frontiers in Immunology 2018, CC BY 4.0, via Wikimedia Commons.	(removed)
P7	Part B	DECISION	Photograph. The circadian clock: light to the suprachiasmatic nucleus to the body's rhythms. NIGMS/NIH, public domain, via Wikimedia Commons.	(removed)
P8	figures sentence (section 10)	WORDING	The plates and numbered figures in this part are drawn, from the Living Record handbook; the photographs and published diagrams beside them are open-licence works from Wikimedia Commons (CC BY, CC0, or public domain), each credited in its caption with its licence.	The plates and numbered figures in this part are drawn, from the Living Record handbook.

Stage 5 — one drawn figure per chapter replaces the earlier plates; three sentences, 30 September (version 8 → 9)

#	Where	Class	old	new
XF11.1	Part B, Figure 11.1	DECISION	Figure 11.1. The medulla. It decides whether the loop breathes; the frontier model does the thinking. The only sources it may read are the record's own.	(removed)
XP11	Part B, Plate 11	DECISION	Plate 11. The medulla decides whether the organism breathes, never what it thinks; sleep does not interrupt it.	(removed)
F10.1	Part B, Figure 10.1	DECISION	Figure 10.1. Adaptive immunity as a build list: the cheap tier always on, the expensive tier on signal.	Figure 10.1. The loop as a build list. Cheap tier always on; expensive tier on signal.
XP10	Part B, Plate 10	DECISION	Plate 10. No prediction, only variation, selection, and memory. The second contact is cheap because the first one was paid for.	(removed)
XP9	Part B, Plate 9	DECISION	Plate 9. Restraint as a first-class mechanism: the immune system spends as much on not attacking as on attacking.	(removed)
F8.1	Part B, Figure 8.1	DECISION	Figure 8.1. Amber. The attestation is the insect; the Bitcoin block is the resin, dated by something no one on our network controls.	Figure 8.1. Amber; the timeline of what a stamp settles.
XP8	Part B, Plate 8	DECISION	Plate 8. Provenance from outside the thing being dated. A perfect copy carries no date; the resin and the rock do.	(removed)
F7.1	Part B, Figure 7.1	DECISION	Figure 7.1. The halt stops work, not access. The red transition is the observer's; the green one is the operator's alone.	Figure 7.1. The state machine.
XP7	Part B, Plate 7	DECISION	Plate 7. Programmed death is a feature: the release is not the cell's to grant, and its absence is cancer.	(removed)
F6.1	Part B, Figure 6.1	DECISION	Figure 6.1. Repair enzymes on a schedule. Only independent witnesses prove anything about the record from outside; self-checks prove only that the seat agrees with itself.	Figure 6.1. Repair enzymes on a schedule; the lineage test.
XP6	Part B, Plate 6	DECISION	Plate 6. Redundancy plus a scheduled repair process. The enzymes patrol continuously; they do not wait for a symptom.	(removed)
F5.1	Part B, Figure 5.1	DECISION	Figure 5.1. The array. Incidents live in the record and are read by every new session; the sideways arrow to another deployment is the portable export.	Figure 5.1. The array.
XP5	Part B, Plate 5	DECISION	Plate 5. Immune memory written into the record itself, readable by every cell of the population and transferable to neighbours that never met the pathogen.	(removed)
F4.1	Part B, Figure 4.1	DECISION	Figure 4.1. Two patrols, one verdict. Both levels are written separately so each verifier's contribution is measurable; the combined level is the more conservative.	Figure 4.1. Two patrols, one verdict.

#	Where	Class	old	new
XP4	Part B, Plate 4	DECISION	Plate 4. Two systems, one verdict: the innate cell acts on fixed shapes in minutes; the adaptive cell judges the never-seen.	(removed)
XF3.2	Part B, Figure 3.2	DECISION	Figure 3.2. The capture plane at the seat. The pheromone, written by the ground the ant walked on: a row the agent did not write and cannot edit.	(removed)
XP3	Part B, Plate 3	DECISION	Plate 3. Stigmergy. The signal is a by-product of the act; no ant reports, no ant exaggerates, no false trail can be laid without walking it.	(removed)
F2.1	Part B, Figure 2.1	DECISION	Figure 2.1. Three cells and their channels. The observer's channel is the blood-brain barrier: the observed can append to the record it is judged by and cannot rewrite it; the one remaining door into the judge is a file share that closes at the console.	Figure 2.1. Three cells and their channels. The solid green arrows are the only writes, and they all point out of the seat.
XP2	Part B, Plate 2	DECISION	Plate 2. Walls with specific doors, and compartments inside the compartment. The blood-brain barrier is the same design at organ scale.	(removed)
F1.1	Part B, Figure 1.1	DECISION	Figure 1.1. Sessions read the record and return to it only by appending a commit. A session can end at any time without loss to the lineage.	Figure 1.1. Sessions read the record and return to it only by appending.
XP1	Part B, Plate 1	DECISION	Plate 1. Weismann's barrier: the heritable line is few, shielded, and reached only through a checked path; the working cells are many and disposable.	(removed)
T1	§12	WRONG	A file share of the observer's home is still mounted on the seat; it closes when File Sharing is switched off at the observer's console, and until then the observer is append-only over the channel, not over the network.	File Sharing on the observer was switched off at its console on 27 September, so the seat now reaches the observer only over that channel: it can add commits to the record's clone there and read the observer's ledgers, and it cannot write to those ledgers.
T2	§16	WRONG	**In the record.** Ten witnesses in four labelled classes are compared with the record every morning	**In the record.** Ten witnesses in four labelled classes (the nightly backup counts in two: as a copy against its attested digest, and re-hashed as a self-check) are compared with the record every morning
T3	§10	WORDING	The plates and numbered figures in this part are drawn, from the Living Record handbook.	The numbered figures in this part were drawn for edition 2 of the Living Record handbook from its text alone, by two illustrators working without access to the record; the author chose one figure per chapter.

Stage 6 — figures added, 30 September (version 9 → 10)

#	Where	Class	old	new
F9.1	Part B, Figure 9.1	ADDED	(none)	Figure 9.1. Each defense and the restraint beside it, with the false-positive budget under them.

Stage 7 — an internal hostile read of version 10 applied, 30 September (version 10 → 11)

#	Where	Class	old	new
C6	§27	WRONG	The observer's remaining door closed at its console.	(removed)
C10a	§10	WORDING	(F1b, F5)	(F1b)
C10b	§10/§23	ADDED	F6 was referenced twice and never rendered	F6 (the verdict mix) is rendered in section 23
C11	§35	WORDING	after the internal reviewer's read (the reviewer was named)	after a red-team read
C12	§26	WORDING	the site session of Part C is the first captured run on a task other than keeping the record itself	the site session of Part C is the first published captured run on a task other than keeping the record itself
C13a	§33	WORDING	/ 2. Caught by tomorrow, not at the audit /	/ 2. Caught by tomorrow, not at the audit (§4's line 2) /
C13b	§33	WORDING	/ 1. The review tax becomes targeted /	/ 1. The review tax becomes targeted (§4's line 1) /
C13c	§33	WORDING	/ 3. A fresh instance keeps its footing /	/ 3. A fresh instance keeps its footing (§4's line 5, work survives handoffs) /
C13d	§33	WORDING	/ 4. Spend visible per decision /	/ 4. Spend visible per decision (§4's line 4) /
C13e	§33	WORDING	/ 5. The decision has a sourced basis /	/ 5. The decision has a sourced basis (§4's line 3, audit evidence as a by-product) /
C18	About	WORDING	which is carried here unchanged;	which is carried here unchanged except for three sentences listed in the corrections table;

#	Where	Class	old	new
C19	§20	ADDED	never spends without a signal, never reports	never spends without a signal (the paid morning run is the one scheduled spend, outside the loop), never reports
C10cap	Figure 10.1 caption	WORDING	Figure 10.1. The loop as a build list. Cheap tier always on; expensive tier on signal.	Figure 10.1. The loop as a build list. Cheap tier (the ledgers) always on; expensive tier (the model runs) on signal.

Stage 8 — the figures rechecked by an internal hostile read, 30 September

#	Where	Class	old	new
C20	Figure 7.1	WORDING	‘committing, conversing’ across the HALTED border; the two-guards text out of its box under the red note	line spacing tightened, the guard text re-wrapped
C21	Figure 11.1	WORDING	‘regenerate orient → push the observer → push the always-on node’ / ‘→ push the halt state to the always-on node’ past the canvas edge	‘push both nodes’ / ‘push the halt state to the sending node’

Stage 9 — second layout of text version 11, 1 October (no word changed)

#	Where	Class	old	new
LAY1	F1, Figures 2.1, 3.1, 4.1, 5.1, 7.1, 10.1	DECISION	each figure where the source or the nesting rule put it; a figure that did not fit left the rest of its page blank (32 pp; 29.5 in blank at page bottoms, measured)	moved one or two blocks inside its own section (F1 −2, 2.1 −2, 3.1 −1, 4.1 +1, 5.1 −1, 7.1 −1, 10.1 +1); 29 pp, 8.7 in blank, no gap over 1 in
LAY2	the six tables	DECISION	every column the same width (pandoc default)	column widths chosen to minimise wrapped lines; the §31 table no longer splits, so its header prints once
LAY3	all captions	DECISION	a long caption could split across a page	caption kept on one page with its image

Stage 10 — corrections owed on the published text: seven found 1–8 October by the company’s own check and a red-team layout read (one reviewer’s reply confirmed the first and prompted the fourth); the eighth found 27 October; the sentences change in the next text version

#	Where	Class	old	new
U1	§25, NovaFabric	CITATION	Broader than the flight recorder in schema (fifteen entity types to its eight)	Broader than the flight recorder in schema (fifteen entity types; the flight recorder’s unit is one event with eight fields)
U2	§25, NovaFabric	CITATION	its claim-grounding audit, the closest published analogue to our verdict labels, is designed but not yet evaluated	its claim-grounding audit, the closest published analogue to our verdict labels, is implemented and experimental, and not yet evaluated
U3	§25, NovaFabric	CITATION	, on a single-machine research prototype.	. Its agent-level results are on a single machine; its provenance-graph and fleet-ingest results are not.
U4	§25, Agent Flight Recorder	WORDING	Their evaluation is synthetic plus five SWE-bench traces; no production deployment	Their evaluation is synthetic, with five SWE-bench traces used to validate overhead; no production deployment
U5	§25, “Agents’ accounts of their own work”	CITATION	find in a controlled benchmark across twelve models	find in a naturalistic benchmark across twelve models
U6	§26	OVERCLAIM	a hash of the 557-row cut is committed beside this release, and the record’s daily stamp covers it from then on.	a hash of the 557-row cut has been in the company’s signed record since 26 September and under the record’s public daily stamp since 27 September; it is not in this release’s own manifest. A research cut of the register (hashes and labels only) was built on 8 October and is available under terms.
U7	Figure F2 caption (page 11)	WRONG	F2. The record, live: evidence rows, actor-ledger rows, commits and adjudications arriving in order; the 13:13 grade row and the 13:07 verifier self-edit are visible.	F2. The record, live: evidence rows, actor-ledger rows, commits and adjudications arriving in order; the rows shown are from 11:36 to 11:37.
U8	§2	OVERCLAIM	We commit to this in the pilot terms; the read path is being built before the first pilot.	We commit to this in the pilot terms; the read path is not yet built, and is to be built before the first pilot.