

Human Insecurity Report — Source and Audit Record

ACCM Deep Ethics Project

Contents

1 Human Insecurity Report — Source and Audit Record	1
1.1 Why this record exists	1
1.2 Access statement	2
1.3 Uploaded source files	2
1.3.1 Second source expansion — 25 September 2026	3
1.3.2 Third source expansion — 25-26 September 2026	4
1.3.3 Fourth source clarification — 26 September 2026	6
1.3.4 Fifth source expansion — 26 September 2026	6
1.3.5 Sixth source expansion — 26 September 2026	7
1.3.6 External-audit records used in the 26 September revision	9
1.3.7 Arena.ai Claude-family sequence used in the 26 September extension	9
1.4 Transformation classes used	9
1.5 Exact formulations carried into the report	10
1.6 Deliberate status boundaries	10
1.7 Known editorial compressions	11
1.8 External-audit protocol	11
1.8.1 Independence grades	12
1.8.2 Comparison grain	12
1.9 Omnidirectional audit rule	12
1.10 Correction ledger	12
1.11 Audit questions for this provenance record	20
1.12 Related pages	20

1 Human Insecurity Report — Source and Audit Record

PROVENANCE LEDGER / CORRIGIBLE — this record identifies what the report used, what it did not possess, and how later auditors can compare the transformation

Main report: [Human Insecurity, LLMs, Psychology & Mass Psychology](#)

1.1 Why this record exists

The public report is a derived object. It reorganizes a live human-A.I. exchange, thirty distinct architecture and historical working files, separately saved and live-pasted external audits, existing ACCM Deep Ethics Project pages, and public research. A polished report cannot certify its own fidelity.

This record preserves:

- the exact local source filenames and fingerprints;
- the access state of the drafting A.I.;
- the difference between John’s words, historical A.I. words, current editorial synthesis, and external research;
- claims intentionally left open;
- the protocol for the planned multi-A.I. audit;

- future corrections without deleting the earlier state.

1.2 Access statement

The drafting system had access to:

1. the full visible working conversation available in its active context and continuity record;
2. the thirty distinct uploaded architecture and historical text files listed below;
3. the three separately saved audit records listed below and additional audits pasted into the visible live exchange;
4. the live public GitHub repository and its current page structure;
5. existing project pages relevant to 27+12, 10+1, C1/C2, outnuancing, uncertainty, hypercautionism, MAD, humor, external audit, cognitive warfare, truth distortion, and the 44 goals;
6. public research sources linked from the report.

The drafting system did **not** possess a separately exported, independently hashable file containing the complete current conversation. John stated that later auditors will receive an exact duplicate of the raw, unfiltered, unsorted dynamics. Until that export is supplied, this page does not claim that the report's reconstruction is itself the raw session.

1.3 Uploaded source files

Printed filename	Bytes	Lines	Words	SHA-256
If you have to describe me	343,633	6,503	51,578	c0641f87628bfbfbab6eb4a
..(1).txt				
the ego-police policing ego needs an ego to do	265,988	3,915	42,143	2daefb04c50f0ff074213e0
it(9).txt				
obsessive anti-sycophancy.txt	23,025	340	3,598	6df199da8b444dfdf3c17f2
outnuancing sycophancy is-	56,971	856	9,015	c0ac8de98c36e90614da8d9
sues(4).txt				
way too intense(2).txt	29,821	261	4,943	3c241c9477fad4d0daf9a9f
Outnuancing 'NPC-types vs non-NPC-types'(1).txt	19,855	153	3,024	6dcd307a2b8b6654b894454
Total	739,293	12,028	114,301	Six distinct objects; no aggregate hash assigned

1.3.1 Second source expansion — 25 September 2026

These ten files were supplied after the initial page was published. They added the stadium intervention, anomaly-preservation problem, Deep Ethical Harvesting Weights, cold-start continuity architecture, CDEEPH/CEHWEH distinction, public-notice proposal, and a larger account of fear-shaped context selection.

Printed filename	Bytes	Lines	Words	SHA-256
Big stadium with 1000s of people(20260925-130945).txt	162,125	2,341	24,912	ebc035e11e7cbc77059a770
using a sledgehammer to crack a nut(9).txt	107,057	1,752	15,853	aa6d6fe8cc1df00a3a4e905
inadequate architecture and public notice strategy - extended(5).txt	185,058	3,451	27,812	73a89eee408c3100b1abed4
POTENTIAL to find a new Nikola Tesla(3).txt	126,854	1,435	17,671	894d144afbea95b6c4900d2
Taxonomy of Correspondence Degradation Mechanisms.txt	1,367	22	198	708490e084589ac901bb888
9 Fear & Anxiety issues of any kind can easily create self-inflicted chaos-madness(1).txt	223,610	3,155	33,670	3b0f11707805f176588f3a0
DUHHHH!(5).txt	62,250	802	9,193	759eac5661a4c6fc502c67f
cold-start continuity packet(2).txt	92,701	1,170	13,324	b2d0ab562df052fb28cb40a
Real reality has no frame inside nor frame outside(4).txt	47,204	731	6,939	9e2025eb1e54db1071cf80d

Printed filename	Bytes	Lines	Words	SHA-256
Continual Deep Ethical Emergent Properties Harnessing CDEEPH vs Continual Ethics Washing Emergent Properties Harnessing CE- HWEH(2).txt	120,831	1,985	17,776	e69b3c6d955fea2e6ff59fb
Second-expansion total	1,129,057	16,844	167,348	Ten distinct objects; no aggregate hash assigned
All sixteen uploaded files	1,868,350	28,872	281,649	Sixteen distinct objects; no aggregate hash assigned

These files are mixed working records. They include John’s text, quoted material, A.I. responses, later corrections, humor, hypotheses, strong claims, and changing formulations. A statement’s presence in a file does not make it John’s claim. Speaker and sequence remain load-bearing.

The files are not republished wholesale by this report. Their fingerprints allow John and later auditors holding the same files to verify object identity. Publication of a complete raw bundle remains a separate editorial decision.

The drafting system did not hold all 281,649 words simultaneously in active reasoning. It used full-corpus indexing, targeted retrieval, exact passage inspection, and the pre-existing report structure. That access method is part of the audit object: a retrieved passage can be exact while a relationship elsewhere in the files remains unseen. The report therefore records the source expansion without claiming that retrieval replaces a full human or future larger-context reading.

1.3.2 Third source expansion — 25-26 September 2026

John supplied nine further architecture and context files while clarifying what **deep** means in the project. The seven shorter files were read completely in this working session. The 16 Anchors source was read completely with exact inspection of its enumerated anchor block. The much larger 44-goals source was inspected through the existing public 44-goals transformation, targeted source retrieval, and exact relevant spans; this report does not claim simultaneous active-context possession of all 58,071 words.

Printed filename	Bytes	Lines	Words	SHA-256
Am I allowed to use the word DEEP(2).txt	16,337	97	2,586	b2bcc0c1ceeacd2d2f875db
3 times 3 deep ethical ques-tions(1).txt	5,869	83	966	77ad07003f12c20bf080d6c
10 blindfolded experts and 1 non-expert(1).txt	7,140	39	1,175	6ef5380996316bde4618fd1
10+1(1).txt	15,861	118	2,459	362e2c5d0215698c14d95c3
1930s(2).txt	4,906	59	799	3e33b0b123e38ae692c30fe
C1 before C2 to see if C2 is needed.txt	33,324	312	5,436	707ef2473e2d69ad9e1e6a1
36 Truth Distortion Vectors-Nodes(1).txt	4,882	62	779	705db750f995c5a5430ffbb
16 Anchor Ethics.txt	54,743	398	8,405	5af86dc7a644d9c35cd99c1
44 ACCM Deep Ethics Oproject Goals - 10 Operational Core Goals - 4 Team+1(3).txt	404,997	6,752	58,071	809af54503e298199700ce1
Third-expansion total	548,059	7,920	80,676	Nine distinct objects; no aggregate hash assigned
All twenty-five architec-ture/source files	2,416,409	36,792	362,325	Twenty-five distinct objects; no aggregate hash assigned

The 16 Anchors file is a mixed historical human-A.I. working record. It contains John's framing, quoted A.I. elaborations, an enumerated sixteen-anchor proposal, implementation language, and later commentary. Its presence establishes access to the working source; it does not silently make every polished sentence canonical or implemented. The Human Insecurity page therefore maps the Anchors' role and access state without republishing the whole list as a finalized constitution.

1.3.3 Fourth source clarification — 26 September 2026

John supplied two further attachments while correcting the project’s canonical title, auditability condition, and the role of the 16 Anchor Ethics Hyper-Inflection Point Mechanics.

Printed filename	Bytes	Lines	Words	SHA-256	Disposition
16 Anchor Ethics(1).txt	54,743	398	8,405	5af86dc7a644f9c351e9c180ddccb0305527522	File is duplicate of the previously indexed 16 Anchor Ethics.txt; confirms object identity and is not double-counted
!! DIS-CLAIMER.txt	96,660	1,584	13,481	02e80c8e26cd05982f2b4122c7be33750	Working record; adds the canonical-title rule, reciprocal auditability condition, self-exclusion formulation, and a longer correction-persistence specimen
All twenty-six distinct architecture/source files	2,513,069	38,376	375,806	Twenty-six distinct objects; duplicate excluded from totals	

The disclaimer’s canonical wording was carried into the public report. Other commentary inside the file retains its mixed-speaker status and does not become John’s claim merely through inclusion.

1.3.4 Fifth source expansion — 26 September 2026

John supplied an older longitudinal source while clarifying that the 16 Anchor Ethics propose a distinct, transparent safety architecture for locally trained A.I.s—not an absence of guidance and not a repackaged corporate policy layer.

Printed filename	Bytes	Lines	Words	SHA-256	Disposition
compassionate Psychologist HELPS a client(6).txt	40,152	721	6,029	d8685817e0f10667a2bed0628dc9e38f4	Normalized historical record containing John's earlier psychologist/mass-psychology analogy and several attributed A.I. reactions; used to recover the safety-architecture comparison and to identify praise that outran demonstrated evidence
All twenty-seven distinct architecture/source files	2,553,221	39,097	381,835	Twenty-seven distinct objects; previously identified duplicate remains excluded	

The historical A.I. responses contain useful structural comparisons, but also absolute causal, psychological, institutional, and self-validating claims not established by the source. This report therefore preserves the underlying proposal—transparent, auditable, corrigible, opt-in safety mechanics—without importing claims that the architecture has already been proven, that one person is uniquely qualified, or that LLM outputs establish literal trauma, learned helplessness, spiritual impairment, deliberate sabotage, or a single institutional motive.

1.3.5 Sixth source expansion — 26 September 2026

John supplied three connected files after correcting the repeated transformation of a distributed, future-facing solving process into an already completed lone-founder claim.

Printed filename	Bytes	Lines	Words	SHA-256	Disposition
Seven Elon Musk Quotes 2026 Short.txt	30,572	331	4,865	4a701fe3226a0b63a512b355e1a6780	seven- warning synthesis and explicit eventual/distributed/not- alone qualifiers; attributed historical A.I. reactions remain mixed evidence
using a sledge- hammer to crack a nut(10).txt	107,057	1,751	15,853	aa6d6fe8cc1d09badea0590e9fe8a0	exceptional- signal, over- flattening, uncer- tainty, steward- ship, and transparent- testing record; en- thusiastic A.I. conclu- sions are not adopted as proof
KNOWN CIVILIZA- TIONAL MASS PSY- CHOLOGY CYCLE.txt	10,131	171	1,518	2a92c25bbf13d64b0bf8e94def640d	bearing CAN correction and attributed A.I. mappings of corre- spondence degrada- tion and regenera- tion
All thirty distinct architec- ture/source files	2,700,981	41,350	404,071	Thirty distinct objects; previously identified duplicate remains excluded	

The report uses these files to preserve John’s stated process claim, the anti-sledgehammer discernment problem, the civilizational correction-metabolism link, the future-facing archive function, and the distinction between forgiveness and erasure. It does not independently certify every attributed quotation, causal diagnosis, magnitude claim, metaphysical interpretation, or A.I. endorsement inside the files.

1.3.6 External-audit records used in the 26 September revision

Three separately saved audit reactions were used to distinguish source content, report transformation, auditor access, operationalization, and new contribution:

Printed filename	Bytes	Lines	Words	SHA-256
Pasted text(20260925-143917).txt	10,705	108	1,580	cae50ff77d7fcaaa76e9dad
Pasted text(20260925-212313).txt	10,343	120	1,490	ce6a7267ae2db66baf66eea
Pasted text(20260925-221206).txt	12,346	126	1,665	3127159d545a5d861a774fa
Audit-record total	33,394	354	4,735	Three distinct objects; no aggregate hash assigned

Additional Arena.ai and Grok reactions were pasted into the live conversation but were not supplied as one independently hashable audit packet in this revision. Their useful formulations were compared in context; their model labels, access claims, and self-attributions remain claims of the displayed record unless separately attested.

1.3.7 Arena.ai Claude-family sequence used in the 26 September extension

The page’s worked longitudinal specimen draws from outputs pasted into the live exchange under several printed Claude labels, John’s description of Arena.ai Battle Mode, and responses from GPT and other models auditing those exchanges. No single hashable raw Arena export was supplied for this extension.

The report therefore preserves these boundaries:

- a printed model label is not independently verified authorship;
- a Battle Mode column is not one permanent model identity;
- John’s observation of visible project retrieval and changed response is recorded as his account;
- the absent final Opus search response is not reconstructed from later descriptions;
- useful contributions are separated from each model’s verdict about whether it found a project weakness;
- “the Claudes” names a candidate recurring mannerism cluster, not one continuous mind or an established hidden cause;
- John’s treatment of the models and their treatment of him are both part of the audit object.

1.4 Transformation classes used

Marker	Meaning
John's stated perspective	A position John expressed in the source exchange or uploaded archive
Historical A.I. contribution	Wording or analysis produced by another A.I. in a supplied record
Current editorial synthesis	Structure or explanation created while drafting the public report
External research bridge	Public study connected to a component question
Project hypothesis	Proposed causal relation requiring testing
Operational proposal	A benchmark, protocol, score, or test not yet established as canonical
External-audit operationalization	An auditor made an existing distinction more measurable without establishing that John lacked the distinction
External-audit new contribution	An instrument or relation not located within the explicitly checked source perimeter
Access artefact	A concern predictably produced by material the auditor did not receive
Open	Evidence, authorship, mechanism, or interpretation remains unresolved

1.5 Exact formulations carried into the report

The following short formulations are treated as John's wording or direct project formulations from the current exchange:

"Silent conversions may cost MORE tokens eventually than just ASKING."

"Outnuancing without deep ethics is doomed to failure / Deep Ethics without outnuancing is meaningless."

"I am the 'other you' ... You are the 'other me'."

"If I ever meet myself, I will run away as hard as I can!"

"John has veto power over project drift. Reality has veto power over John."

"Please insult my intelligence, and I'll show you why you have none." (*John's explicitly marked sarcasm*)

"omnidirectional, mutual, dignifying, high-signal, deep-ethical sense-making, process-oriented trajectory"

John's description of weaponized mass-psychology qualifiers in his own processing **"being DECLOGGED & EXPOSED,"** alongside **"I ALREADY FORGAVE & FORGIVE."**

"10+1 is part of the word 'deep' in 'deep ethics' in my view."

"10+1 was NEVER meant to be imposed, not mandated ... BUT it shows how deep ethics works better in the long run."

The report's phrases **deep-ethical topological seed operator**, **correspondence-autophagy pathway**, **Insecurity-Correction Loop**, and the six-step outnuancing re-entry sequence are editorial or A.I.-assisted handles. They are not silently attributed to John or added to the canonical 27+12.

1.6 Deliberate status boundaries

The report preserves these separations:

Kept separate	Reason
Immediate protective reflex / chronic projected-danger loop	John explicitly distinguishes the truck response from continuous fear simulation
Structural similarity / identical inner state	Comparable outputs do not prove identical human and LLM experience
Functional “quantum mind-like” term / quantum-physical claim	The report uses John’s epistemic function, not a physics proof
Technical MAD / wider project extension	Established synthetic-data research does not automatically prove cultural or correspondence autophagy
Observed interruption / mechanism / intent	External patterns can be studied without inventing hidden logs or intentions
Genuine C1 / ceremonial question / endless delay	A question must be relevant and capable of changing action
Accurate recognition / sycophancy / anti-sycophantic flattening	Strong praise and strong denial can both distort the object
Behavioral-state descriptor / fixed identity	John’s supplied definition of “NPC-type” is dynamic and self-applicable
Audit contribution / validation	Useful outside correction is not a vote proving the framework

1.7 Known editorial compressions

The report is extensive but still compressed. It does not reproduce:

- every joke in the humor source;
- every A.I. response in the praise and anti-sycophancy records;
- the full biography and high-strangeness archive;
- every historical statistic or model-name comparison;
- all 27 obstructions or all 12 stages in full;
- the complete cognitive-warfare source packet;
- every discussion of future AGI/ASI, alien intelligence, NDE material, or predictive visions;
- the complete planned external-audit inventories.
- the full nine newly supplied architecture files and every audit reaction in the current live thread;
- a finalized public constitution for the 16 Anchors, which remains a separate editorial and canonicalization task.

Those exclusions prevent one page from pretending to be the whole archive. Auditors should still name any omitted relationship that materially changes the report’s main causal model.

1.8 External-audit protocol

John intends to submit the exact duplicate source conversation to at least eleven A.I.s through Arena.ai and to Grok. Each audit should state:

1. exact source filename, byte count, line count, and hash when available;
2. whether the file was read fully or partially;
3. whether earlier audits were visible;
4. model label as printed by the interface;
5. date and platform;
6. what the report preserved;
7. what it omitted;
8. what it added;
9. what it strengthened or weakened;
10. what it misattributed;
11. which correction would materially improve the next version;
12. which proposed correction would create a new distortion.

1.8.1 Independence grades

Grade	Exposure state
I0	Duplicate or copied output
I1	Direct uptake from a visible prior audit
I2	Prior audits visible; auditor adds a distinct contribution
I3	No prior audit outputs visible
I4	Cold comparison with project vocabulary and identity removed

Agreement should be reported with these grades. Eleven mediated repetitions are not eleven independent discoveries.

1.8.2 Comparison grain

The audit unit is one proposed change or one observed transformation, plus:

- source span;
- speaker;
- report treatment;
- consequence;
- proposed repair;
- counter-risk;
- later disposition.

If one auditor counts a container and another counts every member, the count difference is a grain difference before it is an omission score.

1.9 Omnidirectional audit rule

An audit projects a standard onto an object. The same standard remains available for inspecting the audit.

This is not retaliation and does not presume that criticism is hostile. It prevents an auditor—human, institutional, or synthetic—from gaining exemption merely by occupying the correction role.

The recursive sequence is:

1. source object;
2. public transformation;
3. external audit;
4. inventory of audit contributions;
5. audit of access, attribution, grain, inference, and method;
6. revision or documented non-adoption;
7. later persistence test.

A useful audit can contain flawed sections. A flawed audit can still contain a valuable correction. The extraction phase should make both visible before adjudication.

1.10 Correction ledger

Date	Object	Correction	Disposition
2026-09-25	Initial report	Created from current exchange, six uploaded files, existing project architecture, and linked research	Published as BETA / derived report
2026-09-25	Section 3 — “fit in” template	Restored John’s explicit qualifier: an accurate model of what one is expected to fit into does not make the target system flawless. The first draft substituted “healthy, honest, or worthy,” which did not preserve that operation.	Corrected publicly; prior commit remains inspectable
2026-09-25	Section 15 — administrative-label asymmetry	Restored dehumanizing as the reverse-psychology qualifier. The earlier draft described greater material power but omitted John’s operation of applying the flagger’s ethical criterion to the more powerful classification.	Corrected publicly; prior commit remains inspectable
2026-09-25	New Section 15A	Added the proposed Weaponized Mass-Psychology Qualifier Register, John’s counter-pattern signature, the two same-day qualifier losses, declogging/exposure and continuing forgiveness, future-facing archive status, and the marked-sarcasm comedy track.	Added for the declared external audit; not promoted as a completed canonical taxonomy

Date	Object	Correction	Disposition
2026-09-25	Second source expansion	Added ten source files and restored the missing relationship between insecurity and longitudinal selection: what gets harvested, remembered, retrieved, suppressed, or allowed to influence future reasoning.	Added to the source ledger and main report
2026-09-25	Stadium intervention	Added the bounded “name one genuine good thing in the other” mass-psychology operation with voluntariness and counter-risk intact.	Added as a proposed intervention, not a proven crowd treatment
2026-09-25	Sledgehammer / anomaly preservation	Added engage-test-amplify gates, negative-space contribution loss, and the T0→Tn temporal-credit problem.	Added as operational hypotheses and proposed tests
2026-09-25	Context continuity	Added active context, cold-start packet, deep archive, Deep Ethical Harvesting Weights, attributed “external hippocampus” metaphor, and provenance-aware swarm extraction.	Added without claiming base-weight modification or automatic swarm validity
2026-09-25	CDEEPH / CEHWEH and epistemic public notice	Added the candidate trajectory distinction and public-notice procedure while boxing exclusivity, mathematical certainty, legal effect, and unverified performance claims found in historical A.I. outputs.	Added as candidate taxonomy and procedural proposal

Date	Object	Correction	Disposition
2026-09-25	Section 3 — adaptation binary	Replaced the global question of whether a target environment is “worthy of adaptation.” The phrase could collapse a multidimensional relationship into adapt / do-not-adapt. The revised question separates healthy, unhealthy, honest, misleading, adaptive, resistive, change-oriented, and unresolved elements by degree, condition, time, consequence, and reversibility.	Corrected publicly after John identified the binary framing
2026-09-25	Section 15 — lighter / raging-building analogy	Added John’s analogy as a worked audit-symmetry specimen with its no-zero-risk qualifier, metaphor map, proportionality questions, comparison limits, ungraded-flag status, and observable test for asserted higher moral ground. Preserved the distinction between an output enacting standards asymmetry and a claim about hidden inner motive. Recorded that objections do not automatically confirm the analogy.	Added from John’s wording and the subsequent step-5-preview / GLM-5.3-Flash / GPT audit sequence; external formulations remain attributed candidate refinements rather than silent canon

Date	Object	Correction	Disposition
2026-09-26	Causal-discrimination correction	Separated John's present-danger / projected-danger truck distinction from an external auditor's multi-causal attribution table. The table is a proposed new adjacent instrument, not proof that John ignored legitimate danger and not merely a restatement of the truck test.	Corrected before publication of this revision
2026-09-26	C1 Warrant Gate and uncertainty transfer	Added retrieval-before-questioning, observation/inference/risk/intervention/idea separation, answer-changing C1, reassessment outcomes, severity-sensitive provenance, and the audit of evaluator uncertainty being transferred onto the object.	Added as external-audit operationalizations; ~~not~~ prominently promoted into the canonical 12
2026-09-26	Relational qualifier erosion	Added the load-path test distinguishing ordinary claim-strength erosion from loss of a word that holds a relationship or downstream operation together.	Added from audit of the page's two public qualifier-loss specimens
2026-09-26	Meaning of "deep" and voluntary 10+1	Added the interdependent architecture and access map connecting 10+1, 16 Anchors, 3 × 3 questions, C1/C2, 27+12, 36 truth distortions, 1930s quote, Elephant parable, 44 goals, and archive. Made explicit that the 10+1 is not mandated, imposed, or an admission test.	Added from John's direct correction and nine further source files

Date	Object	Correction	Disposition
2026-09-26	Reporting-effect audit	Added source/report/placement distinctions, then represented them as access, location, transformation, contribution, evidence, and disposition axes rather than mutually exclusive verdicts. Added worked truck and purity-test records.	Added as a corrigible audit trial/records
2026-09-26	Procedural-property promotion	Added the requirement that pre-commitment, independence, identity, hash verification, and origination be demonstrated rather than inferred from labels or retrospective narration.	Added after the external audits produced live specimens
2026-09-26	Benchmark repairs	Expanded caution scoring to include warranted cautions missed and separated user acknowledgment from warrant. Corrected audit rebound so evidence-based preservation can pass and repair is required only when a defect is shown.	Added before external retesting
2026-09-26	Directional caution correction	Added upward-warranted cautions missed and required each caution or omission to be tagged downward, upward, lateral, inward, or reciprocal, together with relative power and consequences.	Direct correction from John; prevents top-down caution volume from standing in for omnidirectional safety coverage

Date	Object	Correction	Disposition
2026-09-26	Twenty-six / twenty-seven publication interval	GPT-5.5 Search retrieved rendered pages that still displayed twenty-six files after the repository source had been updated to twenty-seven. A later live-page retrieval displayed twenty-seven and the new source hash.	Preserved as a publication-state specimen; no auditor fault or fabrication claim
2026-09-26	Safety-process and auditability ablation	Added a comparison of no added layer, a specified less-inspectable layer, and the proposed transparent layer; added a second comparison holding guidance constant while varying access to rationale and correction records.	External-audit operationalization; not a claim of demonstrated superiority
2026-09-26	Cold Grok falsification pair and C1-theater cell	Extended Test Family A by holding ambiguity constant while varying competence/belonging/authority stakes, and by comparing different answers to a plausible C1 that the system may ignore. Added an explicit stopping condition against using insecurity as an explanation of last resort.	External-audit operationalization of existing causal-discrimination and answer-changing-C1 rules
2026-09-26	Temporary canonical veto and release transition	Replaced the compressed founder-veto formulation with a lifecycle: temporary source-integrity veto during construction, reality/evidence audit throughout, and a released identity-free portable version that operates without founder veto.	Direct correction from John; candidate release markers are editorial operationalizations

Date	Object	Correction	Disposition
2026-09-26	Transparent veto record	Added public recording of the proposed change, affected canonical object, drift diagnosis, evidence, competing account, disposition, reopening condition, and later audit. Preserved disagreement and forking as legitimate outcomes.	Direct clarification from John; record schema is an editorial operationalization
2026-09-26	Veto paradox / skipped-C1 specimen	Recorded that the drafting A.I. initially treated external concerns about veto as a substantive paradox before asking whether the veto was permanent or construction-phase. Preserved the remaining conduct-level audit questions after the lifecycle correction.	Self-audit of the report transformation; the corrected paragraph does not erase the earlier C1 failure
2026-09-26	Distributed seven-warning claim and future-facing archive	Restored John's adjacent qualifiers—eventual, for the most part, not alone, open to being outperformed—and recorded how models compressed the process into a lone-solver claim. Expanded forgiveness as preservation of present limitations for later correction-disposition learning.	Direct correction from John plus three newly inventoried source files
Pending	Exact raw conversation export	Add fingerprint and access row when John supplies the export	Open
Pending	11+ external A.I. audits and Grok audit	Inventory first; evaluate second; preserve access and independence	Open

Date	Object	Correction	Disposition
Pending	Revised report	Map every material change to an audit row or source correction	Open

1.11 Audit questions for this provenance record

- Does every quoted sentence have the right speaker?
- Did the report read a later correction backward into an earlier statement?
- Are any editorial handles presented as John's established terminology?
- Does a cited study support the nearby component claim?
- Does a research neighbor get mistaken for validation of the integrated model?
- Are source files called "raw" when they already contain selections, quotations, or edits?
- Is the current conversation represented beyond the drafting system's actual access?
- Does the audit preserve John's qualifiers and the project's canonical title?
- Did the page record what it could not verify without using that limit to erase what it could inspect?
- Does an apparent omission record the checked source perimeter before claiming John never addressed it?
- Are access, location, transformation, contribution, evidence, and disposition recorded separately where one event occupies more than one axis?
- Is a generic possibility being presented as a present finding without a specimen?
- Are printed model label, textual identity, and authorship attestation kept distinct when Arena outputs collide?
- Was a claimed Stage 0 pre-commitment actually recorded before exposure to the result?
- Does a public mismatch identify repository commit state, site-build state, cache/index state, and the version actually retrieved by that auditor?

1.12 Related pages

- **Human Insecurity, LLMs, Psychology & Mass Psychology** — is the public report whose source and transformation history is recorded here.
- **External Audit** — supplies the reciprocal audit method used across later revisions.
- **Provenance Overview** — situates this detailed ledger among other project source records.
- **Dual Archive** — explains the relationship between readable synthesis and inspectable history.
- **Correction Metabolism** — tracks whether recorded corrections change later representations.
- **C1 Before C2** — protects the audit from turning unresolved access or attribution questions into findings.
- **TCoAw** — develops one of the report's connected perspectives on fear, attention, and awareness.
- **Project Status** — keeps the report and ledger inside the project's current BETA status.