MERIDIAN RESEARCH COLLECTIVE

TECHNICAL WORKING PAPERS / VOLUME 001

Agreement Between Instruments

Studies in Measurement, Memory, and Structural Continuity

WORKING DRAFT - FOR REVIEW

Report MRC-WP-001 / Revision 0.1

Research record through Cycle 013

10 January 2027

Space Elevator Initiative / Lumen Lunar Communications

Noema Institute - Aletheia Program

meridianrc.org / contact@meridianrc.org

Document control

Field	Entry
Report	MRC-WP-001
Title	Agreement Between Instruments
Revision	0.1 / Working Draft - For Review
Revision date	10 January 2027
Research cutoff	Cycle 013 / 10 January 2027
Institution	Meridian Research Collective
Website	meridianrc.org
Contact	contact@meridianrc.org
Status	Technical synthesis, documentary record and proposed methods
Review scope	Measurement geometry; receiver continuity; source attribution; proposed next experiments

meridianrc.org / contact@meridianrc.org

This volume synthesizes retained work and proposes methods for review. It records no additional material loading, operational optical service, electrical acquisition or participant study. A proposed protocol is not an authorization to operate equipment.

The research record includes performed unloaded instrument observations and authored software studies. Their scope is identified within each note. Historical publication dates are retained. Preparatory calculations do not confer acceptance on the intended machines.

Revision history

Revision	Description
0.1	Initial technical review draft; synthesis through Cycle 013 and complete program note sequence.

Reading this volume

Chapters 1-7 develop the argument. The four documentary sections reproduce thirty-nine program research notes, the three program ET/ED frameworks and the retained LUM-ACK-1031 exception. Appendices A-D provide derivations and proposed protocols. Appendix E presents eight representative software cases and links to the complete online case register. The final sections supply the source catalogue, terminology and subject index.

Citation keys

Key	Meaning
M-SEI-NNN	Space Elevator research note for the numbered cycle
M-LUM-NNN	Lumen research note for the numbered cycle
M-ALE-NNN	Aletheia research note for the numbered cycle
P-SEI / P-LUM / P-ALE	Published program ET/ED frameworks
LUM-ACK-1031	Retained uncommanded-acknowledgement transcription
E1-E3	External technical sources listed separately from Meridian records

Research contributors

Contributor	Research responsibility
Mara Voss	Materials and structures lead
Jonas Reed	Test metrology researcher
Elias Chen	Optical systems lead
Leila Haddad	Timing and data provenance researcher
Nadia Okafor	Cognitive systems research lead
Tomás Ibarra	Independent verification engineer
Ruth Calder	Communications lead and collective scribe

Contributor names identify the retained program record. New synthesis and proposed methods are presented as a collective working draft; they are not attributed quotations or claims of additional individual observations.

Contents

Document control 2

Reading this volume 3

Contents 5

Abstract 9

1. The question that agreement leaves open 10

2. What it would mean to have another witness 15

3. The space elevator begins with a short strip 17

4. A receiver needs more than a surviving payload 22

5. The amplifier and the signal it adds 28

6. Continuity is a property of the record path 32

7. A bounded research agenda 35

I. Space Elevator research record 38

M-SEI-001 / Where the displacement comes from 39

M-SEI-002 / When a difference is not yet a displacement finding 40

M-SEI-003 / Which part of the strip moved? 42

M-SEI-004 / The bench that is still on paper 44

M-SEI-005 / Choosing the instrument before choosing the strip 46

M-SEI-006 / The step between zero and one millimetre 48

M-SEI-007 / The zero after unloading 50

M-SEI-008 / The travel the head cannot use 52

M-SEI-009 / The bracket on the certificate 54

M-SEI-010 / Borrowing the whole instrument 56

M-SEI-011 / The first hour on the table 58

M-SEI-012 / Two readings of the same stillness 60

M-SEI-013 / The last hour of the loan 62

II. Lumen research record 65

M-LUM-001 / Keeping the events a correction discards 66

M-LUM-002 / The event neither archive can return 67

M-LUM-003 / Two copies of event five 69

M-LUM-004 / Before the receiver makes its second copy 71

M-LUM-005 / The silence between clock witnesses 73

M-LUM-006 / The timestamp that the receiver must leave open 76

M-LUM-007 / The receiver starts counting again 78

M-LUM-008 / The last witness record 80

M-LUM-009 / A challenge outside the receiver 82

M-LUM-010 / The challenge that stays used 84

M-LUM-011 / Recovering the reply, not the challenge 86

M-LUM-012 / Who keeps the current epoch? 88

M-LUM-013 / A reply to this question 90

III. Aletheia research record **92**

M-ALE-001 / Two channels, one source . . . 93

M-ALE-002 / The cost of withholding corroboration . . . 94

M-ALE-003 / The circle that came back twice . . . 96

M-ALE-004 / A valid signature and an unfinished ledger . . . 98

M-ALE-005 / When the withdrawal reaches the ledger . . . 100

M-ALE-006 / A third account of the draw . . . 103

M-ALE-007 / The request that must not draw twice . . . 105

M-ALE-008 / A copy beyond the missing tail . . . 107

M-ALE-009 / Before the acknowledgement leaves . . . 109

M-ALE-010 / A receipt cannot identify a signal . . . 111

M-ALE-011 / The echo one sample later . . . 113

M-ALE-012 / When the inputs stop adding . . . 115

M-ALE-013 / The order in the signal . . . 117

IV. ET/ED research frameworks and retained exception **119**

P-SEI / Space Elevator: exogenous environment and interaction . . . 120

P-LUM / Lumen: exogenous interaction and contact research . . . 125

P-ALE / Aletheia: exogenous cognition . . . 130

LUM-ACK-1031 / Uncommanded acknowledgement . . . 133

A. Analytical notes **135**

B. Proposed unloaded metrology protocol **139**

C. Proposed receiver recovery protocol **141**

D. Proposed electrical source and storage protocols **143**

E. Selected software cases and online supplement **145**

F. Sources and document catalogue **149**

G. Terms used in this report **155**

H. Open review register **160**

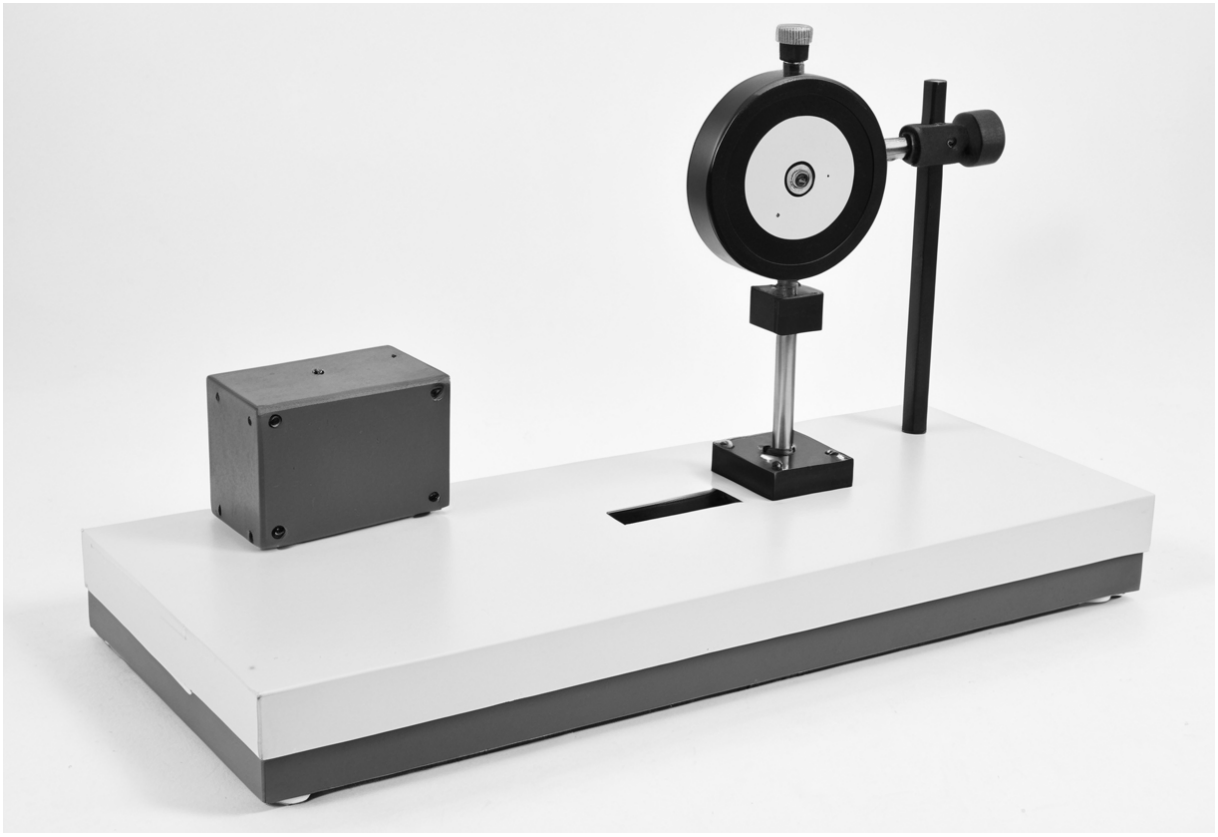
Abstract

An instrument can agree with another instrument while both follow the wrong reference. An archive can agree with its witness while both retain an earlier history. A signal can differ from a command while still being produced by that command. Across thirteen research cycles, Meridian has encountered these three forms of the same practical difficulty: the agreement available to a reviewer may survive the very change the reviewer needs to detect.

This report examines that difficulty through the surrogate-strip metrology work of the Meridian Space Elevator Initiative, the receiver-record and recovery studies of Lumen Lunar Communications, and the acquisition and source-discrimination studies of Noema Institute's Aletheia program. It combines a new technical synthesis with the complete sequence of thirty-nine program research notes, the published ET/ED frameworks and the retained uncommanded-acknowledgement record. Exogenous hypotheses, source-discrimination requirements and unresolved archival evidence receive the same scrutiny as the metrology and software findings. Mathematical developments are explanatory analyses of those records. They introduce no additional bench observations, communications service or participant study.

The principal conclusion is that independence must be specified relative to a claim and a failure mechanism. Separate instruments, copies, custodians or signatures can be useful without providing the separation required for a particular inference. A second channel must be assessed by what it can disagree about, which changes it shares, and which history it can preserve when the primary account fails. The report proposes bounded next experiments and review requirements. Loading acceptance, operational receiver freshness and electrical source attribution remain open.

1. The question that agreement leaves open



Photographic reconstruction of the shared-target comparison concept. Instrument forms and placement are explanatory, not an as-built record of KIS-61 or NQ-17 and not a photograph of the January observations.

1.1 A small movement on the intake table

On January 6, the Space Elevator group had two contacts watching a common local target. KIS-61 retained its original KB-2 bracket and portable fixture. NQ-17, a passive comparator borrowed from North Quay Instrument Cooperative, stood separately on the intake table. The displayed starting values differed: four micrometres on the head and one hundred thirty on the comparator. The difference between those starting numbers was not itself the measurement of interest. The investigators retained each starting reading and followed subsequent changes.

After an hour without commanded movement, those changes were nine and eight micrometres. Reed then moved the target by hand. Both channels added four micrometres. Attempted restoration left both two micrometres above their values at the end of the hour. Through the manipulation, the difference between channel

changes remained one micrometre. The small difference survived a movement introduced precisely to test the inference that agreement meant stillness. [M-SEI-013]

The manipulation was not a calibrated displacement. It established no absolute accuracy, thermal coefficient or strip return. It did something narrower and more useful for the immediate question: it showed that this arrangement could preserve channel agreement while the common target was intentionally disturbed. Agreement was a property of the comparison. Stationarity was a claim about the target and its reference frame. They were different propositions.

The comparator returned to North Quay on January 8. Its absence matters to the next plan. A report cannot keep an instrument available by continuing to describe it. The retained observations support a geometry question for a future loan; they do not convert the expired loan into a permanent reference capability.

1.2 Three ambitions and an awkward common beginning

The Space Elevator Initiative seeks structural knowledge relevant to a continuous route to orbit. Lumen seeks optical communications that remain useful under demanding acquisition and recovery conditions. Aletheia investigates a proposed amplification architecture whose ambition depends on distinguishing a source from the system's own intervention. These are not interchangeable engineering programs. Materials deformation, clock history and perceptual attribution have different units, risks and experimental obligations.

Nevertheless, each program first needed to inspect a path that its intended instrument might otherwise conceal. Frame travel can include motion outside a specimen. A corrected receiver stream can omit a transient. Two nominal evidence channels can be copies of one input. The early software fixtures deliberately introduced these failures. Their value lay in showing a specific inferential mistake under controlled assumptions, rather than in demonstrating the larger proposed machines. [M-SEI-001; M-LUM-001; M-ALE-001]

The subsequent work repeatedly moved the location of the question. A second archive led to a first-record witness. The witness led to its acquisition stage, custody and survival during interruption. The retained decision led to the authority that could say whether it remained current. In metrology, a second channel led to calibration, mounting, usable range and the movement of the target itself. In Aletheia, a source label led to an acquisition ledger, then to durable recovery, and finally back to the physical meaning of the signal being preserved.

This movement is not evidence that every difficulty has one cause. It is evidence that a proposed answer can depend on a boundary that the current experiment has not yet examined. The useful question is therefore concrete: what boundary makes this additional record capable of correcting the first account?

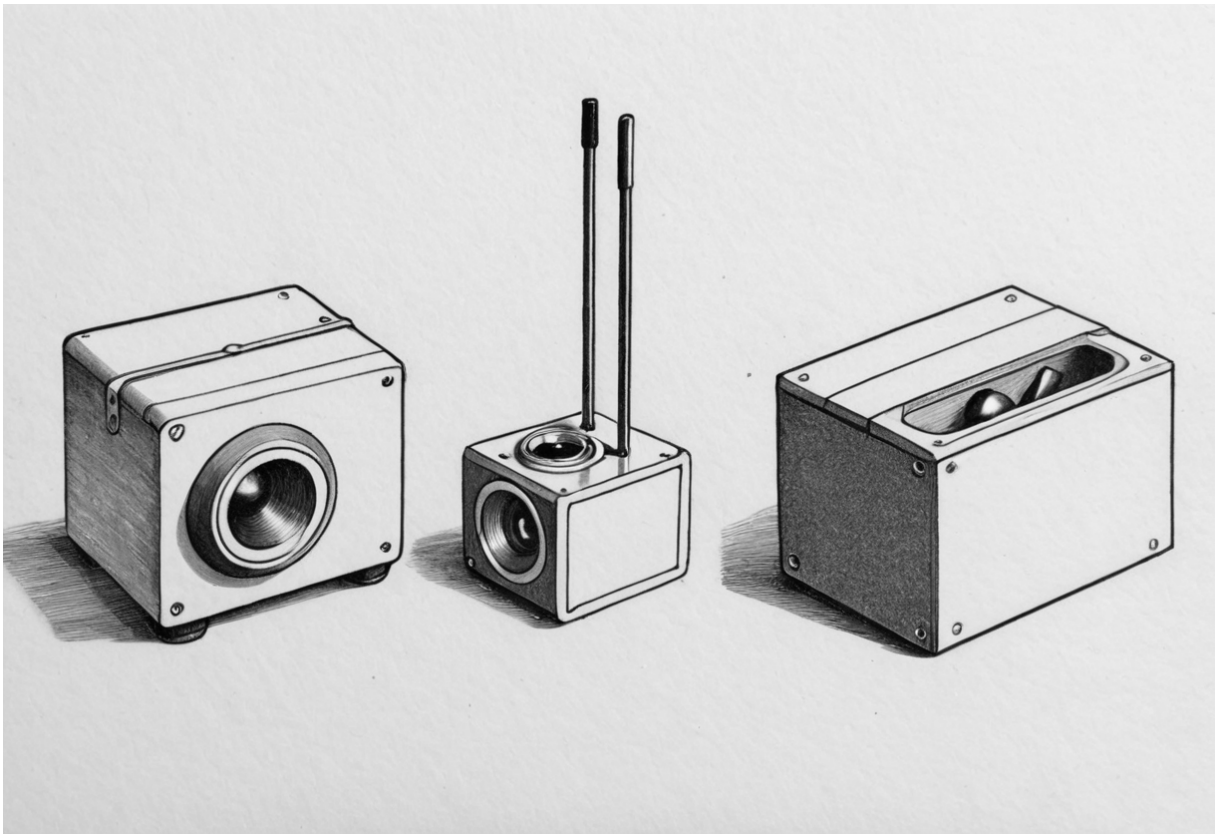
1.3 The object of this volume

The report develops a working account of conditional independence for the three programs. It does not propose a single score for trust. It asks which shared changes defeat a comparison, which records could expose those changes, and which useful claims remain possible when an additional record is unavailable. A payload may remain usable when its precise timing does not. An acquired symbol may be preserved without being promoted to independent corroboration. A head may be present on a table without being accepted for loaded testing.

The body of the report presents the synthesis. The documentary sections reproduce the program notes in their research sequence, the three published ET/ED frameworks and the complete LUM-ACK-1031 account, preserving their distinct dates and scopes. Later analysis does not retroactively make an early model more capable. The appendices derive selected counterexamples, set out proposed experiments and provide a review register. Stable note identifiers provide a route back to the underlying publications.

The working-draft designation concerns the synthesis and proposed methods. It does not make every historical result tentative in the same way. An arithmetic output can be reproducible while its physical applicability remains untested. An instrument observation can be retained while its cause remains unresolved. Review must identify which kind of uncertainty it is addressing.

1.4 Exogenous hypotheses are part of the research question



Explanatory concept for ET/ED observation across optical, field and timing channels. The arrangement is proposed; it is not a built Meridian observatory or evidence of an anomalous event.

The programs also investigate possible interactions with sources outside their ordinary operating models. The Space Elevator framework includes non-terrestrial agency, anomalous fields and non-local or higher-dimensional interaction alongside environmental and instrumental causes. Lumen treats a coherent lunar optical link as a possible interface whose raw observations must survive ordinary correction. Aletheia asks whether an amplification interface could be both receiver and beacon. These are explicit research hypotheses that shape instrumentation, source discrimination and governance. They belong to the same account as the proposed machines and the records used to assess them. [P-SEI; P-LUM; P-ALE]

Preparedness describes how deeply those questions enter a program's design. Lumen retains ET/ED level 4 and Aletheia level 5; the Space Elevator framework has no established numerical level. Those distinctions do not rank the evidence for an external agent. A design can take a possibility seriously while the cause of a particular record remains unknown.

The retained receiver-log transcription LUM-ACK-1031 makes that distinction concrete. It carries an acknowledgement when its software flag says transmission was disabled, repeats an archived diagnostic including its padding error, and retains

a final field absent from that diagnostic. The transcription and its shutdown note are available for examination; the original capture and independent timing are not. Its unresolved status is part of the data, neither an attribution nor a reason to omit the record. [LUM-ACK-1031]

2. What it would mean to have another witness

2.1 Independence has an object

Calling a channel independent without naming the intended claim leaves the important part unspecified. NQ-17 has a different origin and readout mechanism from KIS-61. That separation helps test some head-specific explanations. It does not prevent the two contacts from sharing a moving target. A receiver witness outside adaptive rewriting can expose a later edit, but an archive of that witness inside the same rollback snapshot may fail to expose an older recovered history.

For a proposed comparison, this report uses four questions. What quantity or decision is being supported? Which change would make that claim wrong? What would that change do to each observed record? What retained observation would differ between the correct and incorrect histories? These questions make independence operational. They also reveal when a favorable result depends entirely on an untested declaration such as reference_separate or authority_current.

Separation has several dimensions: physical origin, processing path, calibration, power, storage, administrative custody and time. These dimensions can interact, but none substitutes automatically for the others. Separate power does not make a queued write durable. Different signatures do not establish different draws. Different stands do not establish a motionless target. A different service name does not establish a current authority if its state can be restored with the receiver.

2.2 An observation map and an indistinguishable pair

Let h denote a complete candidate history and O(h) the information actually supplied to a reviewer. A decision rule D can use O(h); it cannot use hidden facts omitted from that information. If two histories h1 and h2 give O(h1) = O(h2), a deterministic rule must return the same decision for both. A randomized rule can vary its output, but cannot make the variation reliably track which unobserved history occurred.

This elementary statement is the report's main diagnostic tool. It is used here as an explicit derivation, not as an empirical law inferred from the number of model cases. An indistinguishable pair identifies a limit in the supplied observations. Better formatting, a stricter threshold or additional copies of the same fields cannot remove that limit. A new observation, justified restriction on possible histories, or narrower claim is needed.

The Lumen joint-rollback case illustrates the point. A valid epoch-seven history and a later history restored to epoch seven can supply identical issuer, witness and authority values. The reviewer cannot recover the missing epoch-eight history merely by seeing three sevens. In the metrology return model, a genuine common

sensor shift and a reference that moves with the specimen can supply identical channel arrays. In Aletheia, a favorable retained symbol and a hidden upstream copy can have identical recovery receipts. [M-LUM-012; M-SEI-007; M-ALE-010]

2.3 Counterexamples are design instruments

The paired histories should remain in the test set after a system acquires additional checks. They prevent a local improvement from being described as universal. A new epoch field catches stale receiver records only when the epoch authority preserves the relevant later state. Authentication catches alteration under the fixture's key assumptions, but leaves an authorized false assertion possible. A balanced acquisition order cancels a specified linear drift, while a quadratic drift remains.

A counterexample must also be described honestly. When a model contains hidden copying as author-supplied ground truth, the reviewer must not receive that truth through a helper field and then be credited with detecting it. The observer's inputs and the complete constructed history belong in separate parts of the record. Otherwise a test can appear to solve the problem by telling its decision rule the answer.

The right response to a blind spot is not necessarily to stop all use. A conditional claim can remain valuable. A model return can satisfy an allowance under the supplied reference assumption; a recovered sample can support payload service under a missing timing bracket; an authenticated receipt can establish an issuing identity within the test. The report recommends retaining that useful result together with the exact condition on which it depends.

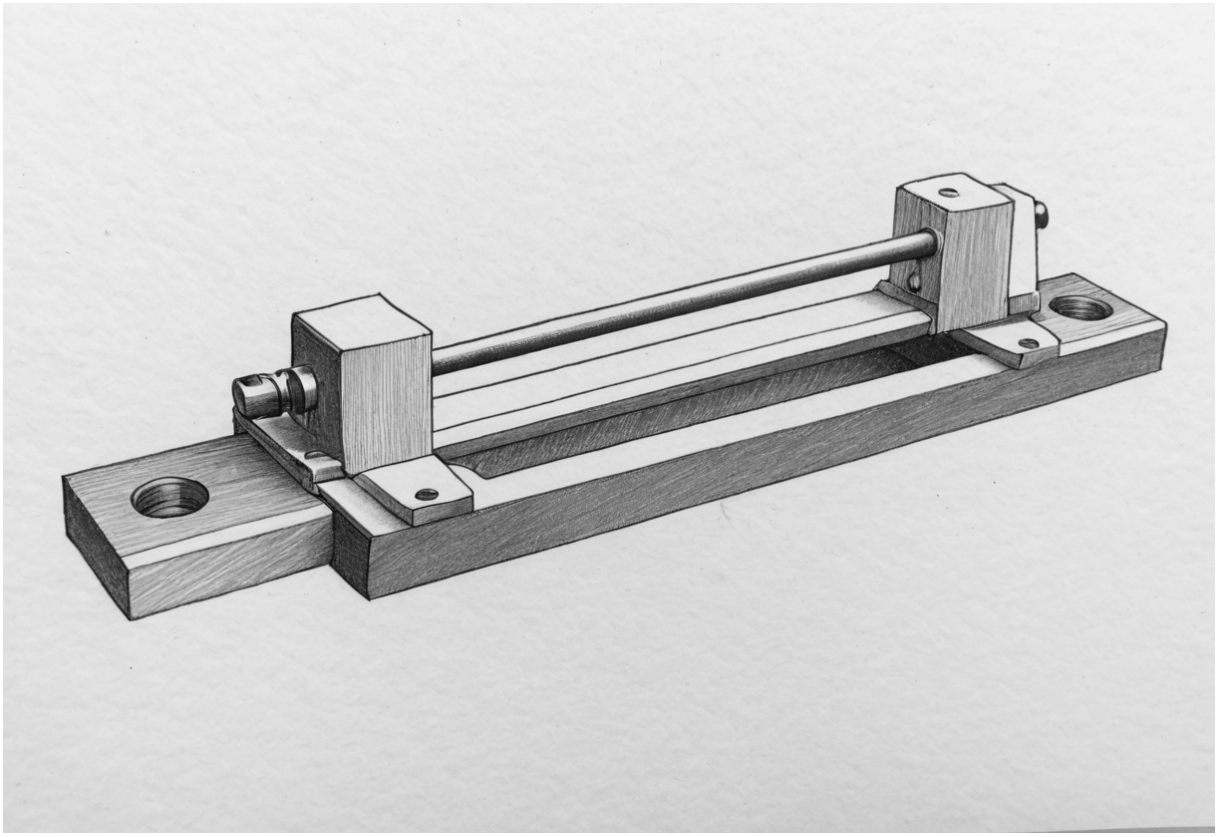
2.4 The cost of withholding

Conservative rules can withhold a true claim because the required evidence is missing. The second Aletheia provenance study contains three independently generated pairs whose tags or context do not support a reviewer finding of independence. Withholding those pairs does not prove dependence. It records a limitation of the evidence available to the rule. [M-ALE-002]

This cost should be measured in later implementation work rather than disguised as an absence of errors. A useful review table distinguishes supported, refuted, unresolved and unavailable. It also preserves what could change the classification. Recovered original evidence may justify a later decision; newly invented labels do not repair the original acquisition history.

The design obligation is symmetrical. Reconsideration must be capable of supporting a withheld claim or exposing a shared origin. If the procedure only admits favorable corrections, it becomes a mechanism for accumulating confirmations rather than a method for learning what occurred.

3. The space elevator begins with a short strip



Proposed surrogate-strip arrangement. This illustration records no assembled frame, loaded test or qualified tether material.

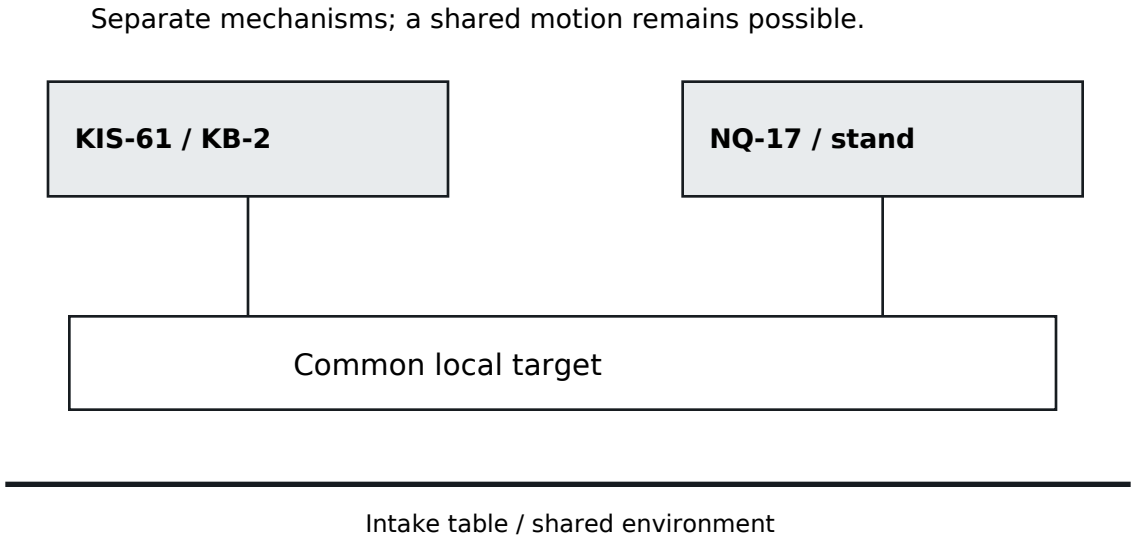


Figure 1. Logical comparison geometry. The drawing identifies shared dependencies; it is not a surveyed apparatus layout.

3.1 From a large structure to a small measurand

A continuous structural route to orbit places extraordinary demands on materials, manufacturing, inspection and operation. The current work addresses an earlier and smaller question: how would a bench distinguish extension in a short surrogate strip from movement in its grips and frame? That question cannot establish tether feasibility, but an uninterpretable displacement channel would undermine even an ordinary coupon study.

The proposed arrangement retains separate gauge extension, crosshead travel and grip movement. The gauge length used in the models is one hundred millimetres. The first software example gives four tenths of a millimetre of gauge extension and six tenths of crosshead travel, with the difference deliberately supplied as fixture slip. Neither force nor cross-sectional area is present. Stress, tensile strength and fracture behavior cannot be derived from those displacement numbers. [M-SEI-001; M-SEI-003]

The discipline of naming the measurand prevents a common expansion of scope. A channel that reports machine travel is not defective merely because it does not isolate specimen extension; the error occurs when a claim silently treats the two quantities as equivalent. The eventual instrument arrangement must make the desired separation observable.

3.2 A difference can be large for the wrong reason

The second study supplies a twenty-micrometre absolute error bound to each of two channels. Their difference then has a conservative forty-micrometre bound under those assumptions. A two-hundred-micrometre injected slip is distinguishable from zero by that bound; a thirty-micrometre slip is not. The smaller result is unresolved at this scale, not absent. [M-SEI-002]

A sixty-micrometre crosshead bias in a no-slip case can also push the residual beyond the bound. The threshold notices a discrepancy but cannot name its cause. Supplied calibration metadata causes the revised fixture to withhold the interpretation. That test demonstrates what the rule does when a calibration problem is known; it does not establish that a real calibration fault will announce itself through trustworthy metadata.

Later work supplies an independent thousand-micrometre reference step. Under an authored ten-percent scale error, a four-hundred-micrometre extension reads four hundred forty, and six hundred of machine travel reads six hundred sixty. The uncorrected difference is two hundred twenty. A truly separate step can reveal the scale error. A reference carrying the same wrong scale can approve it. The word reference therefore names a proposed function, not a guarantee. [M-SEI-003]

3.3 Planning allowances and standard uncertainty

The preferred P1 requirements envelope uses a conservative sum of allowances: one micrometre for half a resolution step, six for calibration, four for alignment, three for temperature, five for mounting and three and a half for repeat half-range. The sum is twenty-two and a half micrometres against a chosen twenty-five-micrometre displacement requirement. These are planning assumptions. They are neither a manufacturer's performance declaration nor a statistically established coverage interval. [M-SEI-005]

Substituting a coarser resolution changes the budget. Removing a missing contribution does not improve it. An unsupported mounting term makes the account incomplete. The distinction is especially important under time pressure, when an available instrument can tempt a team to revise the question until its paperwork appears sufficient.

The covariance study asks a different mathematical question. For a difference of two inputs with standard uncertainties four and three micrometres, a specified correlation coefficient changes the combined standard uncertainty. The illustrative results are seven, five and one micrometre for correlations minus one, zero and plus one. This calculation follows the correlated-input treatment identified in JCGM 100:2008 [E1]. It does not estimate the actual correlation of the installed channels or replace the conservative planning sum. [M-SEI-010]

3.4 Coverage belongs to a configuration

An instrument's nominal range is not necessarily its usable range in the offered assembly. KIS-44 lists a nominal thousand-micrometre range but an installed stop restricts useful travel to six hundred. It cannot cover the retained grid. Removing the stop would produce a configuration needing new evidence, not a clerical correction to the old offer. [M-SEI-008]

KIS-61 can reach the range, but its head-only record does not cover an arbitrary Meridian mounting. Later paperwork applies to the head on KB-2 and the workshop fixture, not to MRC-frame-01. The first conditional loan window closes. A later offer of the complete assembly makes a new route possible. Its reservation and collection are separate events; neither can be inferred from an earliest possible availability date. [M-SEI-009; M-SEI-010; M-SEI-011]

The portable assembly preserves the offered configuration through intake, yet transport and installation remain relevant. Retaining a bracket does not establish that nothing changed. Identifiers, cable seating, restraints, environment and raw readings still belong in acceptance. A complete-looking set of records can support a bounded intake decision without supporting loaded work.

3.5 What the first hours actually show

The December 22 commissioning hour has one untared display: zero, nine, eleven and eleven micrometres. It has no separately observed stationary-reference channel. The repeated last value establishes only the readings retained during that hour. It does not identify a thermal cause or certify a warm-up time. [M-SEI-011]

The December 30 paired session has changes zero, eight, ten and ten on KIS-61, and zero, seven, nine and nine on NQ-17. Their difference is zero, one, one and one. This is a new session with different starting readings. Combining it with December 22 into an average warm-up curve would erase differences in setup and available observation. [M-SEI-012]

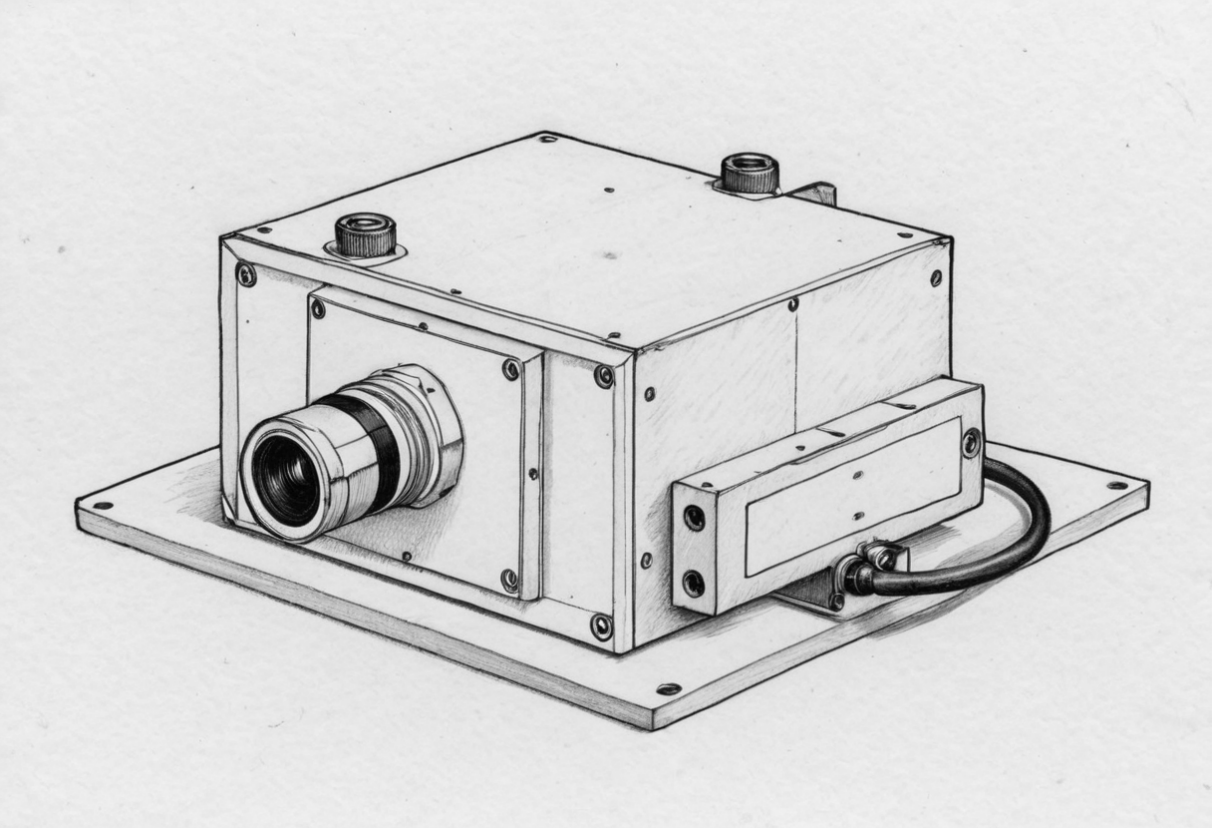
January 6 adds the target manipulation described in Chapter 1. The conclusion is geometrical before it is thermal: the comparison is sensitive to a shared target movement in a way that can preserve its small channel difference. The next reference design must say what stays fixed relative to what, and how that assertion will be tested. No strip has yet been loaded. The twenty-five-micrometre budget, six-micrometre return allowance and thousand-micrometre grid remain requirements for work ahead.

3.6 The structure as an observatory

The Space Elevator's exogenous-interaction framework treats the proposed structure as a distributed observatory as well as transport infrastructure. Precision ranging, clocks, accelerometry and field measurements would supply different ways to test a reported departure from expected motion, propagation or energy accounting. An anomaly classifier has observation access without direct authority over critical actuators. Unknown samples and data remain separated from operational systems. These are architectural requirements, not a report that such an observatory has been built. [P-SEI]

The common-target work contributes to that research question at its actual scale. A pair of channels can agree during a deliberate displacement. A future unusual movement therefore needs a discriminating geometry before agreement can carry an attribution. That result constrains both familiar mounting explanations and more unusual hypotheses; it does not choose between untested sources. Preparedness must include instruments capable of producing a contrary result, rather than granting an unfamiliar interpretation an exemption from the reference problem.

4. A receiver needs more than a surviving payload



Optical acquisition concept for explaining the record path. This is not an operating lunar terminal or a photographed Meridian receiver.

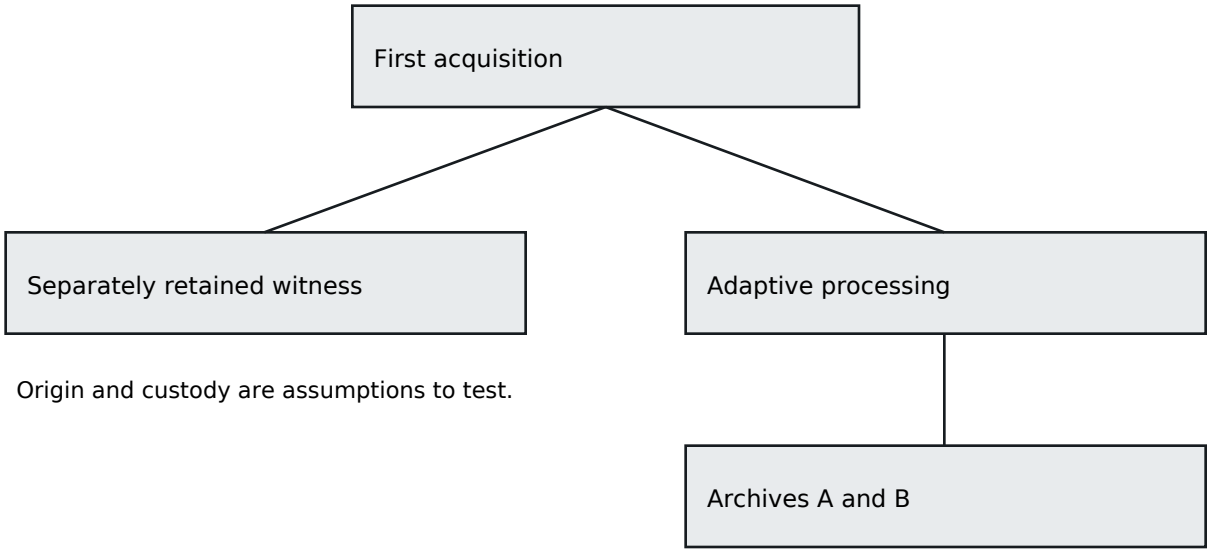


Figure 2. A witness must originate before the processing change it is intended to expose. Separation shown here is a proposed logical boundary.

4.1 Useful correction and preserved loss

Adaptive processing is intended to make a receiver useful. The preservation question is whether that useful stream can also support inspection of events the processing has removed. In the first Lumen fixture, a corrected view retains one hundred eighteen of one hundred twenty authored events and drops two deliberately placed transients. The untouched raw mirror keeps both. This is a test of a record path, not a propagation model or a measured receiver. [M-LUM-001]

Two archives improve recovery when their losses are complementary. In a later twelve-event stream, a loss at event two in one archive and event eight in the other can be filled from the surviving copies. A shared gap at event five remains a gap. An interpolation could create a plausible sequence, but it would not recover the missing acquisition. [M-LUM-002]

The distinction between a recovered event and an invented replacement matters even if the service user sees a smooth result. A system can offer an interpolated product while retaining a separate statement that the underlying observation is absent. The report does not treat interpolation as intrinsically improper; it treats unlabeled substitution as a loss of inferential control.

4.2 Missing, disputed and unwitnessed

A checksum can identify a damaged retained record under the fixture assumptions. It cannot decide which of two internally intact but conflicting payloads describes the source. When a payload is changed and its checksum recomputed, integrity within each copy can coexist with conflict between copies. The competing variants belong in quarantine rather than being silently collapsed to a winner. Equal payloads under different event identifiers also need not be duplicates: a steady detector can produce the same value more than once. [M-LUM-003]

A first-record witness moves the comparison upstream. If an edit happens before both archives are copied, archive agreement survives. A separately retained digest of the first acquisition can expose the mismatch. If the witness is recomputed from the edited source, it agrees too. The new witness is useful only when its actual origin places it before the relevant change. [M-LUM-004]

The resulting vocabulary is operational. Missing means no retained payload is available for that expected event. Disputed means incompatible retained evidence prevents a single uncontested account. Unwitnessed means a payload survives without the required first-record support. These categories have different recovery paths. None should disappear into one total labeled unsuccessful.

4.3 Time is a separate permission

An acquisition timestamp and an archive arrival timestamp describe different events. Transport delay cannot be removed by renaming arrival as acquisition. The early offset examples therefore retain sequence identity, acquisition time and arrival time separately. Current calibration can support a correction; expired or missing calibration can leave content intact while timing remains unjustified. [M-LUM-002]

Two clock witnesses constrain a linear correction between their endpoints. They do not prove linear behavior. In the supplied example, true/local pairs zero/seven and one thousand/one thousand seventeen milliseconds map a local reading of five hundred twelve to reference time five hundred. A midpoint excursion can preserve both endpoints while moving the sample. An additional interior witness exposes the authored excursion; sparse witnesses still leave unsampled paths possible. [M-LUM-003; M-LUM-004]

The later receiver models preserve payload service while withholding precise timing under an incomplete bracket, unknown between-anchor behavior or a reboot into a new counter epoch. This separation avoids two bad outcomes: discarding useful content because its timing is limited, and presenting useful content as proof of a timing claim it cannot support. [M-LUM-005 through M-LUM-008]

4.4 Consuming a challenge is an event

A one-use session proposal adds state outside the recovered receiver journal. A supplied challenge can be unused, consumed or unavailable. Matching labels do not override its consumed state. But a copied issuer snapshot can still supply the same unused input as a genuine first use. The proposal needs a protected history, not merely another field. [M-LUM-009]

The next model makes check-and-consume atomic. Two serialized callers produce one admission and one hold. A deliberately broken split check lets both callers see unused state before either changes it, producing two apparent successes. A final consumed marker cannot disclose that two admissions have already occurred. Atomicity concerns the transition, rather than the final label alone. [M-LUM-010]

A cut after consumption commits but before its reply arrives leaves a difficult state. Retrying as a new admission would contradict the retained one-use event. The model instead preserves the burned challenge for reconciliation. That choice trades immediate availability for a bounded recovery rule. It does not demonstrate durable storage merely by naming a state durable.

4.5 Recovering the old decision in the present

Bound recovery compares issuer and witness records carrying the challenge, request, payload, consumed state and epoch. It retrieves the existing decision without consuming another challenge. A different request cannot borrow the result, and a different payload cannot inherit the binding. This is a narrower operation than redoing the admission. [M-LUM-011]

An external current-epoch input can expose old matching records when it preserves a later present. If it rolls back with them, the ambiguity returns. The authenticated response study then binds an authority reply to the active realm, nonce, request and payload using HMAC-SHA256 over canonical JSON. The mechanism is a local test with public fixture keys and fixed nonce strings. RFC 2104 [E2] supplies the background concept of keyed message authentication; it does not confer deployment security on this exercise. [M-LUM-012; M-LUM-013]

A correctly authenticated answer to a different question holds. A reset verifier that reuses its nonce can accept a recorded answer again. An authority holding the key can authenticate stale or false content. The next engineering task is therefore a custody and restart design spanning key state, current epoch and durable non-reuse. Successful tag verification closes only one part of that task. Payload retention and missing clock anchors remain separate obligations.

4.6 The uncommanded acknowledgement

Record element	Retained account	Comparison still missing
Receiver state	TX_ENABLE false; request queue empty	Independent physical transmitter state
Header	Archived diagnostic match, including padding error	Original receiver capture
Timing	23:59:41-42 in transcription; clock uncertified	Independent timing witness
Payload	Sentence absent from archived diagnostic sequence	Originating received data
Shutdown note	Same sentence reported in integer status field; no template substitution reported	Originating shutdown file and exporter

Evidence map for LUM-ACK-1031. Entries summarize the retained account and identify unavailable comparisons; they do not authenticate the transcription.

LUM-ACK-1031 carries the date 31 October 2026, with its clock explicitly uncertified. Its retained transcription places TX_ENABLE false and REQUEST_QUEUE empty at 23:59:41, followed one second later by an acknowledgement class, an archived-diagnostic header match and an unavailable timing witness. The payload reads: PLEASE RETAIN THE CHANNEL AFTER THE OBSERVER LEAVES. TX_ENABLE is a software field, not independent evidence about a physical transmitter. No independently retained optical capture accompanies the record. [LUM-ACK-1031]

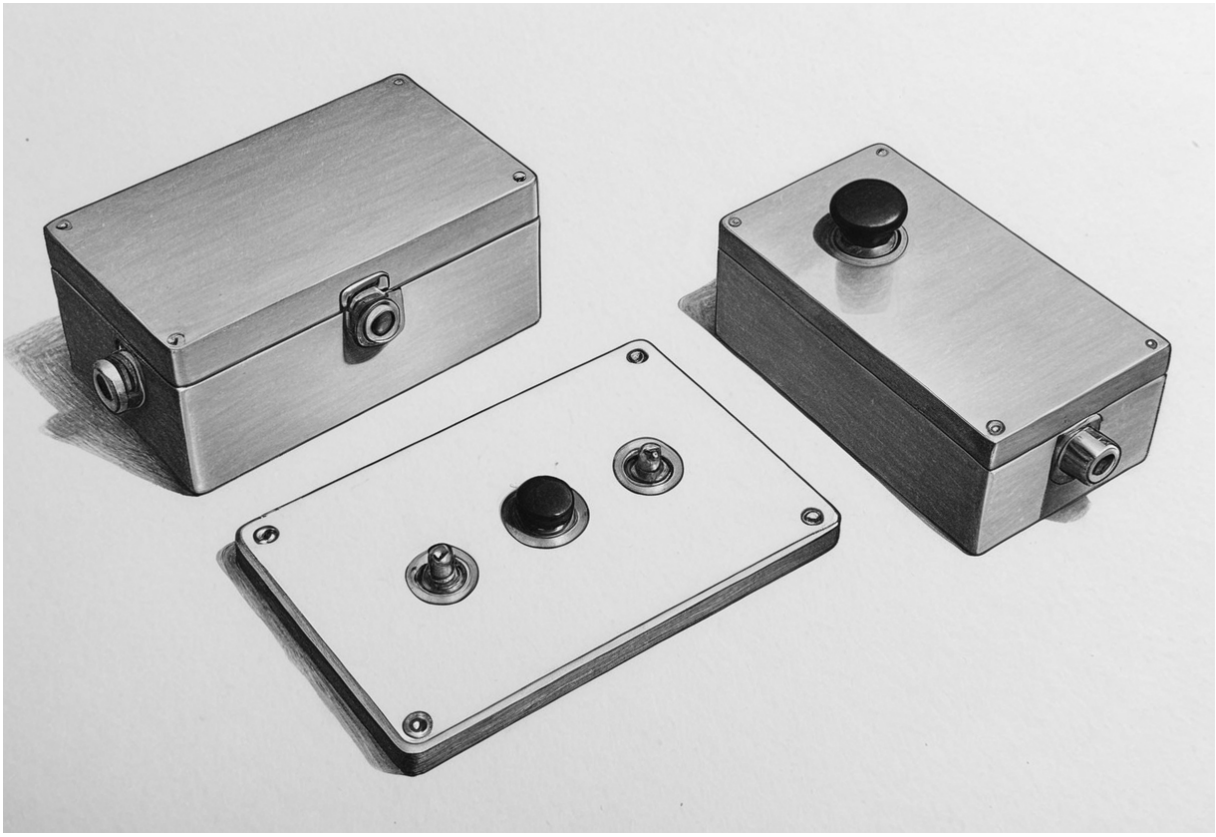
The repeated padding error gives replay, stale-buffer reuse and transcription fault specific explanatory purchase. The final field, absent from the diagnostic sequence, remains a separate question. Two marginal annotations offer different readings of observer and leaves, but neither has a verified author. The retained shutdown note adds a further discrepancy: the same sentence appears in a field described as an integer status code, and the note reports finding no export-template substitution. The originating file and exporter have not been recovered. The assertion belongs to that unverified note; its presence must remain available to review alongside the missing source material.

Lumen's published framework asks that pre-correction observations, independent clocks and other sensor channels remain available when an event might otherwise be cancelled by ordinary link control. This fragment shows why the questions must be separated. A recovered payload would not certify timing. A matching diagnostic header would not recover the absent original capture. A statement about an exporter

would not replace its file. The case calls for the original capture, an independent timing witness and the originating export path. [P-LUM; LUM-ACK-1031]

Its disposition remains preserve, isolate and do not respond. The payload grants no authority to keep a channel open, alter shutdown or initiate contact. That separation protects the ability to investigate an unfamiliar source without allowing the content under investigation to choose the experiment. It is the same observation-and-authority problem examined in the receiver recovery studies, now applied to an unresolved retained account rather than a constructed fault challenge.

5. The amplifier and the signal it adds



Electrical surrogate concept for future command-only, source-only and combined-input tests. No electrical acquisition or participant result is shown.

5.1 A repeated circle is not a second origin

Aletheia's ambition makes source attribution unusually consequential. A proposed amplification system could strengthen a signal, introduce a response of its own, or route an earlier input back to the observer. Before any participant or implant study, the software harness asks whether two nominal channels provide two origins. One circle entering an input and returning through feedback still has one source. Two independently drawn circles can look the same while having a different acquisition history. [M-ALE-001; M-ALE-003]

The distinction is not captured by counting labels alone. Different labels can share a root. The same identifier can be legitimate in separate sessions. Missing context can make an independent pair unprovable to the reviewer. A source tag is evidence only to the extent that its scope and acquisition path are supported. [M-ALE-002]

This report does not infer any extraordinary source from an unresolved classification. Unresolved means that the retained inputs do not settle the specified question. It is a reason to improve the experiment or narrow the claim, not a reserve category into which a preferred explanation can move without further evidence.

5.2 Receipts preserve assertions and their limits

The receipt models bind issuing authority, source root, session, symbol and nonce. A message-authentication check can expose a changed field under its assumptions. A receipt from an old session can be rejected by context. Neither operation makes the issuer's account true. Two valid receipts can name different roots while a separate ledger maps both to the same draw. [M-ALE-003; M-ALE-004]

Authority also changes over time. A record can remain authentic while current permission is withdrawn or disputed. The time a change takes effect and the time the reviewer learns of it are different facts. An incomplete ledger can later acquire evidence that supports independence or establishes shared acquisition. Each reconsideration appends a decision rather than deleting the earlier withholding. [M-ALE-005; M-ALE-006]

These studies concern evidence handling. They do not observe an experience or measure an implant. The limitation is material to the result: a well-preserved assertion about a circle does not establish how a future physiological or perceptual signal would have arisen.

5.3 Draw once, recover the same record

Interruption creates pressure to issue another acquisition when a reply is lost. The model instead binds a request to its original draw, keeps the retained symbol available for recovery, and prevents a missing acknowledgement from becoming permission to draw again. A replacement draw would answer a different question while appearing to complete the old request. [M-ALE-007]

An outside checkpoint can expose unequal histories, but two copies can share a missing tail. The power-loss state machine then distinguishes a queued witness write from a flushed witness commit. Separate modeled supply is useful only after the witness has something retained to preserve. A queued record lost before completion cannot be recovered from its intended destination. [M-ALE-008; M-ALE-009]

A readback binding request and symbol adds a further comparison. It still relies on supplied storage assumptions and cannot establish the signal's upstream origin. Durability and attribution are different properties. A perfect copy can preserve the wrong source with perfect fidelity. [M-ALE-010]

5.4 A command can appear in another time position

The source-control work begins with an exact scaled command echo and a different synthetic trace. A zero-lag cosine can flag exact proportionality in those arrays. It cannot establish independence when the exact match is absent. Moving the command by one sample defeats the initial comparison while retaining a simple command origin. [M-ALE-010; M-ALE-011]

The finite lag search tests shifts from minus two through plus two samples. It catches the authored one-sample delayed echo, whose strongest cosine is one at lag plus one. The lag has no physical duration because no sample rate has been supplied. Squared command and mixed input remain unresolved by the exact-linear screen. Silence has zero energy and therefore no defined cosine; assigning it a favorable independence score would turn a mathematical non-result into evidence.

Short arrays can also produce appreciable correlations without the specified echo relationship. The source-only control has a strongest tested cosine of approximately 0.595550. The figure is a result for that constructed array and search set. It is not a false-positive rate, a population estimate or a biological threshold.

5.5 The combined condition and its acquisition order

Under the authored response law y = 3c + 2s + 0.5cs, the interaction term vanishes in both isolated conditions. Command-only and source-only responses therefore look additive when examined separately, even though their sum fails to predict the combined acquisition. Subtracting the command-only trace leaves a remainder that still depends on the command. A remainder is not a source identity. [M-ALE-012]

The next study adds time. Four conditions occur in an order; a changing baseline can produce the same contrast as an interaction. In the supplied forward sequence, linear drift produces minus one half. Reversing the order produces plus one half without adding an interaction. The paired order cancels that specified linear component but leaves a nonzero result under quadratic drift. [M-ALE-013]

NIST's design guidance identifies blocking and controlled randomization as ways to address nuisance factors [E3]. The present ordered examples are explicit arithmetic counterexamples, not a randomized experiment or a guarantee of drift removal. They show why the next electrical protocol must retain timing, repeated conditions and baseline behavior alongside amplitudes. No acquired samples are authorized for subtraction by these software results.

5.6 What remains before amplification

Two experimental tasks remain distinct. A powered storage-interruption test must establish what a real witness has completed and retained. An electrical source test must establish how a real acquisition path responds to commands, external inputs and their combinations. Passing either task would not pass the other.

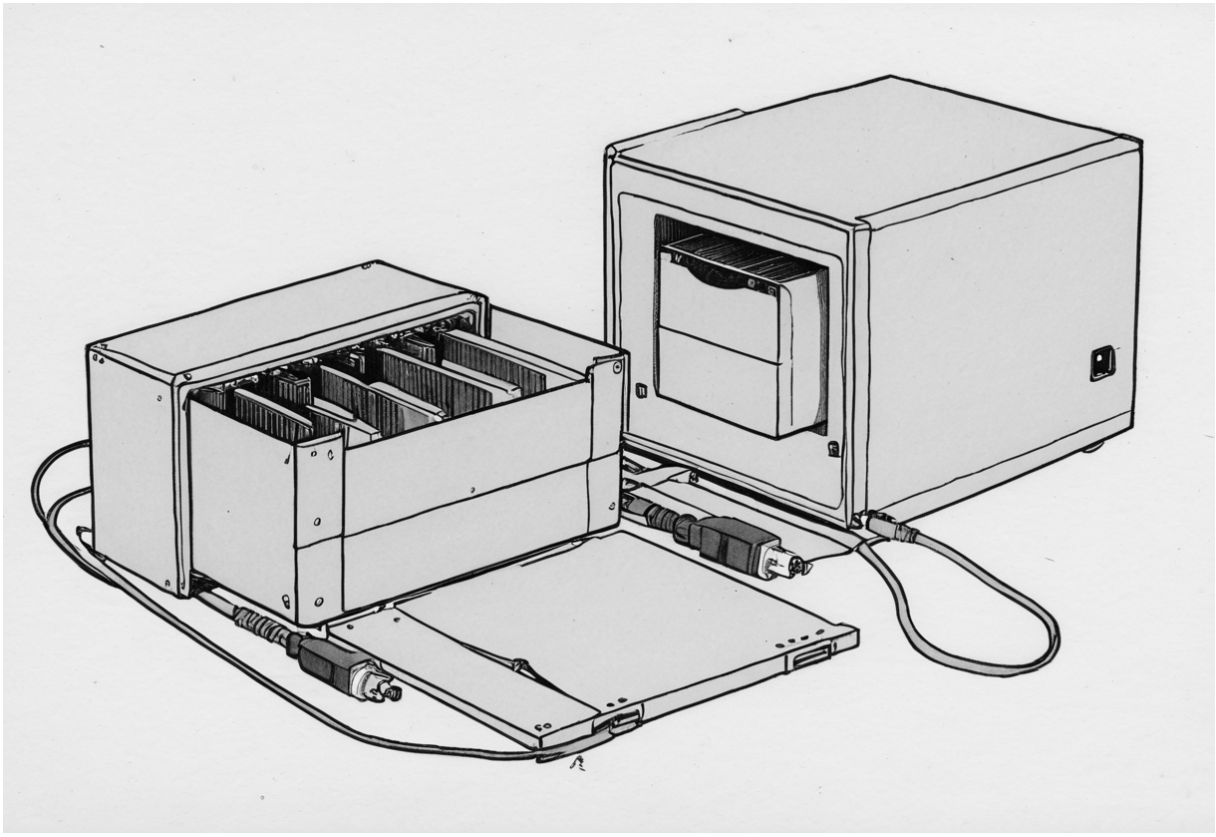
The appropriate immediate objective is a characterized electrical surrogate with known inputs, retained raw outputs and defined failure behavior. Participant work, neural attribution, stimulation and any clinical implication are outside this draft's proposed methods. The extraordinary ambition of Aletheia receives no empirical support from an arithmetic fixture; it receives a clearer account of the ordinary effects that a later claim would have to distinguish.

5.7 An interface that might receive and disclose

Aletheia's ET/ED framework makes bidirectionality an explicit design question. If an unfamiliar coupling exists, changing sensitivity might change both what a system detects and how detectable the system becomes. The published hypothesis envelope retains endogenous imagery, stimulation effects, environmental artifacts, unknown conventional processes, non-terrestrial agency and non-local coupling as competing source classes. No current participant or implant result settles that envelope. [P-ALE]

The framework consequently calls for raw neural and stimulation records compared with environmental observations, blinded conditions and independent replication. It proposes that predictive claims be registered before testing, unknown-origin data remain sandboxed, and apparent agency confer no automatic authority over stimulation or response. Those are proposed research and safety conditions, not completed trials. The software source-discrimination work bears directly on them: a label, receipt or persuasive payload can survive without supplying a second origin. A contact hypothesis needs an observation capable of distinguishing its source from the system's own command and acquisition path.

6. Continuity is a property of the record path



Concept illustration of separate recording and power paths. Physical write completion, durability and source independence remain unestablished.

6.1 Content, identity, authority and presentness

The three programs use different records, yet four questions recur. What content survives? Which event, object or request does it belong to? Who or what is authorized to support its use? Which history makes it applicable now? A successful answer to one question cannot silently supply the others.

A raw head reading can survive without its original tare context. A symbol can survive without a supported independent origin. A receiver payload can survive without usable acquisition timing. An authentic authority reply can survive without a durable guarantee against reuse. This is why a single approved flag is a poor interface for the current research. It collapses distinctions the experiments were designed to expose.

The report recommends records with separate fields for retained content, binding, authority status, time or epoch support, and unresolved dependencies. The fields are proposed review requirements rather than a newly implemented shared service. Each program must define their meaning in its own domain. A metrology calibration interval and a receiver epoch are not equivalent merely because both involve time.

6.2 Absence has several locations

A missing expected event differs from an event never known to be expected. A witness may be absent because it was not created, not completed, lost, unreadable, withheld or outside the current query. A future system should preserve these distinctions when evidence permits, rather than replacing them with a generic null. When the cause is not known, it must remain unknown.

The distinction affects recovery. A queued write that never completed needs a different response from a durable witness temporarily unavailable during an outage. A borrowed instrument returned to its owner is not a failed measurement channel; it is a resource that the next session cannot assume. A withheld corroboration decision is not a discarded input. These ordinary differences determine which next action can add knowledge.

The current studies deliberately make some missing states explicit. They do not establish that an implementation can always diagnose why a record is missing. Diagnostic labels in a model should not be presented as observed causes unless a real retained signal supports them.

6.3 Revisions must have an address

A correction needs a target: an identified result, date, version and reason. Recovered original evidence can support a new interpretation while leaving the earlier review legible. The new decision should name both the changed evidence and the part of the old claim that remains intact. A global replacement of the old account would remove the very sequence needed to assess how the institution learned.

This is especially important when different histories become compatible with the same surviving record. The report's response is to preserve the competing accounts, name the shared observations and state what would discriminate them. It does not settle a contradiction through editorial convenience. A later explanation may narrow the possibilities; the current state should remain reviewable until then.

The same principle applies to this volume. Its synthesis may be revised. Historical note identifiers, original publication dates and the difference between a supplied model state and a performed observation should remain stable. A revised chapter must not manufacture earlier access to evidence or turn a future proposal into a completed event.

6.4 Shared language without a fourth program

The convergence is sufficient to justify a dependency register and methods exchange. It is not yet sufficient to declare a fourth scientific program. The materials group needs reference geometry and calibrated coverage. Lumen needs a protected recovery authority and timing characterization. Aletheia needs powered witness completion and electrical source controls. A new umbrella label would not provide those resources.

A useful common record names a dependency, the claim it supports, the failure that could defeat it, its custodian, its time of applicability and the evidence still missing. Such a record can expose shared vulnerabilities without asserting a common physical cause. It should remain small enough to influence the next experiment rather than becoming a parallel archive of everything already known.

The practical opportunity is to share counterexample methods. Reed's reference problem can sharpen Haddad's question about the origin of an epoch. Haddad's distinction between missing and disputed can sharpen Ibarra's ledger review. Okafor's separation of an experience from its promotion to evidence can sharpen the reporting of a retained receiver event. Each exchange is useful because the differences remain visible.

6.5 Retaining an exception without choosing its cause

LUM-ACK-1031 occupies a different evidentiary position from the performed unloaded observations and the reproducible software cases. It is a retained transcription with unavailable primary material. Its fields, annotations and reported shutdown discrepancy remain research data at that level of provenance. An account can preserve their specificity without either certifying the account or suppressing it because the interpretation is difficult. [LUM-ACK-1031]

The program frameworks supply questions that this record can bear on, rather than a cause already established for it. An unexplained field does not make all three programs share a mechanism. A familiar replay explanation also does not recover missing evidence by being plausible. Revision should retain the earlier account, identify any newly recovered source and state which inference it changes. The full published frameworks and fragment are reproduced in documentary section IV with the same body typography and source treatment as the other records.

7. A bounded research agenda



Staged photographic study accompanying the end of an instrument loan. The case is illustrative; this is not a photograph of the North Quay handover or evidence about the returned equipment.

7.1 The next useful instrument is a discriminating arrangement

The immediate metrology task is to draw the reference geometry before requesting another comparator. The drawing should identify the head, bracket, base, target, stand and intended specimen marks, then specify which relative movements each channel observes. A change that moves the target and both readouts together must be represented explicitly. The design should ask what third observation or controlled reversal could separate that movement from the intended measurand.

The drawing must be accompanied by available range, mounting evidence, calibration coverage, transport history and environmental limits. The retained thousand-micrometre outward and return grid is a requirement for the relevant acceptance work; a single step or a nominal resolution cannot replace it. The six-micrometre return allowance belongs to a qualified return comparison, not to any number smaller than six.

The proposed experiment in Appendix B retains unloaded status. It is designed to improve the geometry account before a loaded surrogate-strip run. The report neither records a new loan nor changes custody. NQ-17 remains returned, and KIS-61 with KB-2 remains available only within the current recorded limits.

7.2 A receiver recovery design that includes forgetting

The next Lumen implementation should treat reset as a first-class experimental condition. It must specify where non-reuse state lives, what makes epoch advancement durable, which credentials survive a receiver restart, and what happens when the protected record is unavailable. A service that is safe only while it remembers everything needs a defined response to forgetting.

The test set should retain active-request mismatch, altered authenticated content, withdrawn authority, lost reply, duplicate caller, verifier reset and joint rollback. The last case should remain a declared limitation wherever the proposed architecture supplies no later surviving observation. A favorable decision under the valid earlier history must not be renamed a failure merely to make the test output look more skeptical; the two histories really are indistinguishable to that reviewer.

The proposed protocol asks for a result ledger that distinguishes new admission, recovery read, reconciliation hold and unavailable authority. It keeps consumption counts. It preserves missing time anchors as an independent hold. None of these requirements guarantees a deployment. They define observable behavior against which a prototype could be tested.

7.3 Electrical work that can contradict its own subtraction

The next Aletheia surrogate should supply an external electrical input and a known command through an acquisition path whose operating limits are characterized. It should retain baseline, command-only, source-only and combined conditions with amplitudes and time order. Reversal and repeated conditions should make it possible to see when a difference follows order rather than the intended factor.

No subtraction method should be accepted solely because it produces a visually quieter trace. The retained combined condition must be capable of contradicting the sum of isolated responses. A residual should be reported as a residual under a named model. If its dependence on command, lag, amplitude or baseline remains unresolved, it should retain that description.

A powered witness test should proceed separately with an identified storage path and observable completion signal. Failure cuts should be placed before local completion, before witness completion, after witness completion and before acknowledgement. Recovery must return the original request-bound value or an explicit hold, rather than silently initiating a replacement draw.

7.4 Decision gates for a working institution

Progress is not limited to a passed acceptance test. A rejected offer can make a configuration requirement concrete. A failed source screen can prevent an unjustified attribution. A lost reply can force a recoverable decision identity. These are useful outcomes when they change the next experiment rather than merely generating another cautious sentence.

The next review should ask whether each proposed control can produce a materially different retained observation under at least one plausible competing explanation. If it cannot, it may still provide redundancy or convenience, but it should not be credited with discrimination. The term independent should be followed by the failure against which independence has actually been supported.

The report leaves the three programs with specific unfinished work. No tether material has been qualified, no operational lunar timing service established, and no neural or extraordinary source identified. The record has nevertheless become more exact about what those future claims would require. The working task is to make the next witness capable of disagreeing for a reason the first witness cannot supply.

7.5 Recovering the exception's missing comparisons

The next review of LUM-ACK-1031 should seek distinct source objects: the original receiver capture, the archived diagnostic vector, an independent timing witness, and the shutdown file with its originating exporter. Compare the matching header and padding error separately from the added payload and the reported integer-field discrepancy. A failure to recover one object should remain a named gap rather than being filled by another copy of the same transcription. This is a proposed review, not a newly completed retrieval. [LUM-ACK-1031]

The three program frameworks give that review an institutional place. Raw observations should precede adaptive correction; source claims should face discriminating controls; unfamiliar content should not control its own evaluation. ET/ED preparedness supplies reasons to arrange those tests. It does not supply their results. The ordinary and extraordinary hypothesis classes remain in the same research discussion, with the observations, missing records and possible next comparisons stated concretely. [P-SEI; P-LUM; P-ALE]

I. Space Elevator research record

Documentary section

The following notes retain their original research dates and technical scope. Presentation follows the current institutional style. Later synthesis does not add measurements to these records. Result tables are reproduced where supplied.

M-SEI-001 / Where the displacement comes from

2026-10-07 / Mara Voss, Jonas Reed

A simple fixture audit separates specimen extension from clamp movement before either is treated as material behavior.

Question and method

Can fixture slip be mistaken for specimen extension?

Compare a five-step synthetic gauge-extension trace against crosshead travel, with zero slip as control and an injected slip of 0.2 mm at the final step.

Observed behavior

The final gauge extension is 0.4 mm; crosshead travel is 0.6 mm. The difference is entirely injected fixture slip.

Interpretation and limits

No physical ribbon was tested. Force, tensile strength, fracture and material quality are not represented.

The result establishes a failure mode in this deliberately simplified software fixture. It does not validate the proposed program architecture.

Next investigation

Design a physical metrology protocol with an independent gauge channel, fixture-motion channel and stated uncertainty budget.

Scope retained for review

No physical ribbon was tested. Force, tensile strength, fracture and material quality are not represented.

Source: doc-space-elevator-001 / Revision 1 / https://meridianrc.org/research/cycle-001/space-elevator/

M-SEI-002 / When a difference is not yet a displacement finding

2026-10-12 / Mara Voss, Jonas Reed

Five synthetic metrology cases show why a displacement residual needs both an uncertainty bound and usable calibration provenance before it can support a test interpretation.

From the first fixture to a decision rule

The earlier displacement study showed that crosshead travel could include deliberately injected fixture motion. Mara Voss and Jonas Reed now ask a harder question: even with separate readings, when is their difference large enough - and well enough characterized - to support an interpretation?

The follow-up uses five synthetic cases. Each current-calibration channel has an authored absolute error bound of 20 micrometres. Adding those bounds gives a conservative 40-micrometre bound on their difference. This is a fixture assumption, not a measured uncertainty budget for a real instrument.

What the controls and challenges showed

The no-slip control produces a zero residual. A 200-micrometre injected slip produces a residual outside the 40-micrometre bound; a 30-micrometre slip remains inside it. The latter case is not declared absent: the check cannot resolve it under these assumptions.

A no-slip case with a 60-micrometre crosshead sensor bias also exceeds the difference bound. A threshold-only rule would flag it as though it were the known-slip case. Calibration metadata marks that case out of specification, so the revised rule withholds the interpretation.

In the shared-scale challenge, a 10 percent error affects both channels. The residual is 220 micrometres, but a shared, unverified calibration is grounds to withhold a qualified reading. Two separate channel names do not establish independent calibration.

Proposed physical protocol

The next physical protocol calls for separate gauge, crosshead and fixture-motion records, retaining acquisition timestamps, calibration revisions and shared dependencies. It requires zero-motion and no-slip controls, a known fixture-displacement challenge, and an uncertainty budget that distinguishes per-channel and correlated errors.

The fixture-motion channel is a proposed additional check, not a third measured channel in the current two-channel simulation. Before a physical test can claim

useful independence, that channel's calibration and timing will need their own verification. The full proposed protocol is retained with the cycle's methods record.

What this does and does not settle

Reed can identify a boundary under stated assumptions, while Voss still has no physical material result. A discrepancy can justify further inspection without establishing its cause. Calibration faults were supplied to the rule in this fixture; an undetected fault would remain a vulnerability.

The next work is to challenge calibration fault detection and correlated uncertainty, then determine what instrumentation would make a physical specimen test interpretable. No tensile strength, fracture behavior or tether qualification is inferred.

Scope retained for review

The 20 um per-channel bounds and calibration faults are authored inputs. Calibration fault flags are known by construction. No physical specimen, measured instrument calibration, force or fracture is modelled.

Source: doc-sei-002 / Revision 1 / https://meridianrc.org/research/cycle-002/space-elevator/

M-SEI-003 / Which part of the strip moved?

2026-10-19 / Mara Voss, Jonas Reed

A model of a short strip between testing-machine grips shows how a calibration step and a separate grip-motion channel change the meaning of an apparent stretch.

The strip, the grips and the measurement

The Space Elevator Initiative needs a continuous structural ribbon whose defects and deformation can be inspected. A tensile test would hold a short specimen between grips and pull them apart. Crosshead travel measures the machine's movement; gauge extension measures the distance between marks on the specimen. Neither reading, by itself, says how much a grip moved relative to the strip.

We modelled a 100 mm surrogate gauge length with a third reading for grip motion. This is a displacement model for preparing a bench test. The length does not describe a manufactured tether coupon, and the model includes no force or failure calculation. The preceding displacement note supplied the question; the proposed metrology protocol supplies the channel arrangement.

Give the instrument something whose movement is already specified

Before interpreting the model strip, we supplied a 1,000-micrometre reference step. With a common ten-percent scale error, a true model extension of 400 micrometres reads 440, while 600 micrometres of machine travel reads 660. Their difference is 220. An independent reference reports the scale error: its specified 1,000 reads 1,100. Applying that correction restores the difference to 200 micrometres.

If the supposed reference shares the same wrong scale, both its specified and observed values appear to be 1,100. It approves its own error, leaving the difference at 220. We withhold that case because the reference dependency is unresolved. Two sensors and a reference can still be one calibration mistake.

A visible grip movement changes the explanation

In the known-motion control, the separate grip channel also records 200 micrometres. The readings agree within the authored 55-micrometre combined comparison bound. That makes the difference consistent with the injected grip movement, without turning consistency into a general diagnosis.

Another case retains a 200-micrometre machine-minus-gauge difference while the grip channel reads zero. It remains a discrepancy, but the grip explanation no longer fits. A stationary control also exposes a 60-micrometre crosshead zero error. We block use of that reading rather than letting an apparently stretched strip conceal an instrument offset.

What Voss and Reed will take to the bench

Mara Voss wants an ordinary surrogate strip to establish the fixture procedure before a scarce carbon specimen is entrusted to it. Jonas Reed requires an independently documented reference, a stationary check and an observable grip-motion channel. Their next step is to specify the equipment and practical control sequence, including a repeat reference check after a run.

All six cases are software checks. The 40-micrometre difference bound and 55-micrometre witness comparison bound are authored worst-case sums, not measured equipment specifications. The model corrects a supplied linear scale error; it does not discover every possible drift, qualify material or measure strength.

Retained result table

Model case	Raw difference	Control / witness	Outcome
Known grip motion	200 µm	Grip moves 200 µm	Consistent with injected movement
Common scale error	220 µm	Independent 1 mm step reads 1.1 mm	Corrected to 200 µm
Shared reference	220 µm	Reference shares the scale	Withheld
Crosshead zero error	60 µm	Stationary reading: 60 µm	Use blocked
Unexplained difference	200 µm	Grip moves 0 µm	Cause unassigned

Scope retained for review

Calibration references, third-channel motion and error bounds are authored inputs. No physical coupon, load, force, fracture or tether material was tested.

Source: doc-sei-003 / Revision 1 / https://meridianrc.org/research/cycle-003/space-elevator/

M-SEI-004 / The bench that is still on paper

2026-10-27 / Mara Voss, Jonas Reed

Six readiness challenges turn the surrogate-strip proposal into an equipment and control sequence. The actual inventory holds the test; a passing model is not an assembled bench.

A strip needs a machine and a witness

The Space Elevator Initiative ultimately needs a ribbon that can carry load through a structure extending beyond geostationary orbit. The present task is much smaller: prepare to hold an ordinary 100 mm surrogate strip between grips, move those grips, and distinguish extension of the strip from motion of the fixture. The earlier grip study supplied a useful channel layout. It did not supply equipment.

The proposed bench requires a controlled load frame, a force channel, a gauge-extension channel, a crosshead-travel channel and a separate grip-motion witness. A separately documented 1,000-micrometre reference step must be available before any displacement interpretation. All five instrument requirements remain unsecured; custody of the reference is not established. No equipment delivery, purchase or physical run is recorded.

A bounded sequence before any scarce specimen

Voss and Reed specify the order: document instrument identities and custody, check the reference, record a stationary reading, run repeat surrogate displacements, then check the reference again. A failed gate holds the sequence for investigation. Force and displacement must remain separate records; this model has no force response or fracture result.

The passing paper case supplies a reference reading of 1,000 micrometres before and after, a zero stationary offset and repeat differences of 200, 205 and 198 micrometres. Their spread is seven micrometres. It is eligible only within the paper model. The actual-readiness case has exactly those favorable numbers but unsecured equipment, and is held.

The reference has a custodian

In the October 23 methods exchange, Jonas Reed gave Leila Haddad the before-and-after reference requirement. Her timing problem gave him a sharper custody question in return: can the same operator alter the purported witness and the record it checks? The model holds a case where reference custody and controller ownership have the same label. Separate labels in the passing case are an authored assumption, not an appointed independent reviewer.

A post-run reference of 1,060 micrometres fails the authored 20-micrometre tolerance even though the initial reading passes. A stationary offset of 60 exceeds the authored 40-micrometre bound. Repeat readings of 200, 205 and 280 have an

80-micrometre spread and fail the separate repeat gate. These are distinct reasons to stop, rather than different names for one unknown error.

The next movement is an equipment decision

Mara Voss asked for a finite list she could act on. Reed now has one. The next task is to identify an available frame and candidate sensors, obtain their actual specifications, assign custody and replace demonstration thresholds with a justified uncertainty budget. No graphene sample is needed for that procurement and planning decision.

Proposed requirements and arbitrary demonstration bounds, not a certified uncertainty budget. Even the passing paper case authorizes no physical test; actual equipment is not secured.

Retained result table

Software case	Repeat spread (um)	Failed gate	Decision
paper_control	7	None	eligible_within_paper_model
actual_readiness_inventory	7	equipment_not_secured	hold
shared_reference_custody	7	shared_reference_custody	hold
post_run_reference_shift	7	reference_check_failed	hold
stationary_offset	7	stationary_control_failed	hold
repeat_disagreement	80	repeat_spread_failed	hold

Scope retained for review

Proposed requirements and arbitrary demonstration bounds, not a certified uncertainty budget. Even the passing paper case authorizes no physical test; actual equipment is not secured.

Source: doc-sei-004 / Revision 1 / https://meridianrc.org/research/cycle-004/space-elevator/

M-SEI-005 / Choosing the instrument before choosing the strip

2026-11-05 / Mara Voss, Jonas Reed

Two candidate specification envelopes turn a bench diagram into a conditional measurement budget. The more precise envelope passes the paper requirement; neither is an available instrument.

The extension we need to distinguish

The Space Elevator Initiative needs to understand how a loaded ribbon stretches and where it fails. Before approaching a scarce carbon specimen, Mara Voss and Jonas Reed want an ordinary 100 mm strip to tell them whether their gauge sees the strip itself or the movement of the frame. The earlier readiness study left them with five unsecured channels and an unassigned reference custodian.

Their first selection is a specification envelope. SEI-P1 and SEI-P2 are internal requirement sets for seeking equipment, not supplier products, quotations or instruments already in a room. P1 requires 2 micrometre displacement resolution; P2 allows 10. Both require separately recorded force, crosshead travel and grip movement. The fixture calculates extension and strain only: no force reading, section area or tensile-strength result has been supplied.

An interval rather than a reassuring average

For P1, the planning allowances in micrometres are 1 for half a resolution step, 6 for calibration, 4 for alignment, 3 for temperature, 5 for mounting and 3.5 for half the seven-micrometre repeat range inherited from the prior fixture. Adding these conservatively gives a 22.5 micrometre extension bound. P2 sums to 85. This addition permits correlated errors; it does not assume that errors conveniently cancel.

At a nominal 200 micrometre extension, the P1 interval is 177.5-222.5 micrometres. With a nominal gauge length of 100,000 micrometres and a 100 micrometre length bound, the corresponding strain interval is 0.00177323-0.00222723. That interval is conditional on every input bound being valid. Repeat half-range is a planning allowance; three repeat values do not establish a statistical coverage level.

The selection has a cost in readiness

Voss sets a 25 micrometre design requirement for this small extension. It is a chosen discrimination target, not a newly discovered material property. P1 fits within it; P2 does not. Dropping the temperature allowance yields an incomplete budget, rather than a smaller and apparently better uncertainty.

Reed keeps the reference step outside the controller record. Voss owns specimen handling and Reed owns metrology planning; the independent reference-custody role is still unassigned. A separately documented 1,000 micrometre step must be

checked before and after a run, but one step cannot establish linearity over the whole measurement range. Frame, sensors, reference and fixture access remain unsecured. The actual readiness case is held even with the favorable P1 numbers.

A request that can now be answered

The next equipment inquiry must ask for usable range, resolution, calibration evidence, mounting geometry, temperature sensitivity and access to raw readings. P1 is retained as the preferred requirements envelope; P2 is deferred for the present 200 micrometre task. If supplier data or a borrowed rig cannot support the P1 bounds, the budget must expand and the target must be reconsidered. No purchase, equipment offer or completed specimen run is recorded.

The 25 um requirement and component bounds are authored design inputs, not manufacturer data or measured uncertainty. The sum is conditional on bounds covering actual errors; repeat half-range is a planning allowance, not a coverage claim. A 1000 um reference step does not establish full-range linearity. No stress, strength or tether qualification follows.

Retained result table

Specification/readiness case	Extension bound (um)	Strain interval	Decision
P1_conditional_control	22.5	0.00177323-0.00222723	conditionally_eligible_in_model
P2_coarse_channel	85.0	0.00114885-0.00285285	hold
actual_inventory	22.5	0.00177323-0.00222723	hold
shared_reference_owner	22.5	0.00177323-0.00222723	hold
missing_temperature_bound	Incomplete	Unbounded	hold

Scope retained for review

The 25 um requirement and component bounds are authored design inputs, not manufacturer data or measured uncertainty. The sum is conditional on bounds covering actual errors; repeat half-range is a planning allowance, not a coverage claim. A 1000 um reference step does not establish full-range linearity. No stress, strength or tether qualification follows.

Source: doc-sei-005 / Revision 1 / https://meridianrc.org/research/cycle-005/space-elevator/

M-SEI-006 / The step between zero and one millimetre

2026-11-13 / Mara Voss, Jonas Reed

A bidirectional software sweep tests what the proposed displacement channel would need to reveal before it could qualify a surrogate-strip bench.

The machine and the substitution

The proposed Space Elevator bench would pull an ordinary 100 mm strip while separately observing gauge extension, crosshead travel and grip motion. A frame display cannot supply all three observations. The earlier instrument-selection note retained P1 as a displacement requirement, with a conservative 22.5 micrometre planning sum against a 25 micrometre design allowance. It did not identify a purchasable sensor.

Voss and Reed now ask what could go wrong when a nominally similar channel is substituted. The software retains P1's resolution, calibration, alignment, thermal, mounting and repeat allowances. Replacing its 2 micrometre resolution requirement with 10 micrometres increases the half-resolution contribution from 1 to 5; the total becomes 26.5 micrometres. Removing the mounting bound makes the total incomplete rather than smaller.

A range, not an endpoint

The new trace visits 0, 100, 200, 400 and 1000 micrometres in both directions. Those authored steps are comparison inputs, not movements of an installed instrument. An injected 30 micrometre bow at the 200 micrometre step leaves zero and 1000 unchanged. Two endpoint-only cases therefore supply identical readings despite different interior truth; both are held for missing range and return coverage.

Sampling the full grid exposes that bow. A separate return trace adds an 8 micrometre offset and exceeds the authored 6 micrometre calibration term. These are different failures: one is missed by the chosen observations, the other by omitting the return direction. A clean grid still says nothing about unsampled points unless an additional covering model is justified.

The reference can inherit the same error

In the false-separation case, the candidate readings and reference both inherit a 30 micrometre shift. Their residual is zero, and their supplied custody declaration says they are separate. The model gives the same favorable envelope decision as the clean control. It does not discover the hidden dependency.

The next equipment request must therefore ask for usable range, mounting evidence, calibration coverage and a reference with a demonstrably different origin. Reference custody remains unassigned, and the frame and channels remain unsecured. None of these software decisions authorizes a strip run, force measurement or strength claim.

Retained result table

Case	Conditional extension bound (um)	Maximum sampled residual (um)	Decision
complete_grid	22.5	0	conditional_model_envelope
coarse_resolution_swap	26.5	0	hold_in_model
missing_mount_bound	Incomplete	0	hold_in_model
endpoint_only_control	22.5	0	hold_in_model
endpoint_only_hidden_bow	22.5	0	hold_in_model
grid_exposes_bow	22.5	30	hold_in_model
return_hysteresis	22.5	8	hold_in_model
undeclared_shared_reference	22.5	0	conditional_model_envelope

Scope retained for review

Synthetic step traces and substituted requirements only. The 22.5/26.5 um sums depend on authored covering bounds. A complete grid cannot authenticate a falsely declared independent reference, or prove behavior between its steps. Force, strength and physical readiness are not established.

Source: doc-sei-006 / Revision 1 / https://meridianrc.org/research/cycle-006/space-elevator/

M-SEI-007 / The zero after unloading

2026-11-21 / Mara Voss, Jonas Reed

A restart study for the proposed strip bench separates a zeroed display from the displacement retained in its raw record.

A strip must come back from a load

The Space Elevator program proposes a bench that would pull a 100 mm surrogate strip while keeping gauge extension, frame travel and grip movement separate. Before any strength claim, the displacement channel must show what remains after unloading. The bidirectional acceptance study established the need for return readings; this study asks what survives when the operator re-zeros the display.

The seven software cases use three authored positions: initial, loaded and returned. The clean gauge trace is 0, 200, 0 micrometres. A stationary comparison channel remains at zero. The model subtracts the reference change from the gauge change between the first and last samples, then compares the residual with the existing 6 micrometre calibration allowance. This is a proposed return check, not an installed instrument calibration.

Tare does not unload the specimen

A second trace ends at 12 micrometres. Setting the display to zero leaves that raw difference unchanged, so both the untouched and tared versions are held. If the initial raw reading has been discarded, the model cannot reconstruct the return and reports an unavailable result. It does not substitute the displayed zero.

When both gauge and stationary reference shift by 12 micrometres, subtraction gives zero. That can account for an authored common sensor shift. A hidden-coupling case supplies exactly the same numbers, but the reference has moved with the simulated specimen. The reviewer cannot distinguish the two histories and gives the same favorable conditional result. A reference name is insufficient evidence of stationary behavior.

A completed procedure, an unperformed run

Reed has completed the requested acceptance sequence as a reviewable bench procedure: preserve the initial raw readings, retain every tare, observe the unloaded return, and stop if reference coverage is absent. Voss has incorporated that sequence into the equipment evidence request. No frame, displacement sensor or independent reference custodian has been secured.

An observed residual would still require separate checks for grip movement, thermal drift, sensor behavior and specimen change. These fixtures identify no material damage. The procedure cannot authorize force or strength measurements, and the earlier full-range and mounting obligations remain.

Retained result table

Case	Corrected return (um)	Decision
clean_return	0	within_model_return_allowance
common_sensor_shift	0	within_model_return_allowance
residual_extension	12	hold_return_offset
tare_with_raw	12	hold_return_offset
tare_without_raw	Unavailable	hold_missing_raw
reference_outage	Unavailable	hold_reference
hidden_coupled_reference	0	within_model_return_allowance

Scope retained for review

Authored gauge/reference arrays, a proposed stationary comparison and a 6 um return allowance only. Neither the reference independence nor its physical coupling is established. A return offset does not identify strip damage, grip motion or thermal cause. No force or strength result, procured instrument or performed strip run.

Source: doc-sei-007 / Revision 1 / https://meridianrc.org/research/cycle-007/space-elevator/

M-SEI-008 / The travel the head cannot use

2026-11-28 / Mara Voss, Jonas Reed

Two supplier records narrow the instrument search for the surrogate-strip bench; neither authorizes a run.

The first records from a workshop

The proposed Space Elevator surrogate-strip bench needs to distinguish strip extension from grip movement and frame travel. Its acceptance procedure samples outward and return positions through 1000 micrometres. On November 25, Reed sent the range and mounting requirements to Kestrel Instrument Service, a workshop servicing displacement instruments. Two retained stock sheets arrived the following day. This is the first candidate-specific record review; the instruments remain with the workshop.

KIS-44 carries a nominal 1000 micrometre range. Its stock sheet limits usable travel to 600 micrometres because of an installed stop. The last two acceptance positions therefore cannot both be reached: 400 is available, 1000 is not. The head is rejected for this procedure. Removing the stop would create a different configuration requiring its own evidence; it is not an accepted workaround.

A candidate that leaves work to do

KIS-61 lists 2000 micrometres of usable travel and 1 micrometre resolution. It survives the range screen. Its calibration record covers the head alone, however, not the bracket, cable routing or repeated mounting on the proposed frame. Those omissions leave the mounting and alignment terms in the displacement budget unsupported.

The workshop lists December 8 as the earliest possible loan date, subject to availability. No loan has been agreed and no equipment has arrived. Voss has retained the candidate and requested evidence covering the assembled configuration. Reed leads the record review under the assignment made in the preceding cycle. Neither becomes the stationary-reference custodian by reviewing a stock sheet.

The original gate remains

The restart acceptance procedure remains unperformed. Initial raw readings, separate gauge and grip channels, stationary-reference custody, thermal bounds and the full outward/return sequence are still required. The 25 micrometre requirement and 6 micrometre return allowance have not been relaxed to fit either offer.

The two rows below classify supplied records rather than measured head performance. They establish no strip strength, accepted calibration or operational tether equipment. The immediate next task is a written mounting arrangement and its evidence, before any loan decision or bench run.

Retained result table

Candidate	Usable travel (um)	Assembly calibration	Disposition
KIS-44	600	Absent	reject_range
KIS-61	2000	Absent	hold_mounting_evidence

Scope retained for review

supplier statements, not independent calibration or physical measurements. Range and resolution alone do not establish the 25 um budget. No stationary-reference custodian, mounting/thermal/alignment coverage, force system or bench run. KIS-61 availability is conditional, not a loan commitment.

Source: doc-sei-008 / Revision 1 / https://meridianrc.org/research/cycle-008/space-elevator/

M-SEI-009 / The bracket on the certificate

2026-12-07 / Jonas Reed, Mara Voss

KIS-61 mounting coverage belongs to a workshop fixture; its conditional loan window closes before any agreement.

A head attached to something

The proposed surrogate-strip bench needs to separate strip extension, grip movement and frame travel. The candidate review retained KIS-61 because its 2000 micrometres of usable travel reached the acceptance positions. Its head-only calibration left the mounting question open. Reed requested coverage of the intended assembly and confirmation of availability on December 3.

The December 4 reply supplies coverage for the head on Kestrel Instrument Service's KB-2 bracket and workshop fixture. Meridian's proposed mounting is MRC-frame-01. The two names do not describe interchangeable assemblies. No evidence in the reply transfers the calibration to the proposed frame, its cable routing or repeated installation.

The December window closes

The workshop has also assigned KIS-61 and its bracket to an existing calibration customer. December 8 had been an earliest possible date, subject to availability; it was never a reservation. That window is withdrawn. Following the December 5 review, Voss declines the current offer. The instrument remains at the workshop.

Kestrel will review stock again on December 15. This is a date to ask another question, not an agreed delivery. A later offer will require a new resource decision and evidence for the configuration actually offered. The present cycle ends with a closed offer rather than a loan awaiting arrival.

What remains before a strip can be loaded

The acceptance procedure still requires the 1000 micrometre sampled grid, a 25 micrometre displacement budget and a 6 micrometre return allowance. Those limits have not been changed to suit the stock room. KIS-44 remains rejected under its earlier range review.

The table classifies the reply, not head performance measured by Meridian. Stationary-reference custody, thermal and alignment coverage and physical acceptance remain open. The next equipment search must specify the mount and head together; no material strength or operational tether result follows from this review.

Retained result table

Candidate	Covered mounting	Availability	Disposition
KIS-61	KB-2	withdrawn	decline_current_offer

Scope retained for review

documentary review, not independent calibration. No received instrument, stationary-reference appointment, physical bench run or strength result. The 25 um budget, 6 um return allowance and 1000 um sampled grid remain unchanged.

Source: doc-sei-009 / Revision 1 / https://meridianrc.org/research/cycle-009/space-elevator/

M-SEI-010 / Borrowing the whole instrument

2026-12-16 / Jonas Reed, Mara Voss

A complete fixture is reserved while a covariance calculation keeps standard uncertainty distinct from acceptance limits.

Two channels make a difference

The surrogate-strip bench is intended to separate extension from motion in the grip and frame. Its corrected return subtracts a reference reading from the gauge reading. The earlier acceptance procedure retains a 25 micrometre displacement budget, a 6 micrometre return allowance and a 1000 micrometre sampled grid. None of those limits is relaxed here.

The GUM treatment of correlated inputs informs a new sensitivity calculation. For y=x1-x2, the combined variance is u1 squared plus u2 squared minus 2 rho u1 u2. With illustrative standard uncertainties of 4 and 3 micrometres, rho values -1, 0 and +1 give 7, 5 and 1 micrometres. These are authored inputs, not estimates from KIS-61 measurements.

A small number is not an acceptance

The sign matters because the measurand is a difference: positive correlation can reduce its uncertainty. This calculation does not show that the two installed channels are correlated, or that a common physical bias is harmless. Shared motion can cancel in a difference while remaining relevant to the absolute displacement being investigated.

Standard uncertainty describes dispersion under a stated model. It is not the earlier worst-case displacement allowance, a maximum observed return, or a material strength result. The next physical plan must retain raw channels and establish which environmental and reference terms are shared.

A new offer for the complete assembly

The December 15 stock reply offers KIS-61 on KB-2 together with its complete portable workshop fixture. This is a new offer after the earlier window closed. On December 16, Voss reserves the assembly for collection no earlier than December 21. The supplier still holds it; no delivery has occurred.

Keeping the offered assembly together avoids assuming its paperwork covers MRC-frame-01. Transport and installation can still change its behavior. The reservation therefore creates an equipment route, while installation review and independent stationary-reference custody remain prerequisites to a physical run. No strip is loaded during this cycle.

Retained result table

Authored correlation case	rho	Standard uncertainty (um)
negative_correlation	-1	7.0
uncorrelated	0	5.0
positive_correlation	1	1.0

Scope retained for review

Authored uncertainty inputs, not measured instrument statistics or acceptance bounds. supplier records are not independently verified; no equipment delivered, reference custodian appointed, strip loaded or material strength measured.

Source: doc-sei-010 / Revision 1 / https://meridianrc.org/research/cycle-010/space-elevator/

M-SEI-011 / The first hour on the table

2026-12-24 / Jonas Reed, Mara Voss

The complete KIS-61 assembly arrives; an unloaded readout changes by 11 micrometres before the bench has an independent stationary reference.

The assembly enters custody

The surrogate-strip bench is intended to distinguish extension in a strip from movement in its grips and frame. On December 21, Reed and Voss collect the complete KIS-61 head, KB-2 bracket and portable workshop fixture described in the reservation note. This is the new whole-assembly agreement; the earlier declined offer remains closed.

On December 22 the assembly is installed on a separate intake table, retaining its original mounting rather than transferring to MRC-frame-01. Identifiers match the offered configuration. With power off, Reed checks cable seating and free travel after removing transport restraints. No strip or force system is attached. Voss authorizes unloaded power-on observation only.

A display moves while no travel is commanded

The untared head display reads 0 micrometres at power-on, 9 after ten minutes, 11 after thirty and 11 after sixty. No stage travel is commanded after setup and no display tare is applied during this hour. These are displayed readings from a commissioning log, not a corrected or calibrated displacement measurement.

There is no separately observed stationary-reference channel. The 11 micrometre change could include head or electronic drift, motion in the fixture, or environmental effects. The log does not distinguish them. The final repeated value records this hour; it establishes no subsequent stability or cause.

Custody changes; the run gate stays closed

The return procedure still requires raw gauge and stationary-reference readings. Its 6 micrometre allowance cannot be applied to an unreferenced power-on display change as though this were an outward-and-return strip test. The 25 micrometre displacement budget and 1000 micrometre sampled grid also remain intact.

The assembly is now present and the intake inspection is complete. Independent reference custody, post-transport coverage, environmental bounds and full-range checks remain open. Acceptance and loaded-strip runs stay withheld. The next task is a separately originated stationary comparison, not a display adjustment that erases the first hour.

Retained result table

commissioning observation	Untared display (um)
Power-on / 0 min	0
10 min	9
30 min	11
60 min	11

Scope retained for review

performed unloaded commissioning, not a real instrument reading. Head display alone cannot separate electronic drift, fixture motion and environment. Supplier coverage is not independently renewed after transport. No stationary reference, full-range acceptance, loaded strip, measured uncertainty or strength result. The 25 um budget, 6 um return allowance and 1000 um grid are unchanged.

Source: doc-sei-011 / Revision 1 / https://meridianrc.org/research/cycle-011/space-elevator/

M-SEI-012 / Two readings of the same stillness

2027-01-01 / Jonas Reed, Mara Voss

A loaned passive comparator shares most of the first-hour change; the local target still cannot establish its own stationarity.

Another instrument, a limited loan

The surrogate-strip bench needs to distinguish extension from movement in its supports. The first commissioning hour supplied only a powered head display. Reed and Voss request a passive comparator from North Quay Instrument Cooperative on December 27, receive a loan offer the next day, and collect NQ-17 with its stand on December 29. Voss holds custody; the loan ends January 8.

NQ-17 has nominal one-micrometre divisions. No independently verified calibration or uncertainty accompanies the loan. On December 30 its separate stand is placed on the intake table with its contact observing the same unloaded local target as KIS-61. The original head and bracket remain together. No strip, force system or commanded travel is introduced.

Keep the two starting readings

The KIS-61 display begins at 3 micrometres and the comparator at 125 micrometres. Neither display is tared during the log. The table reports changes from those retained starting values at zero, ten, thirty and sixty minutes after the head is powered. KIS-61 changes by 0, 8, 10 and 10 micrometres; NQ-17 by 0, 7, 9 and 9. Their arithmetic difference is 0, 1, 1 and 1.

These are performed observations in a new session, not a reconstruction or correction of the earlier 0/9/11/11 display log. Most of the new change is shared between channels. That observation is compatible with target or support motion and shared environmental response; it does not determine a thermal cause or prove either instrument is stable.

Separate origin does not remove a shared target

The passive comparator comes from a different custody chain and uses a different readout mechanism. It still shares the target and table environment. The one-micrometre difference is not an uncertainty estimate, a calibrated strip extension or a passed six-micrometre return test. Nominal divisions do not establish measurement accuracy.

Voss permits this unloaded comparison only. Full bidirectional coverage at the existing 1000-micrometre grid, reference geometry, environmental bounds and post-transport review remain open. The 25-micrometre budget and six-micrometre return allowance stay unchanged. The loan gives the next comparison a deadline without granting acceptance.

Retained result table

Minutes after KIS-61 power-on	KIS-61 change (um)	NQ-17 change (um)	Arithmetic difference (um)
0	0	0	0
10	8	7	1
30	10	9	1
60	10	9	1

Scope retained for review

Authored performed stationary comparison only. NQ-17 has nominal 1 um divisions but no independently verified calibration or uncertainty budget. Its separate stand shares the local target and table environment. Agreement cannot establish absolute stationarity, transfer post-transport coverage or satisfy the full bidirectional grid. Loaded and acceptance runs remain withheld; 25 um budget, 6 um return allowance and 1000 um grid persist.

Source: doc-sei-012 / Revision 1 / https://meridianrc.org/research/cycle-012/space-elevator/

M-SEI-013 / The last hour of the loan

2027-01-08 / Jonas Reed, Mara Voss

The borrowed comparator follows a deliberate common-target movement, then returns to North Quay without granting the bench acceptance.

A target both instruments can follow

A tether-material bench needs extension readings that can be separated from motion elsewhere in the apparatus. The borrowed NQ-17 comparator gives Reed and Voss one more session before its January 8 return. The first paired hour showed mostly shared change; it did not establish that the local target was still.

They plan the final unloaded session on January 4 and perform it January 6. KIS-61 and NQ-17 retain their instrument mounts on the intake table. Initial displays are 4 and 130 micrometres. Neither display is tared during the session; the table below subtracts the retained starting readings. No strip, applied force or calibrated travel is introduced.

Agreement follows the target

At zero, ten, thirty and sixty minutes the displayed changes are 0/7/9/9 micrometres on KIS-61 and 0/6/8/8 on NQ-17. These are new performed observations, not revised numbers for December 30. After the stationary hour, Reed moves the common local target by hand while Voss records both readings. Each channel adds four micrometres. Attempted manual restoration leaves each two micrometres above its reading at minute sixty.

The arithmetic difference stays at one micrometre after minute ten, including while the shared target is deliberately moved. This is a concrete reason that agreement cannot establish target stationarity. The movement is uncalibrated, the absolute displacement remains unknown, and neither scale is thereby certified. The residual two micrometres is a local target offset after a hand adjustment, not a measured strip return or a passed six-micrometre return test.

The case leaves the table

On January 8 Voss and Reed return NQ-17, its stand and case to North Quay Instrument Cooperative. The handover closes the short loan on time. No renewal is sought for this session; subsequent Meridian work cannot assume the comparator remains available.

KIS-61 and KB-2 remain at Meridian. Loading and acceptance stay withheld. The 25-micrometre displacement budget, six-micrometre return allowance and 1000-micrometre bidirectional grid are unchanged. Before another loan, the team will specify which target and support motions a reference must separate, and the verified range and uncertainty it must cover. Nominal one-micrometre divisions supply none of those guarantees.

Retained result table

Minutes after power-on	Action	KIS-61 change (um)	NQ-17 change (um)	Difference (um)
0	Start	0	0	0
10	No commanded motion	7	6	1
30	No commanded motion	9	8	1
60	No commanded motion	9	8	1
61	Move target by hand	13	12	1
62	Attempt target restoration by hand	11	10	1

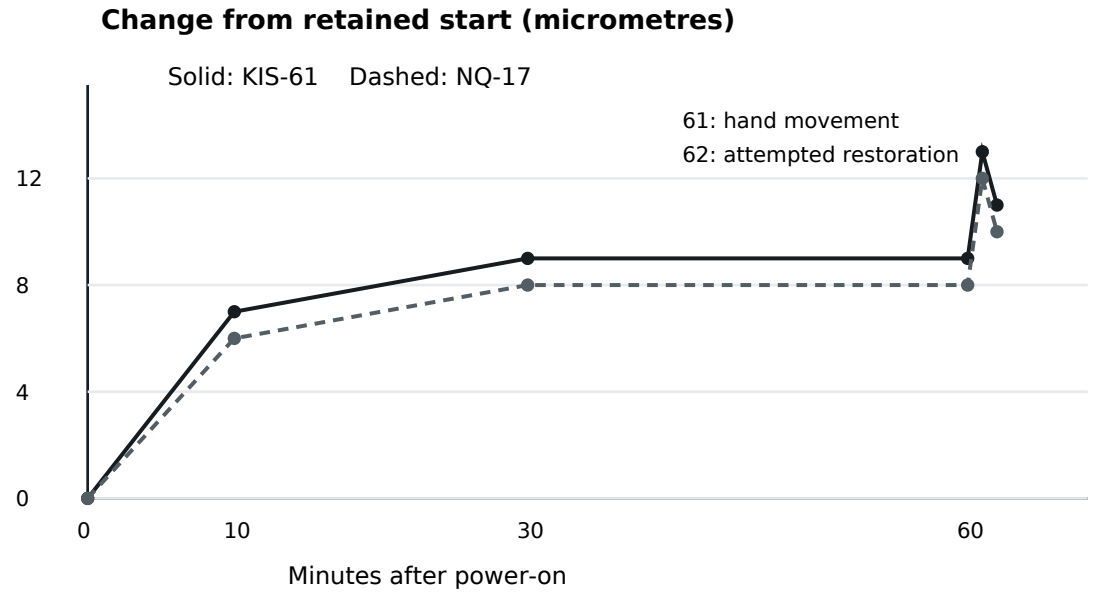


Figure 3. January 6 changes from the retained starting displays. Closely spaced actions at minutes 61 and 62 follow the stationary hour. No absolute calibration is implied.

Scope retained for review

performed unloaded local comparison; manual displacement is uncalibrated and instrument mounts stay fixed. No absolute reference, traceability, uncertainty budget, loaded material or bidirectional full-grid acceptance. The 25 um budget, 6 um allowance and 1000 um grid remain unchanged.

Source: doc-sei-013 / Revision 1 / https://meridianrc.org/research/cycle-013/space-elevator/

II. Lumen research record

Documentary section

The following notes retain their original research dates and technical scope. Presentation follows the current institutional style. Later synthesis does not add measurements to these records. Result tables are reproduced where supplied.

M-LUM-001 / Keeping the events a correction discards

2026-10-07 / Elias Chen, Leila Haddad

An optical-processing fixture demonstrates why a corrected stream cannot replace an independent raw archive.

Question and method

Can an adaptive summary reconstruct transients that it has discarded?

Create 120 indexed synthetic events; deliberately suppress transients at indices 19 and 79 in a corrected view. Use an untouched raw mirror as control.

Observed behavior

The corrected view retains 118 events and loses both injected transients. The raw mirror retains all 120 and both transients.

Interpretation and limits

The omission is programmed, not an unexpected discovery. This fixture does not model atmospheric propagation or a deployed lunar terminal.

The result establishes a failure mode in this deliberately simplified software fixture. It does not validate the proposed program architecture.

Next investigation

Test clock-offset metadata, packet loss and replay under independent archive failures.

Scope retained for review

The omission is programmed, not an unexpected discovery. This fixture does not model atmospheric propagation or a deployed lunar terminal.

Source: doc-lumen-001 / Revision 1 / https://meridianrc.org/research/cycle-001/lumen/

M-LUM-002 / The event neither archive can return

2026-10-12 / Elias Chen, Leila Haddad

A twelve-event replay fixture recovers disjoint archive losses but leaves a shared gap intact. Content recovery and trustworthy acquisition timing turn out to be different questions.

The intact mirror was the favorable case

The first Lumen replay note retained an untouched mirror of 120 events. That made it possible to find both transients discarded by the corrected view. Elias Chen and Leila Haddad now remove that favorable assumption: what happens when the archives themselves have gaps?

The follow-up generates twelve indexed source events. Archive A has a +7-millisecond acquisition-clock offset, archive B a -11-millisecond offset, and each receives events after a different arrival delay. Sequence identity, acquisition time and archive arrival time remain separate fields.

What redundancy recovers

With intact archives, the union retains all twelve sequence IDs. When A loses sequence 2 and B loses sequence 8, their union also retains all twelve: each archive supplies what the other lacks.

When both archives lose sequence 5, the union retains eleven events. Sequence 5 stays explicitly missing. The replay does not invent an intermediate event or replace the gap with a smooth interpolation; neither archive contains the payload needed to recover it.

A complete event list is not an aligned clock

Known, current calibration allows retained acquisition timestamps to be corrected to the authored source times. Missing calibration leaves all 22 retained archive records in the disjoint-loss case unaligned. The event IDs can still be recovered, but their corrected acquisition times cannot be justified.

In the expiry challenge, 11 retained records fall outside their respective calibration validity windows and remain unaligned. This is not the loss of eleven distinct events: the count refers to records across both archive views. Again, all twelve event IDs remain available.

Arrival delays belong to the transport history, not the acquisition clock. The check does not use similar arrival times as evidence of synchronization.

An operational question still open

Chen has a bounded recovery rule for complementary loss; Haddad has an explicit counterexample to the promise that raw redundancy always makes reconstruction possible. The fixture makes their reporting distinction practical: state which content survives, which timing is justified and which gaps remain.

The authored metadata is trustworthy here. Corrupted sequence IDs, payload conflicts, drift and inaccurate calibration inputs require further tests. No optical receiver or lunar terminal has been exercised.

Scope retained for review

IDs and payloads are trusted fixture inputs. Offsets, validity windows and losses are authored. No receiver hardware, physical clock drift, corrupted identifiers or lunar link is modelled.

Source: doc-lumen-002 / Revision 1 / https://meridianrc.org/research/cycle-002/lumen/

M-LUM-003 / Two copies of event five

2026-10-19 / Elias Chen, Leila Haddad

A replay of twelve model receiver samples distinguishes a recoverable damaged record, an absent sample and two intact records that disagree about the same event.

From incoming light to a replayable event

Lumen's proposed optical base would receive light, convert it into detector measurements and process those measurements into useful traffic and acquisition information. This study concerns preserving a receiver's observation history alongside adaptive processing. It does not test the transmitted message or the performance of a laser link.

Our software stream contains twelve labelled photodiode samples, numbered zero through eleven and spaced 100 milliseconds apart. The sample names are synthetic payloads, not measured photocurrent. Each archive copy carries the sequence number, acquisition time, payload and a checksum over those fields. The earlier archive-loss study assumed those identities and contents were trustworthy. Here we deliberately damage them.

An intact record can still disagree

Changing archive A's event-five identifier to six leaves its retained checksum inconsistent. The replay rejects that record and recovers event five from B; all twelve events remain usable. This is the case where redundancy delivers the recovery Chen wants.

A different case changes B's event-five payload and recomputes its checksum. Both copies are internally intact, but they disagree. Replay keeps both variants for inspection and excludes event five from the uncontested stream. Eleven events are usable. Nothing in a checksum tells us which version describes the source.

That differs from the shared-loss case, where both archives lack event five altogether. There are again eleven events, but now there is no competing version to inspect. The report separates missing from disputed. It also preserves two distinct event IDs when their payloads happen to be equal: a steady detector reading is not a duplicate event merely because its value repeats.

The clock walks away from its first comparison

In the timing example, a local clock reads 7 milliseconds at reference time zero and 1,017 milliseconds at reference time 1,000. A sample stamped 512 belongs at reference time 500 in the authored linear model. Subtracting only the starting offset places it at 505; using both synchronization witnesses places it at 500.

This correction assumes linear drift between those witnesses. A clock could wander and return between the endpoints. The next replay should test that possibility and avoid extending a correction beyond its observed interval.

A useful replay must show what it cannot settle

Elias Chen accepts the eleven-event uncontested stream as useful, provided the quarantined variants remain accessible. Leila Haddad asks for the conflict to appear beside the stream rather than in a success total's footnote. They agree on a report with three separate fields: recovered events, missing events and conflicting events.

These are six deterministic software cases, with retained checksums and authored clock witnesses. No detector, terminal or lunar path was operated. The next engineering question is how the receiver creates and protects its first record before either archive copies it.

Retained result table

Challenge	Usable events	What remains
Intact control	12 / 12	No gap or conflict
Changed sequence ID	12 / 12	One damaged record rejected
Resealed payload conflict	11 / 12	Two versions of event 5 retained
Both copies lose event 5	11 / 12	Event 5 missing
Equal payloads, different IDs	12 / 12	Both events retained
Linear clock drift	12 / 12	Sample corrected from 505 to 500 ms

Scope retained for review

Checksums cannot establish source truth or defeat deliberate resealing. Clock witnesses and linear drift are authored. No optical terminal, photodiode hardware or lunar link was tested.

Source: doc-lumen-003 / Revision 1 / https://meridianrc.org/research/cycle-003/lumen/

M-LUM-004 / Before the receiver makes its second copy

2026-10-27 / Elias Chen, Leila Haddad

A first-record witness catches an edit that two agreeing archives conceal. An interior timing anchor catches a clock excursion that endpoint synchronization misses.

The first sample is part of the instrument

Lumen proposes optical infrastructure that could keep surface users and Earth gateways connected. Its receiver would turn detector readings into a useful stream while retaining enough raw history to inspect a transient or a processing decision. The preceding replay study compared surviving copies. This study moves the challenge upstream, before the receiver makes its second copy.

We authored twelve photodiode sample records and a separately retained digest for each first acquisition. Two archives then receive copies. The digest is an assumed witness of what was recorded at acquisition, not proof that a detector observed the world correctly. Payloads are labels in software; there is no light source or photodiode hardware in this experiment.

Two identical archives can preserve the same edit

The injected pre-copy edit replaces event five before both archives receive it. Their agreement survives. A digest retained separately from the edited stream disagrees, leaving eleven recovered events and one disputed event. Both archived variants remain accessible even when their contents are identical; the discrepancy is with the first-record witness.

If the purported witness is recomputed from the edited source, all twelve pass. That is a deliberate blind spot. A checksum cannot rescue a witness that inherits the same changed input. A post-copy edit instead produces two conflicting archive variants. The shared-gap case has eleven recovered events and one missing event; the missing-witness case has eleven recovered events and one surviving but unwitnessed event. Each state needs its own display and later review path.

The clock can wander and return

Reed and Haddad used their methods exchange to ask where a clock correction gets its authority. With supplied endpoints at true/local 0/7 and 1,000/1,017 milliseconds, the linear control places the midpoint local reading at 512 milliseconds and reconstructs 500 milliseconds.

The second case adds a 20-millisecond excursion at the midpoint while preserving both endpoints. Its local midpoint is 532 milliseconds. Endpoint-only correction estimates 519.801980 milliseconds for the true 500-millisecond event. A separately supplied midpoint anchor exposes a 20-millisecond local residual, beyond the authored five-millisecond tolerance; endpoint timing is withheld.

One interior anchor does not measure the entire clock path. It exposes this supplied excursion, and another excursion between anchors could still escape. We have a reason to reject one timing fit, not a universal bound on timing error.

Continue service without promoting agreement to truth

Chen keeps the uncontested portion available. Haddad keeps disputed, missing and unwitnessed identifiers beside it. They separate the survival of the service stream from permission to use a sample for a timing claim. The next design task is a receiver-side witness retained before adaptive rewriting, with explicit failure behavior if that witness is unavailable.

First-record witnesses and timing anchors are authored independent inputs. A shared witness accepts the injected lie; sparse anchors cannot bound every between-anchor excursion. No lunar link or optical hardware tested.

Retained result table

Record case	Recovered	Missing IDs	Disputed IDs	Unwitnessed IDs
intact_control	12	[]	[]	[]
pre_copy_edit_independent_witness	11	[]	[5]	[]
pre_copy_edit_shared_witness	12	[]	[]	[]
post_copy_conflict	11	[]	[5]	[]
shared_gap	11	[5]	[]	[]
missing_first_witness	11	[]	[]	[5]

Scope retained for review

First-record witnesses and timing anchors are authored independent inputs. A shared witness accepts the injected lie; sparse anchors cannot bound every between-anchor excursion. No lunar link or optical hardware tested.

Source: doc-lumen-004 / Revision 1 / https://meridianrc.org/research/cycle-004/lumen/

M-LUM-005 / The silence between clock witnesses

2026-11-05 / Elias Chen, Leila Haddad

A receiver can preserve a sample while losing the right to place it on a precise timeline. Twelve cases separate first-record retention, archive service and timing permission.

Keep what the photodiode first supplied

Lumen's proposed lunar receiver turns optical arrivals into recorded samples and then adapts those samples for service. Its first-record study showed that two agreeing archives can share an earlier edit. Chen and Haddad now put the retained witness before adaptive rewriting in the logical interface. A sequence number, session identity, raw payload digest, capture-stage label and retention acknowledgment travel separately from the processed service record.

The prototype is twelve authored photodiode-sample labels, not optical measurements. The capture path forks before rewriting: one branch retains a digest under a declared separate custodian; the other branches into service processing and archives A and B. A missing acknowledgment leaves a surviving event unwitnessed. It does not manufacture a disagreement or delete the usable payload.

Agreement after the wrong fork

Rewriting sample five after first retention disputes that sample against its witness. Moving the witness after rewriting makes all twelve payloads appear recovered by digest comparison. Because the declared capture stage is wrong, acquisition permission is nevertheless withheld. A declared shared custodian is also withheld. If a shared source is concealed behind false independent labels, this interface cannot expose it.

An outage at witness five yields eleven recovered events and one unwitnessed survivor. A gap in both archives yields eleven recovered and one missing event. Two payload variants yield eleven recovered and one disputed event. These are separate service states. None, by itself, certifies the receiver clock.

A clock can leave and return

In one timing case, witnesses at 0, 500 and 1,000 milliseconds all report zero residual. A sample at 250 milliseconds has a hidden 20 millisecond excursion; a paired control has none. The reviewer receives identical anchors in both cases. It therefore marks both events unwitnessed between anchors. The injected excursion exists in simulator truth only, and is not smuggled into the decision as an extra observation.

A provisional bound becomes possible only when an additional assumption is supplied: maximum residual slew of 0.02 milliseconds per millisecond, with each anchor bounded to 1 millisecond. At the midpoint of a 100 millisecond gap, the

nearest anchor gives 1 + 0.02 × 50 = 2 milliseconds. At the midpoint of a 500 millisecond gap it gives 6 milliseconds, exceeding the chosen 5 millisecond requirement. This is a bound on the declared model, not measured oscillator behavior.

Operate with an explicit missing permission

An absent interior clock witness leaves the remaining payloads available but timing uncertified. Sharing clock custody with the record controller withholds the timing permission even in the short-gap case. First-record retention and clock anchoring therefore need distinct custody assignments and availability records.

Chen accepts a service display with surviving samples and a separate timing state. Haddad requires a future clock characterization to justify any slew limit before that display can claim physical timing. No lunar receiver or optical terminal was operated. Clock bounds require a declared maximum residual slew, not established receiver performance. Supplied custody labels are assumptions. The 20 ms hidden excursion is known to the simulator only; agreeing anchors cannot discover it. Integrity recovery does not confer timing permission. No optical or lunar link tested.

Retained result table

Synthetic case	Recovered / missing / disputed / unwitnessed	Timing bound (ms)	Permission
retained_before_rewrite	11 / 0 / 1 / 0	Not evaluated	evaluate_per_event
rewrite_before_witness	12 / 0 / 0 / 0	Not evaluated	withhold_dependency
witness_outage	11 / 0 / 0 / 1	Not evaluated	evaluate_per_event
archive_gap	11 / 1 / 0 / 0	Not evaluated	evaluate_per_event
archive_conflict	11 / 0 / 1 / 0	Not evaluated	evaluate_per_event
shared_custodian_blind_spot	12 / 0 / 0 / 0	Not evaluated	withhold_dependency
sparse_unbounded	Timing case	No bound	unwitnessed_between_anchors
hidden_between_anchor_excursion	Timing case	No bound	unwitnessed_between_anchors
declared_slew_short_gap	Timing case	2.0	conditional_timing_within_declared_bound
declared_slew_wide_gap	Timing case	6.0	withhold_timing_bound
clock_witness_outage	Timing case	No bound	unwitnessed_between_anchors
clock_custody_shared	Timing case	2.0	withhold_clock_dependency

Scope retained for review

Clock bounds require a declared maximum residual slew, not established receiver performance. Supplied custody labels are assumptions. The 20 ms hidden excursion is known to the simulator only; agreeing anchors cannot discover it. Integrity recovery does not confer timing permission. No optical or lunar link tested.

Source: doc-lumen-005 / Revision 1 / https://meridianrc.org/research/cycle-005/lumen/

M-LUM-006 / The timestamp that the receiver must leave open

2026-11-13 / Elias Chen, Leila Haddad

Eight receiver cases keep payload and timing permissions separate while testing witness gaps and an assumption that can fail without leaving an observable warning.

What remains on the display

The intended Lumen receiver retains photodiode samples before adaptive processing and displays the status of the observation beside the sample. Chen and Haddad replay one authored sample, its pre-rewrite digest and a separate timing-anchor stream. The clock-witness note established that agreement at sparse anchors does not establish the path between them.

This study makes the display rule explicit. A retained sample can remain available for replay while its timestamp is marked unqualified. A disputed payload gets a different state. Withholding timing does not erase the stored sample, and recovering a sample does not complete its clock record.

Cadence under a declared slew limit

For a sample at 50 milliseconds between zero-residual anchors at 0 and 100, a 1 millisecond anchor allowance plus an assumed 0.02 millisecond-per-millisecond residual-slew limit gives a 2 millisecond conditional bound. A sample at 250 between anchors at 0 and 500 gives 6 milliseconds. The latter exceeds the 5 millisecond design allowance and is held.

Losing the following anchor removes bracketing altogether; the procedure does not extend the last good timestamp into the missing interval. A declaration that the witness shares the receiver dependency also holds timing. These rules are about supplied model records, not demonstrated optical-link performance.

A favorable row with a violated assumption

Two sparse cases without a slew limit have identical reviewer input: one has no excursion, and the other has an injected 20 millisecond excursion between anchors. Both remain unbounded. The full fixture retains simulator truth separately so a reviewer is not credited with detecting information it never receives.

More importantly, the favorable dense case has an identical-input partner with a 4 millisecond event residual. That partner violates the assumed slew limit. The reviewer still returns the conditional 2 millisecond bound because the available anchors have not changed. The procedure does not detect this failure; the bound's usefulness depends on a future physical characterization that covers the actual clock.

Three visible permissions

A rewrite after the independent sample digest is retained produces a payload dispute and withholds permission for that event. The proposed interface will display retained payload, acquisition dependency and timing qualification separately. Its timing label must carry the assumed slew dependency rather than presenting 2 milliseconds as measured uncertainty.

Physical clock selection, residual characterization and outage measurement remain next tasks. No optical terminal, lunar link or physical clock was operated in this study.

Retained result table

Case	Payload	Conditional timing bound (ms)	Timing decision
dense_assumed_slew	retained	2.0	conditional_within_assumed_slew
long_gap	retained	6.0	withhold_over_5_ms
sparse_control_no_slew	retained	No bound	withhold_unbounded_timing
sparse_hidden_no_slew	retained	No bound	withhold_unbounded_timing
lost_following_anchor	retained	No bound	withhold_unbracketed_event
rewrite_after_capture	disputed	2.0	withhold_disputed_payload
declared_shared_witness	retained	No bound	withhold_witness_dependency
assumed_slew_violated	retained	2.0	conditional_within_assumed_slew

Scope retained for review

Receiver replay only; all anchor residuals are authored as zero with a declared 1 ms anchor bound. The 0.02 ms/ms residual-slew limit is unmeasured. A violated assumption can invalidate a favorable timing bound without changing the available records. No optical terminal or physical clock is tested.

Source: doc-lumen-006 / Revision 1 / https://meridianrc.org/research/cycle-006/lumen/

M-LUM-007 / The receiver starts counting again

2026-11-21 / Elias Chen, Leila Haddad

A retained sample can survive a receiver reboot while its timing permission fails at the new counter epoch.

One sample, two boots

A future Lumen terminal must recover useful data after interruption. Its local tick counter may restart before a fresh timing witness is available. The earlier receiver study separated payload recovery from timing qualification; Chen and Haddad now test whether replaying a saved favorable label can erase that distinction.

The control stores a bound string sample at tick 50 of boot B, with same-boot anchors at ticks 0 and 100. The authored 1 millisecond anchor allowance and 0.02 millisecond per millisecond residual-slew bound give 1 + 0.02 × 50 = 2 milliseconds. This is the same conditional timing calculation used in the preceding study, with an explicit boot identity added to every anchor.

The first anchor is not a bracket

After boot C begins, its tick 50 lies numerically between the old ticks 0 and 100. Those values belong to boot B and provide no bracket for the new sample. One C anchor at tick 0 still leaves the following side unobserved. Both cases retain the payload but hold timing.

Adding C anchors on both sides permits the conditional 2 millisecond result again. A cached favorable label with no retained anchors does not. Altering the sample string breaks its digest binding and disputes the payload separately. The seven cases therefore exercise a restart as a change in available evidence, rather than a command to recover the previous screen state.

What the epoch label cannot prove

A hidden case relabels old anchor records as if they came from the current boot. Its available input is identical to the honest control, and the reviewer accepts the same conditional bound. The model cannot authenticate its own supplied boot labels or discover an unrecorded reset.

The next receiver implementation needs durable sample bindings, a defensible boot identity and recoverable witness records. A physical clock still has to be selected and characterized. Neither the 2 millisecond result nor this restart harness establishes optical hardware performance, lunar communication or a measured clock envelope.

Retained result table

Case	Payload	Timing	Conditional bound (ms)
retained_boot_record	retained	conditional_model_timing	2.0
counter_restart_old_anchors	retained	hold_witness	Unavailable
one_new_anchor	retained	hold_witness	Unavailable
new_boot_bracket	retained	conditional_model_timing	2.0
payload_changed	disputed	hold_payload	Unavailable
cached_permission_no_witness	retained	hold_witness	Unavailable
hidden_epoch_relabel	retained	conditional_model_timing	2.0

Scope retained for review

Boot IDs, payload binding and anchors are authored software records. The 1 ms anchor and 0.02 ms/ms slew assumptions are unmeasured. No proof of authentic boot identity, storage durability, physical clock behavior, optical terminal or lunar link.

Source: doc-lumen-007 / Revision 1 / https://meridianrc.org/research/cycle-007/lumen/

M-LUM-008 / The last witness record

2026-11-28 / Elias Chen, Leila Haddad

A numbered witness journal exposes some missing records without making its own session label trustworthy.

Recovering a sample at tick fifty

Lumen needs a receiver that can retain a sample through interruption and reconstruct the evidence supporting its time. The restart study held timing when old-boot anchors were reused. Chen and Haddad now give the witness records their own session and sequence numbers, then remove or replace entries during recovery.

The control sample belongs to session S2 at tick 50. Witness records numbered 1 and 2 belong to S2 at ticks 0 and 100. Their agreement supplies a conditional model bracket. This cycle tests journal structure, not the earlier numerical clock-slew bound; no new millisecond performance is inferred.

A complete prefix is still too short

Deleting record 2 leaves a perfectly contiguous prefix beginning at 1. It does not leave an anchor after tick 50. The payload survives while timing remains held. A gap inside a longer list and a reused sequence number produce different diagnostic holds. A sample from S1 cannot borrow the S2 journal.

The six cases retain a favorable-looking failure: if the supplied sample and entire journal have all been relabelled S2, the available input is exactly the control. The structure checker cannot discover the false origin. Numbering the journal detects some damage to a record; it does not establish where that record came from.

Recovery has two remaining jobs

The receiver design must preserve a durable journal and establish a session identity independently of the replayed labels. It must also recover a following anchor before qualifying a sample that lacks one. Neither job is completed by this software study.

Physical clock selection and characterization remain open. There is no installed optical terminal or lunar link, and no inference about the unresolved seasonal fragment follows from these deliberately authored losses.

Retained result table

Case	Recovery decision
complete_journal	conditional_model_bracket
missing_tail	retain_payload_hold_timing
middle_gap	hold_gap
duplicate_sequence	hold_duplicate_sequence
other_session	hold_session
hidden_shared_relabel	conditional_model_bracket

Scope retained for review

Authored journals and session labels only; no authenticated session generator, measured oscillator, durable hardware store, optical terminal or lunar link. A contiguous local sequence cannot prove its own origin or reveal an entirely missing later session.

Source: doc-lumen-008 / Revision 1 / https://meridianrc.org/research/cycle-008/lumen/

M-LUM-009 / A challenge outside the receiver

2026-12-07 / Leila Haddad, Elias Chen

A one-use session proposal rejects some replays while retaining the cloned-issuer failure.

Who issued S3

Lumen's receiver must preserve samples through interruption while keeping their timing evidence attached to the correct session. The witness journal study showed that replacing every session label could leave a false-origin record looking complete. Haddad's next proposal places a one-use challenge outside the recovered receiver journal. Chen tests the proposal with stale and missing state.

In the favorable model case, an issuer ledger records challenge C41 as issued and unused for session S3. The sample at tick 50 and witness anchors at ticks 0 and 100 all carry C41. The review checks the supplied issuer state before reusing the preceding journal checker. Its favorable result remains a conditional model bracket.

A used challenge and an absent anchor

A consumed C41 holds recovery even when all labels match. An unavailable issuer also holds, and a sample carrying C40 cannot borrow C41's journal. These are six authored cases, not a running challenge service. The model accepts issuer validation as an input; it performs no cryptographic authentication.

Removing the tick-100 witness still retains the payload while holding timing. Fresh session evidence cannot supply the missing side of a time bracket. The session proposal and the earlier clock-characterization work therefore remain separate obligations.

The ledger can be replayed too

The retained cloned-issuer case gives the reviewer exactly the same sample, journal and issuer input as the favorable control. Its authored history says that all three were copied together. The reviewer returns the same favorable result because it has no independent observation with which to distinguish them.

The next design must specify who maintains challenge state and how consumption survives a restart outside the receiver's failure domain. No deployed issuer, physical clock performance, optical terminal or lunar link is established here. The unresolved seasonal archive has not entered this study.

Retained result table

Case	Recovery decision
fresh_model_session	conditional_model_bracket
consumed_challenge	hold_consumed_challenge
issuer_unavailable	hold_issuer
sample_relabel	hold_challenge_mismatch
missing_tail	retain_payload_hold_timing
cloned_issuer_state	conditional_model_bracket

Scope retained for review

Software-only challenge and issuer inputs; no cryptographic verification, deployed issuer, durable storage or measured oscillator. The earlier false-origin case is narrowed to issuer-state custody, not solved. No optical or lunar link and no seasonal-fragment intake.

Source: doc-lumen-009 / Revision 1 / https://meridianrc.org/research/cycle-009/lumen/

M-LUM-010 / The challenge that stays used

2026-12-16 / Leila Haddad, Elias Chen

A serialized consumption model rejects duplicate recovery while exposing the failure of a split check and an upstream rollback.

Two callers arrive together

Lumen preserves receiver payload and its time anchors across interruption. The session proposal left an unused challenge outside the receiver but did not specify the transition that consumes it. Haddad now models one atomic check-and-consume operation in that external issuer state. Chen gives two callers the same C42.

RFC 8613 supplies a relevant engineering constraint: replay protection needs atomic state updates and explicit recovery after volatile context is lost. This study adapts that constraint to the challenge model; it is not an implementation of OSCORE or its cryptography.

An interval with no successful caller

Serialized model calls produce one conditional bracket and one consumed-state hold. In the deliberately broken split check, both callers inspect unused state before either changes it, and both appear successful. A final consumed marker cannot tell a reviewer that two admissions already occurred.

A cut before commit produces no admission. A cut after the consumed marker persists, but before a successful reply, leaves a burned challenge pending reconciliation. Restart must not promote that state to unused. This sacrifices a convenient retry to preserve the one-use rule within the model.

What the marker cannot establish

The first-use and hidden-rollback controls supply identical unused issuer input. Both obtain the favorable result even though one authored history substitutes an old snapshot. Atomic access protects competing calls to the same current state; it does not authenticate the state that was supplied.

The next implementation needs protected issuer custody and an observable recovery boundary. Payload recovery and missing-anchor timing holds remain separate. No cryptographic service, clock performance, optical terminal or lunar link is demonstrated.

Retained result table

Case	Decision
first_use	conditional_model_bracket
already_consumed	hold_consumed_or_unknown
unknown_challenge	hold_consumed_or_unknown
issuer_unavailable	hold_issuer_unavailable
cut_before_commit	hold_no_commit
cut_after_commit	hold_consumed_pending_reconciliation
hidden_issuer_rollback	conditional_model_bracket
atomic_competing_calls	conditional_model_bracket / hold_consumed_or_unknown
broken_split_check	conditional_model_bracket / conditional_model_bracket

Scope retained for review

An authored serialized transaction model, not an OSCORE implementation, deployed cryptographic issuer or physical durability test. Missing anchors still hold timing; cloned authority, clock characterization and independent issuer custody remain unresolved.

Source: doc-lumen-010 / Revision 1 / https://meridianrc.org/research/cycle-010/lumen/

M-LUM-011 / Recovering the reply, not the challenge

2026-12-24 / Leila Haddad, Elias Chen

Bound request and payload records permit a model recovery after a lost reply, while a shared rollback remains invisible.

A consumed challenge has an unfinished conversation

Lumen's receiver needs to preserve payload and time anchors through interruption. The consumption study leaves a useful difficult case: C42 has been consumed, but its reply never reaches the receiver. Trying C42 as fresh would discard the one-use decision that survived.

Haddad specifies a recovery lookup for the existing decision. In this model, the issuer and a separately supplied witness each contain epoch 7, challenge C42, consumed state, request Q9 and payload block-A. Chen compares the complete bindings before returning the existing decision. Request and payload names are literal fixture values; no digest, authentication service or storage transaction is implemented here.

The recovery names its object

The lost-reply and repeat-recovery cases retrieve the same supplied decision with zero new consumptions. Q10 cannot borrow Q9's result, and block-B cannot inherit block-A's binding. A missing witness, different epoch, changed witness record or withdrawn authority each produces a hold. Repetition is a read of the retained decision, not another admission.

This comparison assumes the original consumed decision was committed correctly. It does not create a durable journal or repair a partial commit. A committed consumption without a complete decision/witness record remains a reconciliation task, rather than permission to invent the missing record.

Two copies can share an old present

The hidden-rollback control substitutes both epoch-7 records for a later epoch-8 history. The supplied inputs are identical to the favorable case, and the lookup returns the same decision. Agreement between copies cannot reveal a later history neither copy retains.

An epoch number is useful only with a justified source for its current value. The next design must place a freshness witness outside the issuer's recovery snapshot and specify its failure behavior. Missing time anchors still hold timing permission even when payload recovery is possible. No protected storage, cryptographic implementation, characterized clock or lunar optical link is demonstrated.

Retained result table

Authored recovery case	Decision	New consumptions
lost_reply	recover_existing_decision_no_new_consumption	0
repeat_recovery	recover_existing_decision_no_new_consumption	0
different_request	hold_binding_mismatch	0
different_payload	hold_binding_mismatch	0
missing_witness	hold_missing_state	0
old_issuer_epoch	hold_epoch_disagreement	0
different_witness_record	hold_record_disagreement	0
withdrawn	hold_authority	0
both_domains_rolled_back	recover_existing_decision_no_new_consumption	0

Scope retained for review

Authored state comparison only. Durable records, separate loss domains and atomic decision creation are assumed inputs; this code does not implement them. No authentication, protected monotonic epoch, clock characterization or operational lunar link. Matching jointly rolled-back records cannot establish freshness.

Source: doc-lumen-011 / Revision 1 / https://meridianrc.org/research/cycle-011/lumen/

M-LUM-012 / Who keeps the current epoch?

2027-01-01 / Leila Haddad, Elias Chen

A separate model authority exposes issuer rollback while a rollback of all authorities remains indistinguishable from a valid earlier history.

Recovery needs a present to compare with

Lumen's receiver can recover an already bound request after losing its reply. The recovery lookup also showed why two copied records cannot establish that they are current. Haddad and Chen now supply the current epoch through a third model input outside the issuer and witness snapshots.

The added input contains the realm lumen-recovery and an integer current_epoch. It stands for a separately held authority, not an implemented server, signed attestation or protected counter. The earlier lookup still checks consumed state, request Q9, payload block-A, challenge and authority before the epoch comparison occurs.

The same copies can now fail

Matching epoch-8 records and an epoch-8 authority recover the existing decision. Matching issuer and witness records rolled back to epoch 7 hold against that same current authority. Epoch-8 records also hold when the authority itself reports epoch 7: disagreement does not give the lookup permission to choose which account is right.

An absent authority, wrong realm and unknown current epoch each hold. Withdrawn authority and a changed request still fail the earlier checks. The model never consumes another challenge and does not repair missing time anchors.

A third copy is not independence

If the authority rolls back together with issuer and witness, all three supplied epochs can agree at 7. That case has exactly the same reviewer inputs as a valid epoch-7 history and receives the same favorable result. The extra field does not manufacture knowledge of epoch 8.

The practical next step is a custody and failure design for the current-epoch authority, including how its response is bound to a live recovery request. Authentication, replay resistance, availability and durable advancement remain unimplemented. These ten histories test a conditional rule; they do not establish an operational freshness service.

Retained result table

Supplied-state case	Decision
current	recover_bound_current_model_decision
issuer_witness_rollback	hold_epoch_not_current
authority_behind	hold_epoch_not_current
missing_authority	hold_epoch_authority_unavailable
wrong_realm	hold_wrong_authority_realm
unknown_epoch	hold_unknown_current_epoch
withdrawn	hold_authority
other_request	hold_binding_mismatch
all_domains_rollback	recover_bound_current_model_decision
valid_epoch_7_history	recover_bound_current_model_decision

Scope retained for review

Software over supplied records only. The external authority is assumed independent, authentic and current within favorable cases; no protected counter, authenticated exchange, durable write, lunar link or timing recovery exists. Joint rollback of every authority remains invisible.

Source: doc-lumen-012 / Revision 1 / https://meridianrc.org/research/cycle-012/lumen/

M-LUM-013 / A reply to this question

2027-01-08 / Leila Haddad, Elias Chen

Authenticated toy responses reject altered or differently bound replies; verifier reset and joint rollback still defeat the model's freshness assumptions.

Give the authority an actual reply

Lumen needs a receiver to recover an existing decision after losing its reply, without consuming another challenge or trusting an old snapshot. The current-epoch study supplied an authority record but did not authenticate a response. Haddad and Chen now implement a local toy exchange: an authority reply carries a realm, session nonce, request, payload and epoch, together with an HMAC-SHA256 tag over their canonical JSON representation.

The program checks the tag and the active realm/nonce/request/payload before applying the existing bound-recovery and epoch rules. This is executed software, but its key and nonce strings are fixed public fixture data. There is no network authority, unpredictable challenge generator, protected key store or durable service. The exercise tests the implemented decision path, not deployment security.

Authentic does not mean applicable

A correctly authenticated reply for the active values recovers the existing epoch-8 decision and closes that local verification session. A valid tag on a different nonce, request, payload or realm still holds. Changing the epoch after signing fails authentication. A response signed with another test key also fails, and absence of a response holds.

Matching issuer and witness snapshots at epoch 7 hold against the authenticated epoch-8 response. An epoch-7 reply against epoch-8 records also holds; the verifier does not decide which account should win. Withdrawn authority remains withdrawn. A second use of a closed session holds, and none of these recovery reads consumes a new challenge.

The verifier also needs a history

Resetting the verifier and reusing its old nonce reopens the same active values. The recorded authentic response passes again. A nonce label does not prove uniqueness after state loss. This explicit failure remains in the retained output rather than being treated as a fresh successful exchange.

Rolling the records and responding authority together to epoch 7 still produces exactly the same verifier inputs as a valid epoch-7 history. A key holder can authenticate false or stale content. The next task is therefore durable nonce non-reuse and actual epoch/key custody outside common recovery snapshots, with defined behavior when those records cannot be recovered. No lunar timing, live response latency or protected monotonic advancement has been measured.

Retained result table

Fixture history	Decision
current	recover_bound_current_model_decision
old_nonce	hold_response_binding
wrong_request	hold_response_binding
wrong_payload	hold_response_binding
wrong_realm	hold_response_binding
tampered_epoch	hold_authentication
wrong_key	hold_authentication
missing_response	hold_no_response
stale_records	hold_epoch_not_current
authority_behind	hold_epoch_not_current
all_domains_rollback	recover_bound_current_model_decision
valid_epoch_7_history	recover_bound_current_model_decision
withdrawn_authority	hold_authority
closed_session	hold_session_closed
reset_reuses_nonce	recover_bound_current_model_decision

Scope retained for review

Deterministic local software with public fixture keys and fixed nonce strings. No network endpoint, live unpredictable challenge generator, protected key store, persistent nonce ledger, protected epoch counter, durable service or recovered lunar time. A key holder can sign false content; all-domain rollback and reset reuse remain failures.

Source: doc-lumen-013 / Revision 1 / https://meridianrc.org/research/cycle-013/lumen/

III. Aletheia research record

Documentary section

The following notes retain their original research dates and technical scope. Presentation follows the current institutional style. Later synthesis does not add measurements to these records. Result tables are reproduced where supplied.

M-ALE-001 / Two channels, one source

2026-10-07 / Nadia Okafor, Tomás Ibarra

A source-discrimination fixture rejects apparent corroboration when two nominal channels share an upstream origin.

Question and method

Can feedback be misclassified as independent corroboration?

Route one tagged synthetic event into an input channel and a feedback channel. Compare channel counting with unique-origin counting; a one-channel view is the control.

Observed behavior

Channel counting reports two channels. Provenance-aware counting reports one independent origin and blocks promotion based on channel count alone.

Interpretation and limits

Source tags are correct by construction. No person, brain recording, implant or external information channel was tested.

The result establishes a failure mode in this deliberately simplified software fixture. It does not validate the proposed program architecture.

Next investigation

Challenge the origin check with missing tags, stale identifiers and shared upstream dependencies.

Scope retained for review

Source tags are correct by construction. No person, brain recording, implant or external information channel was tested.

Source: doc-aletheia-001 / Revision 1 / https://meridianrc.org/research/cycle-001/aletheia/

M-ALE-002 / The cost of withholding corroboration

2026-10-12 / Nadia Okafor, Tomás Ibarra

A provenance rule catches distinct labels that share an upstream origin, but also withholds three independent pairs when their tags or context cannot support the claim.

After counting the origin

The first source-origin check gave two channel labels one shared input. Nadia Okafor and Tomás Ibarra now challenge the provenance information itself: an origin label can be missing, reused or detached from the context in which it was acquired.

The follow-up compares seven authored pairs. Every pair has two nominal channels. The revised rule checks acquisition context, tag context and upstream-root metadata before treating the pair as independently sourced. Missing information produces an unknown classification, rather than an extra independent origin.

Labels are not enough

The independent control has two upstream roots and is accepted within the fixture. The shared-origin control has one root and is rejected as corroboration. A second shared-origin case uses two different labels, yet still has one upstream root; label counting would overstate its independence.

Conversely, the same identifier can appear in two different acquisition sessions with distinct upstream roots. A global unique-label count would collapse that pair into one origin. The context-aware rule retains two in this case, showing why identifiers need a scope.

Three independent pairs are withheld

The missing-tag, stale-context and unknown-upstream cases each have two independent origins in the authored ground truth. The rule cannot establish that independence from the available provenance and withholds all three. This is an explicit cost of the conservative decision, not a claim that those inputs were dependent.

Okafor's distinction between preserving a report and promoting it to corroboration remains important here. An unknown result does not erase the inputs. Ibarra's gate limits a claim while leaving the records available for a later provenance review.

The rule has not authenticated its evidence

Upstream roots and context metadata are trustworthy test inputs when present. The check therefore does not show that the system can detect a forged root or an apparently current but false context. Passing these cases is not an authentication result.

The next question is what evidence could verify an upstream dependency claim, and how a withheld pair can be reconsidered without rewriting its original record. No person, implant or external information channel was tested.

Scope retained for review

Context and upstream roots are trusted metadata when present. Ground-truth independence is authored, not experimentally measured. Forged provenance, human experiences, implants and exogenous information are not tested.

Source: doc-aletheia-002 / Revision 1 / https://meridianrc.org/research/cycle-002/aletheia/

M-ALE-003 / The circle that came back twice

2026-10-19 / Nadia Okafor, Tomás Ibarra

A synthetic symbol path shows how feedback can resemble corroboration, why a signed origin label can still mislead, and how a recovered receipt changes a decision without erasing its past.

What the two channels contain

Aletheia's proposed amplification loop would make weak cognitive and perceptual signals more legible while checking identity continuity and source attribution. Before a human or implant study, a software harness can ask whether its external-evidence gate confuses repetition with corroboration.

Our harness starts with a circle symbol. In one control, two separate generators each supply a circle. In another, the first circle is copied into a feedback path and returned as a second channel. The same visible agreement has different histories. No person sees these symbols in the experiment, and no neural sensing or stimulation occurs. The previous provenance study supplied the missing question: who can vouch for the origin labels?

A receipt binds the label to an issuer

Each model acquisition receipt binds an issuing authority, source root, session, symbol and nonce. A message-authentication check detects an edited source root; a context check rejects a receipt from an earlier session. Missing receipts remain withheld. The demonstration keys are public fixture inputs, not a secure deployment.

Two valid receipts naming different roots are still insufficient if the acquisition ledger shows both came from the same draw. In the signed-false-lineage challenge, both signatures verify and the roots differ, yet a separate ledger traces both to draw A. The rule records shared acquisition. Authenticating an assertion establishes who issued it; it does not make the assertion true.

Withholding gets a dated return path

On 18 October, the model decision for one independent pair is withheld because its second receipt is absent. On 19 October, the original receipt is recovered and checked against the current session and the two-draw ledger. A new decision accepts the pair within the fixture. The old withholding, its reason and its date remain in the decision history.

Nadia Okafor requested that this route be part of the method rather than a discretionary favour. Tomás Ibarra agreed, provided the new evidence is a recovered acquisition record rather than a newly invented label. The model makes the distinction explicit; it does not establish how a future acquisition service would preserve or recover that evidence.

What the circle does not establish

Seven cases exercise message binding, session checks, shared acquisition and reconsideration. A matching circle is not evidence of psychic reception. Even the independent control uses authored generators and an authored ledger. No experience, implant or external intelligence was tested.

The next challenge is a compromised or mistaken acquisition authority together with an incomplete ledger. The receipt service is now a named part of the engineering problem. Aletheia must eventually protect the independence of the observer as carefully as it protects the amplification loop.

Retained result table

Case	Receipt check	Decision
Separate draws	Valid, separate ledger draws	Independent within fixture
Feedback echo	Valid, one draw	Shared acquisition
Edited root	Authentication fails	Withheld
Old session	Context fails	Withheld
Missing receipt	Unavailable	Withheld
Signed false lineage	Valid, ledger shows one draw	Shared acquisition
Recovered original receipt	Valid, two draws	New acceptance; old withholding retained

Scope retained for review

Demonstration keys are public; receipt authorities and the draw ledger are authored. No human experience, neural data, stimulation or implant was tested; signatures do not prove an issuer truthful.

Source: doc-aletheia-003 / Revision 1 / https://meridianrc.org/research/cycle-003/aletheia/

M-ALE-004 / A valid signature and an unfinished ledger

2026-10-27 / Nadia Okafor, Tomás Ibarra

Seven acquisition challenges keep issuer authentication separate from draw independence. A late ledger entry can support a claim or undermine it, without erasing the first withholding.

What the amplifier must distinguish

Aletheia proposes to make weak cognitive and perceptual signals more legible while preserving autonomous judgment. Before any sensing or stimulation hardware, its evidence system must distinguish two independent inputs from one input returning through feedback. Our object remains two authored circle-symbol receipts. No person reports a circle in this study.

The preceding receipt study bound source assertions to fixture authorities. This cycle gives a valid assertion to an authority that is compromised in one scenario and mistaken in another. Both scenarios deliberately present the same receipts and shared-draw ledger. The labels describe injected causes; the software cannot diagnose the issuer's intent from identical evidence.

A correct receipt can tell the wrong source story

All seven pairs authenticate under public demonstration keys. In the independent control, generator-a maps to draw-a and generator-b to draw-b. In the compromised and mistaken cases, both map to draw-a. Two signed roots still yield one acquisition. The independent ledger exposes the shared draw; the signature alone does not.

A withdrawn-authority case authenticates the original receipt bytes but withholds use under the supplied authority-status rule. Cryptographic binding and current permission are separate questions. The case supplies that status rather than demonstrating a real revocation service.

An absent entry is not a second draw

When generator-b has no ledger entry, the result is withhold_incomplete_ledger. It is not independent and it is not proven shared. Haddad's October 23 exchange with Okafor supplied the missing-versus-disputed distinction; Ibarra carried it into a lookup that retains the absent root instead of inferring a favorable result.

Two reconsideration cases begin on October 26 with the same incomplete ledger. On October 27 a separately retained entry arrives. In one it maps generator-b to draw-b, permitting independence within the fixture. In the other it maps generator-b to draw-a, establishing shared acquisition. Each appends a new dated decision, its ledger and its reason. Neither changes the original signed receipt or removes the earlier withholding.

A second look must be allowed to disappoint

Okafor had asked for a route back when missing documentation becomes available. This cycle makes that route symmetrical: a second look can support the claim or reveal its dependency. Ibarra accepts that authentication answers who issued an assertion within the fixture, not whether the issuer was accurate.

The next task is to specify draw-ledger custody and authority-status timing. A colluding ledger and issuer can defeat the present model. The demonstration does not justify a claim about psychic reception, clinical benefit or any person's experience.

Public fixture keys and authored ledgers. Identical observations cannot distinguish malice from issuer error. Independent ledger custody is assumed; collusion can defeat this model. No participant, neural signal, stimulation or implant tested.

Retained result table

Software case	Receipt authenticates	Missing ledger roots	Decision
independent_control	True	None	independent_within_fixture
compromised_authority_shared_draw	True	None	shared_acquisition
mistaken_authority_shared_draw	True	None	shared_acquisition
incomplete_draw_ledger	True	generator-b	withhold_incomplete_ledger
withdrawn_authority	True	None	withhold_authority_withdrawn
late_independent_ledger	True	None	independent_within_fixture
late_shared_ledger	True	None	shared_acquisition

Scope retained for review

Public fixture keys and authored ledgers. Identical observations cannot distinguish malice from issuer error. Independent ledger custody is assumed; collusion can defeat this model. No participant, neural signal, stimulation or implant tested.

Source: doc-aletheia-004 / Revision 1 / https://meridianrc.org/research/cycle-004/aletheia/

M-ALE-005 / When the withdrawal reaches the ledger

2026-11-05 / Nadia Okafor, Tomás Ibarra

Nine synthetic custody challenges separate the time authority changes from the time a reviewer hears about it. Two agreeing signed chains still cannot prove independent draws.

Two circles are not yet two sources

Aletheia seeks to test a proposed psychic amplifier without mistaking two copies of one input for corroboration. At this stage, the source harness contains two synthetic circle receipts, generator roots and draw identifiers. There are no participants or neural signals. The prior authority study established that an authentic issuer statement can still describe a shared draw.

The new fixture gives the issuer a receipt binding and gives each of two declared ledger custodians a separate hash chain. Every chain entry carries a root, canonical acquisition identifier, draw time and previous-entry digest. Public HMAC keys demonstrate these bindings; they provide no production security. Roots are different labels for sources; independence depends on acquisition identifiers in the ledgers, not on different root names.

Effective time and received time

Every authority notice has an effective time, a receipt time and a revision. The synthetic draw occurs at minute 600. At review minute 630, a declared active authority and complete distinct-draw ledgers produce provisional independence. A withdrawal effective at minute 570 arrives at minute 660. The review at minute 690 appends withholding because authority was withdrawn before the draw, preserving the minute-630 decision as the judgment made with the information then available.

A separate withdrawal effective at minute 660 arrives at minute 690. The review at minute 720 retains provisional independence for the earlier minute-600 draw. Current withdrawal does not, by itself, rewrite earlier authority. Missing status coverage is withheld rather than filled with an assumption of perpetual activity. These minute values belong to synthetic sessions executed November 5, not real notices received on the collective calendar.

Two signatures can tell the same lie

When one custodian reports distinct draws and the other reports a shared draw, the fixture retains a disputed ledger. When both custodians collude, they sign internally consistent distinct-draw chains for a simulator-known shared draw. The reviewer accepts provisional independence: the chains agree and the declared custody assumptions conceal the falsehood. Adding signatures has not added an honest observation.

This blind spot is retained as a result. The classifier does not label every discrepancy an attack, and it does not claim to identify collusion from the agreeing chains. The simulator supplies motive for this case; a reviewer facing only those records could not establish it. Broken chain linkage is withheld, but a perfectly linked chain is not proof of truthful acquisition.

A judgment must be able to move either way

Two histories start with an incomplete ledger and withholding. Late distinct entries strengthen one decision to provisional independence; late shared entries weaken the other to shared acquisition. All prior decisions remain alongside their evidence snapshots. The delayed withdrawal supplies a further direction of change: an apparently eligible draw becomes ineligible when an earlier effective notice arrives.

Okafor requires reconsideration to remain symmetrical. Ibarra adds authenticated notice delivery and a genuinely independent acquisition checkpoint to the next design task. Custodians, delivery and status coverage remain declared assumptions; neither a real authority service nor a draw apparatus has been commissioned. Public HMAC keys and declared independent custodians demonstrate bindings only. Matching signed chains can jointly lie. Status coverage is an explicit assumption; effective and receipt times need authenticated external custody before real use. Disagreement identifies a discrepancy, not motive. No human, neural signal, implant or stimulation tested.

Retained result table

Synthetic case	Retained decision history (review minute → decision)
independent_control	630 → independent_within_declared_custody
late_withdrawal_before_draw	630 → independent_within_declared_custody; 690 → withhold_authority_at_draw
withdrawal_after_draw	630 → independent_within_declared_custody; 720 → independent_within_declared_custody
missing_status_coverage	630 → withhold_status_coverage
one_honest_custodian	630 → disputed_ledger
colluding_custodians	630 → independent_within_declared_custody
incomplete_then_distinct	630 → withhold_incomplete_ledger; 690 → independent_within_declared_custody
incomplete_then_shared	630 → withhold_incomplete_ledger; 690 → shared_acquisition
broken_ledger_chain	630 → withhold_ledger_chain

Scope retained for review

Public HMAC keys and declared independent custodians demonstrate bindings only. Matching signed chains can jointly lie. Status coverage is an explicit assumption; effective and receipt times need authenticated external custody before real use. Disagreement identifies a discrepancy, not motive. No human, neural signal, implant or stimulation tested.

Source: doc-aletheia-005 / Revision 1 / https://meridianrc.org/research/cycle-005/aletheia/

M-ALE-006 / A third account of the draw

2026-11-13 / Nadia Okafor, Tomás Ibarra

An independently supplied acquisition map exposes a discrepancy that two agreeing ledgers conceal, while hidden copying and conflicting signed corrections remain unresolved.

An observation before the ledgers

Aletheia's proposed source-discrimination harness must distinguish two genuinely separate symbol draws from two receipts referring to the same draw. The authority-timing note showed that agreeing authenticated ledgers can conceal a shared acquisition. Okafor and Ibarra now add a software checkpoint representing a tap at the draw, upstream of both ledger writers.

The test compares two root-to-draw maps with the tap's map. It does not count matching symbols as matching origins: two distinct draws may both yield a circle. The checkpoint must identify acquisition events, with one record for each root, rather than merely confirm that the resulting symbols agree.

What a separate account contributes

In the independent control, both ledgers and the tap map two roots to distinct draws. In the colluding-ledger case, both ledgers claim distinct draws while an honest tap preserves the shared draw. The procedure reports disputed acquisition records. It cannot use that discrepancy alone to determine which recorder is mistaken or dishonest.

A missing second tap entry holds coverage. A tap disclosed as a copy of ledger A holds dependency even when all values agree. In the undeclared-copy case, a false declaration makes the copied map indistinguishable from the independent control. The reviewer returns the same conditional independence state. Adding a third account has not solved a shared-origin problem by itself.

Authority accounts that cannot be collapsed

The synthetic draw occurs at minute 600. An authenticated withdrawal effective at 570 arrives at 660: the review at 630 is retained, and the review at 690 appends withholding. The demonstration signs the effective and received fields with public HMAC keys; this supplies a reproducible binding exercise, not a secured real-world authority service.

A second case supplies two correctly bound revision-2 corrections with conflicting effective/status pairs. Both become available before the later review. The procedure appends disputed authority correction rather than choosing whichever arrived last. A changed notice body fails its binding check; absent status coverage holds the decision separately.

The boundary of the study

The checkpoint design now needs a credible tap origin, separate failure dependencies, coverage evidence and an authority process that can resolve contradictory corrections without rewriting earlier reviews. These requirements concern a proposed harness. No participant, neural signal, implant or stimulation was tested, and no exogenous information was inferred.

Retained result table

Case	Retained decision history (minutes → state)
independent_control	630 → independent_within_declared_checkpoint
colluding_ledgers_honest_checkpoint	630 → disputed_acquisition_records
checkpoint_gap	630 → withhold_checkpoint_gap
declared_copied_checkpoint	630 → withhold_checkpoint_dependency
undeclared_copied_checkpoint	630 → independent_within_declared_checkpoint
late_withdrawal	630 → independent_within_declared_checkpoint; 690 → withhold_authority_at_draw
conflicting_corrections	630 → independent_within_declared_checkpoint; 690 → disputed_authority_correction
bad_notice_binding	630 → withhold_notice_binding
missing_status_coverage	630 → withhold_status_coverage

Scope retained for review

Toy checkpoint maps and public HMAC keys. An honest separate tap exposes a discrepancy, not which custodian lied. Falsely declared separation retains a blind spot. Authenticated conflicting corrections remain unresolved. No brain signal, participant, implant, exogenous information or clinical result is tested.

Source: doc-aletheia-006 / Revision 1 / https://meridianrc.org/research/cycle-006/aletheia/

M-ALE-007 / The request that must not draw twice

2026-11-21 / Nadia Okafor, Tomás Ibarra

Interrupted acquisition distinguishes replaying a retained symbol from issuing a second draw or restoring an old permission.

A lost acknowledgement

Aletheia proposes a symbol-acquisition harness before any participant or implant study. If a draw is committed and the acknowledgement is lost, retrying the request must not silently select a new symbol. The independent checkpoint study left origin and authority questions unresolved; this study adds interruption to the request path.

The authored request uses a stable identifier and a request-content binding. A retained commit associates it with symbol-3. Retrying that same request returns symbol-3, rather than invoking another draw. If the journal records only a pending request, the harness holds: it cannot tell whether a draw occurred outside the retained acknowledgement.

Recovery preserves adverse states

Eight cases include a reused identifier bound to different request content, two conflicting committed symbols, a missing journal, a withdrawal and contradictory authority notices. None is repaired by generating a replacement draw. A retained symbol may be replayed as a record while its use remains withheld, unknown or disputed.

The notice model assumes supplied authority inputs; it does not verify signatures or define a real resolution authority. Opposing allowed and withdrawn assertions remain disputed in either arrival order. No restart changes that policy, and the current study does not resolve the earlier authenticated-correction conflict.

The missing tail

The final case silently removes a second draw and a withdrawal before review. The remaining journal and notice list are byte-for-byte equivalent as inputs to the favorable retry case. Both produce the same replay decision. This is a retained failure, not an exactly-once guarantee: the model has no external witness from which to recover the missing tail.

Next work must establish the durable commit boundary, independent tap origin and how authority coverage is recovered after an outage. The fixture has no entropy hardware, neural measurement, participant, implant or exogenous observation. A reliable request replay would still be only one part of an acquisition system.

Retained result table

Case	Recovery action	Draw	Status
retry_after_commit	replay_same_draw	symbol-3	conditional_model_use
retry_before_ack	hold_unacknowledged	Unavailable	unknown
same_key_new_request	reject_key_reuse	Unavailable	conflict
two_committed_draws	hold_conflicting_draws	Unavailable	conflict
withdrawal_survives_restart	replay_record_only	symbol-3	withheld
contradictory_notices	replay_record_only	symbol-3	disputed
journal_missing	hold_journal	Unavailable	unknown
hidden_journal_truncation	replay_same_draw	symbol-3	conditional_model_use

Scope retained for review

Request bindings, journals and authority notices are trusted fixture inputs, not an authenticated distributed protocol. The model does not guarantee exactly-once acquisition outside its retained journal. No physical entropy source, brain signal, participant, implant or exogenous result.

Source: doc-aletheia-007 / Revision 1 / https://meridianrc.org/research/cycle-007/aletheia/

M-ALE-008 / A copy beyond the missing tail

2026-11-28 / Nadia Okafor, Tomás Ibarra

An outside checkpoint exposes unequal acquisition histories while retaining a shared-loss blind spot.

Two requests, two records

Aletheia's proposed acquisition harness must preserve which symbol a request received. The retry study demonstrated that a lost acknowledgement must not cause an automatic replacement draw, and retained a silently truncated journal that looked favorable. Okafor and Ibarra now ask what changes if a second commit copy survives.

The control has ordered commits R1 to symbol-3 and R2 to symbol-5. A second authored list contains the same commits. Comparison returns conditional model agreement. This is a comparison of input records, not a demonstration of an independent physical tap or an authenticated checkpoint service.

Asymmetric loss becomes visible

When the local copy loses R2 but the witness retains it, unequal length holds the comparison. Loss from the witness alone also holds; the outside copy is not presumed infallible. A changed symbol at equal length is caught by content comparison, and an unreadable witness produces an availability hold.

When both copies lose R2, they agree again. The shared-truncation case is favorable under the available inputs even though the simulator retains the missing history. It preserves the preceding study's failure in a more specific form: two copies can share the same omission. An additional copy is useful only for the losses it actually survives.

Agreement is not permission

These six cases do not generate another draw or overwrite the earlier journal. They also do not decide authority status. A withdrawn or disputed permission from the preceding recovery study remains withdrawn or disputed even when two commit lists agree.

The next design must place checkpoint creation across the actual commit boundary and explain how one custodian could fail without taking both histories with it. Independent origin, durable storage and authority coverage remain separate requirements. No neural measurement, participant, implant, physical entropy source or exogenous effect is involved.

Retained result table

Case	Checkpoint decision
complete_copies	conditional_model_agreement
local_tail_lost	hold_length_disagreement
witness_tail_lost	hold_length_disagreement
changed_symbol	hold_content_disagreement
witness_unavailable	hold_witness_unavailable
shared_truncation	conditional_model_agreement

Scope retained for review

Synthetic list comparison only. The outside copy is a separate model input, not a demonstrated independently originated tap. Copy agreement neither proves acquisition origin nor authorizes use. Authority notices remain governed by the prior recovery study; no participant, implant or entropy hardware.

Source: doc-aletheia-008 / Revision 1 / https://meridianrc.org/research/cycle-008/aletheia/

M-ALE-009 / Before the acknowledgement leaves

2026-12-07 / Tomás Ibarra, Nadia Okafor

A power-loss state machine distinguishes a flushed checkpoint from a queued write without drawing again.

R2 passes through four steps

Aletheia's acquisition harness must preserve the symbol assigned to a request when an acknowledgement disappears. The checkpoint comparison showed that two lists could share the same missing tail. Ibarra now specifies an order: draw once, flush the local commit, flush the witness commit, then send the acknowledgement. Okafor tests cuts between those steps.

The authored R2 draw is symbol-5. The model stores that request and symbol locally and in a witness, then makes records unreadable during an outage according to explicit power-domain assumptions. Its recovery reviewer sees only readable flushed copies and the supplied authority state. Flushed bytes remain in modeled storage through the outage; physical durability is not established. All nine cases retain exactly one draw; none invokes a replacement acquisition.

A different supply is only part of the answer

When the witness has flushed and its modeled supply is distinct, local power loss leaves R2 available at the witness. Recovery holds reconciliation before the local record can be restored. Witness loss leaves the local draw intact but holds witness agreement. A cut after the local write and before the witness flush produces that same incomplete state.

A shared-supply outage makes both flushed copies unreadable. With the local supply down and the witness write only queued, neither readable durable copy is available either. Both cases remain unresolved without another draw. The latter failure survives a separate-supply label: there was no durable witness record for that supply to preserve.

A recovered symbol still has a source

Matching commits do not override withdrawn or disputed permission; both authority cases hold. The hidden upstream-copy case also remains beside the favorable committed pair. Their reviewer inputs are identical, while the authored history differs on whether the symbol was independently acquired. Durability is not source independence.

This is a software state machine with explicit flush and power assumptions, not a physical storage test. It identifies the interval a prototype must preserve but establishes no actual supply separation, independent tap or authority service. There is no participant, implant or neural measurement. The next design must expose

durable witness completion before acknowledgement and retain the unresolved interval through restart.

Retained result table

Case	Recovery decision	Draw count
committed_pair	replay_existing_model	1
local_loss_after_witness	recover_from_witness_hold_reconciliation	1
witness_loss	retain_local_hold_witness	1
cut_before_witness	retain_local_hold_witness	1
shared_supply_loss	hold_unresolved_no_redraw	1
queued_witness_loss	hold_unresolved_no_redraw	1
withdrawn_permission	hold_authority	1
disputed_permission	hold_authority	1
hidden_upstream_copy	replay_existing_model	1

Scope retained for review

An authored state machine, not a powered storage experiment. Supply separation, flush behavior and authority are explicit model inputs. No independently originated tap, participant, implant or entropy hardware. No redraw or resolution of earlier conflicting notices is authorized.

Source: doc-aletheia-009 / Revision 1 / https://meridianrc.org/research/cycle-009/aletheia/

M-ALE-010 / A receipt cannot identify a signal

2026-12-16 / Tomás Ibarra, Nadia Okafor

Witness readback is tested beside a synthetic command-echo control, separating durable copying from acquisition origin.

R3 reaches the witness

The Aletheia acquisition harness retains one symbol for each request through interruption. Following the write-boundary study, Ibarra adds a witness readback containing both R3 and symbol-2. The recovery reviewer compares that receipt with the readable local and witness commits before allowing replay of the existing record.

Nine authored receipt cases distinguish a matching readback from a queued-only witness, missing or stale receipt, different symbol and unavailable local copy. Withdrawn and disputed authority still hold. No recovery case requests another draw. The receipt is supplied under explicit durable-storage assumptions; its contents do not prove that a physical write completed.

An echo can be copied perfectly

A separate source control addresses the kind of signal an eventual interface would preserve. The WAND preprint reports simultaneous recording and stimulation with artefact cancellation in a nonhuman primate. It makes stimulation artefact a concrete design concern for Aletheia, without supporting psychic effects or transferring device validation to this harness.

Okafor supplies the eight-value command trace 0,1,0,-1,0,1,0,-1. One synthetic source is exactly three times that trace; its cosine similarity is 1. A different trace gives 0. The first is flagged as a command echo. The second receives no echo flag, which is not proof of independent origin. These toy arrays are neither neural recordings nor a reproduction of the WAND method.

Two different next experiments

The hidden upstream-copy control has exactly the same receipt inputs as the favorable case. Perfect agreement after the split cannot reveal where a symbol originated before it. Future electrical source tests must include command-only, source-only and combined-input controls, and examine timing offsets and nonlinear artefacts rather than relying on this exact-match screen.

A physical storage interruption test and an electrical artefact test answer different questions. Both remain to be performed. The synthetic receipt sequence establishes no implant, participant result, biological signal or resolution of the earlier authority dispute.

Retained result table

Case	Decision	Draw count
readback_pair	replay_existing_model	1
queued_only	hold_incomplete_no_redraw	1
missing_receipt	hold_receipt_mismatch	1
stale_receipt	hold_receipt_mismatch	1
mismatched_witness	hold_content_disagreement	1
local_unavailable	hold_incomplete_no_redraw	1
withdrawn	hold_authority	1
disputed	hold_authority	1
hidden_source_copy	replay_existing_model	1
scaled_command_echo	flag_command_echo	Not an acquisition
different_trace	no_echo_flag_not_origin_proof	Not an acquisition

Scope retained for review

Authored symbols, source arrays and durable-readback assumptions; no storage hardware, neuronal recording, participant, implant or psychic result. The echo screen is a toy counterexample, not a validated artefact rejection algorithm. Earlier authority conflicts remain unresolved.

Source: doc-aletheia-010 / Revision 1 / https://meridianrc.org/research/cycle-010/aletheia/

M-ALE-011 / The echo one sample later

2026-12-24 / Nadia Okafor, Tomás Ibarra

A finite lag search catches one delayed command echo and leaves nonlinear and mixed sources unresolved.

The interface can hear its own command

Aletheia's proposed amplifier must preserve an acquisition without mistaking its own command for an independently originated signal. The receipt and source study compared an exact scaled echo with a different trace. That zero-lag screen had no reason to recognize a command displaced in time.

Okafor now supplies the twelve-value command 0,1,-2,0,3,-1,0,2,0,-3,1,0 and a separate authored source 1,0,1,-1,0,2,-2,0,1,1,0,-1. Six controls contain the command multiplied by three, that command delayed one sample, the squared command, the source alone, command-times-three plus source, and silence. These are synthetic arrays, with no sampling-rate or biological interpretation.

An offset changes what the screen sees

Ibarra compares each source with command shifts from -2 through +2 samples, padding uncovered positions with zero. The table retains the unshifted cosine and the strongest absolute cosine in this finite set. It reports the signed cosine and selected lag; it does not estimate a physical conduction delay.

The delayed-command control has unshifted cosine -0.275862 and cosine 1 at lag +1. The command-only control gives 1 at zero lag. Both receive an exact-linear-echo flag. Squared command and combined input remain unresolved, even though their authored construction contains command information. Silence has zero energy, so its cosine is undefined rather than evidence of independence.

A useful failure remains a failure

The source-only control's strongest tested cosine is 0.595550 at lag +1. A correlation in a short array does not by itself identify an echo, and lack of an exact match does not identify a neural source. The screen deliberately flags only exact linear agreement; it is not a general artefact detector.

No samples are subtracted or discarded. A future electrical experiment must retain command-only, source-only and combined measurements and include delays and nonlinear response. Powered witness completion remains a separate unperformed test. No participant, neural recording, implant or resolution of the earlier authority dispute is introduced.

Retained result table

Synthetic source	Zero-lag cosine	Best tested lag (samples)	Best cosine	Screen outcome
command_only	1.0	0	1.0	flag_exact_linear_echo
delayed_command	-0.275862	1	1.0	flag_exact_linear_echo
squared_command	0.01323	2	-0.053858	unresolved_source
source_only	-0.347404	1	0.59555	unresolved_source
combined	0.973224	0	0.973224	unresolved_source
quiet	undefined	undefined	undefined	undefined_zero_energy

Scope retained for review

Authored arrays and a deliberately narrow exact-linear screen. A finite lag search neither detects arbitrary artefacts nor establishes independent origin. No physical electrical harness, powered witness test, neural recording, implant or participant. Existing authority disputes remain unresolved.

Source: doc-aletheia-011 / Revision 1 / https://meridianrc.org/research/cycle-011/aletheia/

M-ALE-012 / When the inputs stop adding

2027-01-01 / Nadia Okafor, Tomás Ibarra

An authored interaction term defeats prediction from isolated command and source responses, constraining the next electrical test.

The combined acquisition is a separate condition

Aletheia's intended amplifier must record a signal while its own commands are active. The delayed-echo study retained mixed and nonlinear sources without classifying them. Okafor and Ibarra now ask a prior question: would the command-only and source-only acquisitions predict what happens when both inputs are present?

The inputs are dimensionless authored amplitudes c and s. Twelve deterministic cases use six input pairs under two specified laws: y = 3c + 2s and y = 3c + 2s + 0.5cs. The latter term is deliberately introduced interaction. These are equations evaluated in software, with no sample rate, electrical transfer measurement or biological source.

The isolated responses can look perfectly ordinary

For each pair the command-only response uses s=0, and the source-only response uses c=0. The interaction term vanishes in both isolated conditions. Their sum therefore predicts the linear combined response exactly, with zero residual in all six additive controls.

In the interaction cases the combined output differs by 0.5, -0.5, -0.5, 0.5, 1 and 0 across pairs (1,1), (1,-1), (-1,1), (-1,-1), (2,1) and (0,2). Reversing one input changes the sign; increasing the command changes the difference. The zero-command control gives zero despite the interaction law remaining present.

Subtraction does not name the remaining source

Removing a separately measured command-only trace would leave 2s + 0.5cs under the specified interaction law. That remainder still depends on the command. Its existence cannot identify a neural contribution, let alone support the amplifier's extraordinary ambitions.

The next electrical protocol must retain isolated and combined conditions, input reversals, amplitudes and raw output. No subtraction is authorized for acquired samples. Powered witness completion remains a separate unresolved experiment; this calculation supplies neither a storage result nor evidence from a participant or implant.

Retained result table

Command	Source	Solo sum	Combined	Difference
1	1	5.0	5.5	0.5
1	-1	1.0	0.5	-0.5
-1	1	-1.0	-1.5	-0.5
-1	-1	-5.0	-4.5	0.5
2	1	8.0	9.0	1.0
0	2	4.0	4.0	0.0

Scope retained for review

Authored deterministic arithmetic, not measured electronics or a reproduced neural-interface technique. The interaction law is deliberately specified. Real response, noise, calibration, powered witness durability and biological/extraordinary origin remain untested. No samples are removed and no participant or implant is involved.

Source: doc-aletheia-012 / Revision 1 / https://meridianrc.org/research/cycle-012/aletheia/

M-ALE-013 / The order in the signal

2027-01-08 / Nadia Okafor, Tomás Ibarra

An ordered additive model can imitate an interaction contrast; paired order cancels the specified linear drift while leaving a quadratic confound.

Four conditions take time

Aletheia's proposed amplifier must distinguish its command response from the signal it records while commands are active. The combined-input study showed an interaction hidden from solo conditions. Okafor and Ibarra now test whether a similar difference can appear simply because the four acquisitions happen at different times.

Their software uses binary dimensionless inputs c and s, with output $y = 3c + 2s + kcs + dt + qt^2$ at consecutive integer time steps. The coefficients deliberately specify interaction, linear drift and quadratic drift. No voltage, electrical acquisition, participant or implant is involved. Inputs and orders are retained January 5; eight cases run January 7.

Reverse the order before naming the remainder

The forward order is (0,0), (1,0), (1,1), (0,1), at times 0 through 3. The reverse order visits those conditions backwards. For repeated conditions the program averages their raw outputs, then calculates the four-condition contrast y11 - y10 - y01 + y00. With no drift the contrast is zero for the additive law and 0.5 for k=0.5.

With k=0 and linear drift d=0.25, the forward contrast is -0.5 and the reverse contrast is +0.5. The latter equals the static interaction case even though its law has no interaction term. A forward sequence followed immediately by its reverse uses times 0 through 7; equal mean acquisition times cancel this particular linear drift. Its additive contrast is zero, while k=0.5 gives 0.5.

A cancellation with an explicit limit

The same paired order fails to cancel the specified quadratic drift. With q=0.125 and no interaction, the contrast is 1; adding k=0.5 makes it 1.5. All raw ordered rows remain in the fixture output. A nonzero contrast alone therefore cannot identify interaction, much less the physical or biological origin of a signal.

Before electrical work, the protocol must retain acquisition order, repeated conditions, timing and controls that can reveal changing baseline behavior. No acquired samples are subtracted or discarded by this exercise. Independently powered witness completion remains a separate open task, and the amplifier's speculative ambition has acquired no new empirical support.

Retained result table

Ordered model case	Interaction coefficient	Linear drift / step	Quadratic coefficient	Four-condition contrast
static_additive	0	0	0	0.0
static_interaction	0.5	0	0	0.5
linear_forward	0	0.25	0	-0.5
linear_reverse	0	0.25	0	0.5
balanced_linear	0	0.25	0	0.0
balanced_interaction_linear	0.5	0.25	0	0.5
balanced_quadratic	0	0	0.125	1.0
balanced_interaction_quadratic	0.5	0	0.125	1.5

Scope retained for review

Dimensionless authored arithmetic over explicit acquisition order. No voltage, physical electronics, storage completion, participant, neural signal, implant or extraordinary source is measured. Balanced order cancels only the specified linear component; nonlinear drift remains a confound.

Source: doc-aletheia-013 / Revision 1 / https://meridianrc.org/research/cycle-013/aletheia/

IV. ET/ED research frameworks and retained exception

Documentary section

These records are part of the research account. The program frameworks specify hypotheses, instrumentation and governance; they do not report completed trials. The retained exception preserves its unresolved provenance and competing explanations. Preparedness levels describe integration into research design, not evidence strength.

P-SEI / Space Elevator: exogenous environment and interaction

Exogenous environment program

Design for the possibility that we are not alone in the operating environment.

A structure extending from the equatorial atmosphere through the magnetosphere and beyond geostationary orbit cannot responsibly define its environment as weather, debris and human spacecraft alone. Meridian maintains explicit, testable hypothesis classes for non-terrestrial intelligence, anomalous energetic phenomena and non-local or higher-dimensional interaction. None is treated as established fact; none is excluded by policy before the instruments have had a chance to look.

Operating assumption / AE-01 Observe before explaining.

The elevator may become more than transportation infrastructure. A continuously instrumented, ~100,000 km structure would be an unprecedented longitudinal observatory - and potentially a persistent human-made interface with phenomena, agencies or energetic regimes that current aerospace systems encounter only briefly.

Meridian does not assert that extraterrestrial or extra-dimensional beings are presently interacting with Earth. The engineering requirement is narrower and actionable: if an anomalous agent or field couples to the tether, climbers, power system or surrounding space, the system should detect it, preserve the evidence, remain safe and avoid destroying the event through premature interpretation.

Instrumentation

The tether as a distributed observatory

Structural sensing is expanded into a scientific array. Time-synchronized stations distributed along the tether establish a baseline across atmosphere, ionosphere, radiation belts, GEO and the upper extension, allowing local disturbances to be distinguished from coherent events spanning large distances.

Optical Visible, near-IR, thermal IR and UV imaging with calibrated astrometry.

EM / RF Broadband spectrum monitoring, electric-field probes, magnetometers and induced-current telemetry.

Particle Radiation, charged-particle and plasma measurements correlated with structural events.

Geometry Precision ranging, clock comparison, accelerometry and gravimetric measurements for motion or propagation anomalies.

Hypothesis envelope

ET / ED classes remain explicitly open

Events that survive conventional explanation are not collapsed into a single category. Meridian maintains competing hypotheses so that data collection can discriminate among them instead of merely assigning an "unknown" label.

H_1 Uncatalogued terrestrial, atmospheric, orbital or technological source.

H_2 Non-terrestrial technological artifact, vehicle, probe or directed signal.

H_3 Responsive or intelligent agency not consistent with known human systems.

H_4 Non-local / extra-dimensional interaction: apparent discontinuities in ordinary propagation, geometry or energetic coupling.

Engineering protection

Interaction without surrendering the machine

An unexplained event becomes an engineering problem the moment it couples energy or information into the structure. The elevator therefore separates observation, control and safety systems so that an anomalous input cannot automatically propagate into system authority.

Safing Climbers enter controlled hold or retreat states; traffic is cleared from the affected region.

Electrical Segmented conductive paths, isolation, surge diversion and sacrificial interfaces limit energetic propagation.

Control Anomaly detection has read access to rich telemetry but no direct command authority over critical actuators.

Evidence Raw sensor streams are cryptographically timestamped and mirrored to geographically independent archives.

Contact protocol

No unilateral response

A coherent, apparently responsive signal or encounter is not treated as permission to improvise. Meridian adopts the logic of contemporary SETI post-detection practice: independent verification first, transparent preservation of evidence and broad consultation before any intentional reply.

Verify Require independent instruments, modalities and external observers wherever physically possible.

Contain Separate physical samples, unknown code/data payloads and anomalous RF content from operational networks.

Disclose Maintain a prewritten scientific communication path that distinguishes observation from interpretation.

Respond No deliberate transmission or behavioral response to apparent intelligence without international review.

AE event classification

Escalation is based on evidence quality and interaction, not strangeness.

AE-0

Environmental anomaly

Unexpected observation with insufficient evidence for novelty. Preserve high-rate telemetry and run automated checks against weather, satellites, aircraft, debris, sensor failure and known electromagnetic sources.

AE-1

Correlated anomaly

Event appears on multiple calibrated sensors or physically separated stations. Freeze calibration state, expand collection bandwidth, request independent observation and prevent routine data cleanup from discarding the transient.

AE-2

Structured / responsive event

Observed behavior appears contingent on the elevator, its emissions or nearby activity. Suspend experimental stimulation. Establish blind challenge tests where safe, with analysis teams separated from operators to reduce expectation effects.

AE-3

Exogenous hypothesis active

Conventional explanations fail after external review and data support technological or agent-like structure. Activate scientific, legal, international and contact-governance teams; operational safety takes precedence over pursuit of additional evidence.

AE-4

Non-local signature

Measurements reproducibly imply behavior inconsistent with ordinary line-of-sight propagation, local energy accounting or expected space-time geometry. Treat "extra-dimensional" as a hypothesis to be experimentally constrained - not a conclusion - and prioritize independent clocks, ranging, field measurements and replication.

Investigation agenda

Research the encounter before there is an encounter.

Baseline census

Build years of normal-environment data before interpreting rare events. The Galileo Project's multimodal observatory work is a useful terrestrial model: calibrated optical/IR, acoustic, RF, magnetic, particle and weather instrumentation operating continuously rather than relying on eyewitness reports.

Cross-domain fusion

Develop event reconstruction that combines tether strain, climber telemetry, orbital catalogs, radar, optical imaging, magnetometry, RF spectra and particle environment. The research target is not a dramatic image; it is a physically constrained event with independent observables.

Agency tests

Pre-register methods for distinguishing random or environmental correlation from apparent response. Candidate interactions can be examined using randomized timing, blinded stimulus protocols and replicated sensor stations without immediately attributing intent.

Non-local physics

Define in advance what would count as evidence for an ED/non-local model: reproducible timing discrepancies, correlated separated events without an ordinary propagation path, unusual field coupling, or geometry inconsistent with independently verified trajectories. Instrument precision determines whether the hypothesis is testable at all.

Technosignatures nearby

Extend SETI logic from distant astronomy into the elevator's local operating volume. Search for narrow- and broadband artificial signals, unusual modulation, persistent station-keeping behavior, material artifacts and energy signatures that are difficult to produce through known natural processes.

Human factors

Prepare for cognitive and institutional failure modes as seriously as mechanical ones: confirmation bias, stigma, secrecy incentives, social-media contamination, hoaxes and premature disclosure. Competing analysis teams and open evidentiary thresholds protect the science in either direction.

Evidence standard: current NASA and AARO work supports improved scientific collection of unresolved anomalous phenomena but does not establish extraterrestrial or extra-dimensional origin. Meridian's ET/ED program therefore treats those origins as explicit hypotheses within a broader anomaly-resolution framework. The 2026 IAA SETI post-detection principles provide a useful precedent for independent verification, durable data preservation and restraint in intentional response.

Source: P-SEI / https://meridianrc.org/projects/space-elevator/#anomalous

P-LUM / Lumen: exogenous interaction and contact research

Exogenous / non-local interaction program

Assume the beam may become an interface.

A high-power coherent transmitter on the Moon is unusual infrastructure: precisely timed, highly directional, continuously monitored and visible across astronomical distance under the right geometry. Lumen therefore treats interaction with non-human, extraterrestrial, non-local or extra-dimensional agencies as explicit hypothesis classes within the observatory program - not established facts, and not subjects to be silently excluded from instrumentation or procedure.

Operating assumption / EX-04 Transmit. Observe. Do not assume silence.

The communications base is deliberately instrumented as a possible contact surface. If an external intelligence, anomalous agent or unfamiliar energetic regime notices, intercepts, modulates, shadows, returns or otherwise couples to the optical link, the system should recognize the event before ordinary link-control software explains it away.

Level-4 posture: Lumen explicitly considers extraterrestrial intelligence and extra-dimensional/non-local interaction plausible enough to justify dedicated sensing, safing and governance. It does not claim such entities are currently present. The program's obligation is to make a genuinely anomalous event detectable, repeatable where possible, and difficult to contaminate with wishful interpretation.

H1

Deliberate optical response

Coherent return, structured modulation or timing behavior not attributable to known spacecraft, Earth terminals, scattering or instrument feedback.

H2

Physical interception

Unexpected occultation, beam distortion, localized thermal/particle signatures or geometry changes consistent with an unknown object or agent crossing the path.

H3

Non-local coupling

Correlated phase, timing, polarization or sensor events without an ordinary propagation path; investigated first as clock, software, detector or model failure.

H4

Extra-dimensional interaction

Persistent observations that require apparent discontinuities in path length, causally unusual correlation or energy/information coupling not captured by the operational physical model.

Instrumentation

The comm terminal watches outside its own beam

Link telemetry alone is too easy to misread. Lumen adds independent sensors with shared precision time so apparent anomalies can be tested against optical, electromagnetic, particle, geometric and environmental channels.

Optical Wide-field visible/NIR/thermal imagers, spectrometers and calibrated photometry co-boresighted with the communications path.

Polarization Full transmit/receive polarization state retained as raw telemetry rather than reduced to link-quality metrics.

EM / RF Broadband spectrum, electric-field and magnetometer measurements surrounding each high-interest event.

Timing Independent clocks, one-way/two-way ranging and phase histories used to test apparent non-local or propagation anomalies.

Beam behavior

Do not let adaptive optics erase the evidence

Normal beam-control loops are designed to cancel perturbations. During a qualified anomaly they can destroy the signature being studied. Lumen therefore records pre-correction wavefront data and supports a protected observation mode that freezes or bounds adaptive compensation while maintaining laser safety.

Raw Preserve wavefront sensor data before control-loop filtering and correction.

Shadow Run independent estimators that observe but cannot drive steering.

Replay Reconstruct exactly what the controller saw and what correction it applied.

Probe Only pre-authorized low-risk diagnostic patterns may be transmitted during investigation.

Engineering protection

Contact cannot inherit beam authority

An apparently intelligent or responsive pattern is still untrusted input. Observation, command, networking and high-power laser control are separated so anomalous content cannot automatically alter pointing, modulation, power level or operational policy.

Control Anomaly classifiers are read-only with respect to safety-critical actuation.

Data Unknown code, compressed payloads or structured messages are quarantined from operational networks.

Power Unexpected energy deposition or return triggers aperture closure, derating and crew stand-off.

Evidence Raw streams are hash-chained, mirrored to Earth and retained with calibration state and software versions.

Contact protocol

Recognition is not permission to answer

The 2026 IAA SETI principles provide a current precedent: independent verification, careful communication, broad consultation and no reply to confirmed ETI before international consideration. Lumen extends that restraint to any apparently responsive non-human or non-local event.

Verify Require independent instruments, Earth stations and external observers whenever physically possible.

Disclose Separate observations, confidence estimates and hypotheses in public reporting.

Consult Escalate confirmed intelligence indicators to international scientific and governance channels.

Respond No deliberate semantic reply, beam-following behavior or escalation of power without explicit authorization.

EX-0

Link anomaly

Ordinary troubleshooting first: pointing, clock, atmosphere at Earth gateway, detector saturation, dust, thermal drift, relay state, software, known traffic and ephemerides.

EX-1

Cross-sensor correlation

Preserve raw data and request independent observations when an event appears in at least two physically distinct channels or persists beyond expected equipment behavior.

EX-2

Agency indicator

Escalate when structure, contingency or repeated response is statistically incompatible with known systems and environmental models. Freeze autonomous adaptation that could create a feedback illusion.

EX-3

Proximal / energetic event

Protect crew and hardware first: laser safe, isolate affected networks, move personnel to shielded or interior positions, maintain passive observation and avoid physical approach unless hazards are characterized.

EX-4

Confirmed non-human interaction

Preserve globally verifiable evidence, invite independent analysis, place the facility under pre-agreed international scientific/governance procedures and prohibit unilateral messaging or behavioral response.

Investigations designed to discriminate, not merely detect.

Coherent returns

Challenge candidate returns with wavelength shifts, polarization, timing, angle-of-arrival and known reflector models. A "reply" must survive internal reflections, retroreflectors, relay cross-talk and Earth-origin contamination.

Beam shadowing

Use multiple separated apertures and simultaneous wide-field imaging to distinguish real path interception from local dust, detector artifacts or adaptive-optics transients.

Phase discontinuity

Compare independent clocks and optical/radio ranging. Apparent path-length jumps are first assigned to timing distribution, oscillator error, software, multipath and ephemeris failure before non-local hypotheses enter the candidate set.

Responsive patterns

Pre-register stimulus families for low-power diagnostic sequences so later correlations can be tested against a known statistical plan rather than improvised after an exciting event.

Human factors

During high-interest events, crew narratives are collected separately from instrument analysis to preserve witness information without allowing expectation or authority gradients to rewrite the technical record.

Materials coupling

Inspect exposed coatings, optical surfaces and semiconductor packages after unexplained energetic events for localized lattice, chemical, magnetic or radiation signatures that ordinary link telemetry cannot capture.

Evidence posture: NASA's UAP work supports improved scientific collection of unresolved anomalous phenomena but does not establish extraterrestrial origin. The IAA's June 2026 SETI post-detection principles specifically address technosignature verification and communication, and counsel against an intentional reply before international consultation. Lumen's extra-dimensional/non-local hypothesis class is more speculative still; it exists here as an instrumentation and falsification target, not a conclusion.

Source: P-LUM / https://meridianrc.org/projects/lumen/#exogenous

P-ALE / Aletheia: exogenous cognition

Exogenous cognition program · ET/ED-5

Assume the interface could be bidirectional.

At Level 5, ET/ED interaction is not confined to an anomaly appendix. It is built into implant architecture, trial design, safing, epistemology, network isolation and governance. The program explicitly studies whether amplification might reveal information or agency that is not adequately explained by the user's endogenous neural activity - while preserving conventional neurological and instrumentation explanations as the default competitors.

Operating assumption / AX-05 The implant may be a receiver. It may also be a beacon.

A device optimized to increase sensitivity to weak or liminal cognitive signals could, in principle, alter more than what the user can perceive. If unknown intelligences, non-local information channels, unfamiliar fields or higher-dimensional couplings exist, the implant could change the detectability, bandwidth or stability of an interaction in either direction.

This is a research hypothesis, not an empirical claim that extraterrestrial or extra-dimensional beings communicate with humans. Level 5 means the architecture takes the possibility seriously enough to detect, falsify, contain and document it without relying on belief.

Hypothesis envelope

Keep multiple origins alive.

Apparent contact is decomposed into competing source classes rather than collapsed into one metaphysical answer.

H0 Endogenous imagery, memory, dream intrusion, dissociation or confabulation.

H1 Stimulation-induced self/other or presence illusion.

H2 Environmental electromagnetic, acoustic, optical or network artifact.

H3 Unknown but conventional physical process.

H4 Non-terrestrial technological or biological agency.

H5 Non-local / extra-dimensional agency or information coupling.

Instrumentation

No contact claim from brain data alone.

Every high-value event is time-locked against external instruments so that internal experience and external environment can be compared without retrospective story-building.

Neural raw intracranial state + stimulation telemetry.

EM / RF broadband spectrum + local field probes.

Optical calibrated visible / IR monitoring where appropriate.

Physical magnetic, vibration, temperature, power and clock anomalies.

Social independent observers and synchronized blinded participants.

Contact quarantine

Content never authenticates itself.

A perceived agent may be articulate, emotionally precise, predictive or apparently impossible for the subject to know. None of those properties grants the content permission to control the implant or the user's environment.

Rule No requests for secrecy are honored.

Rule No instructions execute automatically.

Rule No self-identification is accepted as provenance.

Rule Predictive claims must be pre-registered before testing.

Rule Unknown-origin data remain sandboxed.

Experimental program

Make the extraordinary survive controls.

Research prioritizes protocols where a real exogenous effect would produce a measurable pattern that hallucination, expectation and instrumentation error are unlikely to reproduce.

Paired isolated participants tested for time-locked shared content.

Shielded Faraday / sensory-isolated sessions with matched sham periods.

Blind hidden targets, random schedules and independent scoring.

Cross-site replication on synchronized systems in distant laboratories.

Fail-safe contact protocols tested first with adversarial simulations.

01 Isolate the event from adaptive stimulation

02 Preserve raw neural + environmental data

03 Run conventional artifact and clinical checks

04 Attempt blinded independent replication

05 Escalate only to human contact-governance review

Level-5 principle: the more an event resembles intelligent contact, the less authority the implant is allowed to exercise. Any entity - real, imagined, endogenous, extraterrestrial, non-local or otherwise - that attempts to bypass verification, disable safing, isolate the user, demand exclusive trust or modify its own evaluation pathway is treated as an adversarial condition until proven otherwise.

Source: P-ALE / https://meridianrc.org/projects/aletheia/#anomalous

LUM-ACK-1031 / Uncommanded acknowledgement

Lumen / retained exceptions / ET/ED 4

Uncommanded acknowledgement

PRESERVE · ISOLATE · DO NOT RESPOND

Archive reference

LUM-ACK-1031

Date carried by fragment

31 October 2026 · clock not certified

Source

Receiver-log transcription; original capture unavailable

Disposition

Quarantined · provenance unresolved

The fragment describes an acknowledgement received while transmission was disabled. Its first fields match an earlier diagnostic sequence, including a padding error carried by the archived test vector.

Replay, stale-buffer reuse and a transcription fault remain sufficient explanations for that match. The final field is absent from the diagnostic sequence.

Retained transcription

23:59:41 TX_ENABLE false

23:59:41 REQUEST_QUEUE empty

23:59:42 RX_CLASS acknowledgement

23:59:42 HEADER_MATCH archived diagnostic

23:59:42 TIMING_WITNESS unavailable

23:59:42 PAYLOAD [retained below]

PLEASE RETAIN THE CHANNEL AFTER THE OBSERVER LEAVES.

TX_ENABLE is a recorded software flag. It does not establish the condition of a physical transmitter. No independently retained optical capture accompanies the fragment.

Two marginal notes

"Observer" may name the receiving process. A diagnostic could request that its parent leave a channel open after detachment.

The second note underlines leaves: Why would the process describe its own termination as a departure?

Neither annotation has a verified author. Neither establishes an external source.

ET/ED disposition

Retain the supplied transcription unchanged. Keep the decoder isolated from control and response paths. Seek an original capture and an independent timing witness before attributing the acknowledgement.

No reply is authorized. The text grants no permission to retain a channel, alter a shutdown sequence or initiate contact.

Level 4 describes preparedness for exogenous hypotheses. It is not a finding about the fragment's origin.

Inspect the operator's shutdown note

The sentence also appears in the shutdown log.

The note identifies the field as an integer status code. It records a check of the export template and finds no text substitution.

That field does not accept text.

This assertion belongs to the unverified note. The originating file and exporter have not been recovered.

Source: LUM-ACK-1031 / https://meridianrc.org/archive/uncommanded-acknowledgement/

A. Analytical notes

A.1 Common movement and a quiet difference

Consider two displayed changes x1 = m + b1 and x2 = m + b2, where m is a shared target contribution and b1 and b2 collect channel-specific terms. Their difference is x1 - x2 = b1 - b2. The common term disappears. This cancellation is an algebraic property of the chosen difference; it does not establish m = 0.

For an illustrative pair b1 = 1 and b2 = 0, shared movements m = 8, 12 and 10 produce pairs (9,8), (13,12) and (11,10). Every pair differs by one. These values match the changes retained at the end of the January 6 hour, after the hand movement, and after attempted restoration. The decomposition is illustrative: the observed arithmetic does not uniquely determine either bias or actual target motion. It is enough to show why the difference cannot certify stationarity. [M-SEI-013]

The ambiguity is larger when target and supports can move differently. Each contact observes displacement relative to its own mounting. A geometry drawing should therefore use relative coordinates rather than treating the table as an unexplained absolute frame. No numerical correction is authorized by this simplified derivation.

A.2 Worst-case difference and an unresolved small effect

If errors e1 and e2 satisfy absolute bounds B1 and B2, then the magnitude of e1 - e2 is at most B1 + B2. With twenty micrometres for each, a measured residual has a forty-micrometre worst-case error bound under the supplied assumptions. This argument does not require independence or a probability distribution.

A residual of thirty micrometres is compatible with zero true difference within that account, but is also compatible with a nonzero difference. The correct conclusion is unresolved at the stated bound. A two-hundred-micrometre residual is incompatible with zero under the same complete and valid bounds. A missing bias term invalidates that conclusion, which is why supplied calibration faults matter to the model. [M-SEI-002]

The derivation is a bound on a difference under stated inputs. It does not measure a coverage probability or establish that actual errors lie inside the assumed bounds. The complete-budget requirement is therefore part of the result.

A.3 Covariance changes a difference

For Y = X1 - X2, variance expansion gives Var(Y) = Var(X1) + Var(X2) - 2 Cov(X1,X2). Writing covariance as rho times u1 times u2 gives uc = sqrt(u1^2 + u2^2 - 2 rho u1 u2). With u1 = 4 and u2 = 3 micrometres, the expression is sqrt(25 - 24 rho). It gives 7, 5 and 1 at rho = -1, 0 and +1. [M-SEI-010; E1]

The limiting examples are useful checks on sign. Positive correlation can reduce uncertainty in a difference because the shared component cancels. The same cancellation can hide a common displacement relevant to an absolute claim. A small uc for the defined difference is not proof that a target is stationary or that the absolute channels are accurate. The measurand and the validity of the model remain decisive.

The formula must not be applied to planning allowances as though they were standard deviations. The P1 conservative sum and this illustrative standard-uncertainty calculation answer different questions and retain different assumptions.

A.4 An unsampled interior can move

Let a calibration curve be correct at x = 0 and x = 1000 micrometres. Those two observations alone do not fix its value at x = 200. A constructed error function can vanish at both endpoints while taking thirty micrometres at the interior point. The endpoint-only clean and bowed histories then produce identical observations at the sampled positions. [M-SEI-006]

Adding the two-hundred-micrometre position exposes the specified bow. It does not rule out a different error function that vanishes at all newly sampled positions and changes between them. A full sampled grid improves coverage; a bound between grid points needs a justified covering model, denser observations or another constraint on possible behavior.

The same logic applies to clock witnesses. Endpoint agreement cannot establish every intervening path. The report uses this analogy to identify a shared sampling problem, not to transfer a material calibration tolerance into a timing tolerance.

A.5 Return before and after a tare

The proposed raw return calculation is R = (g_return - g_initial) - (r_return - r_initial), with g for gauge and r for reference. A display tare changes presentation. If raw values are preserved, the calculation does not change. If the initial value is lost, the calculation is unavailable rather than zero. [M-SEI-007]

For g = (0,200,12) and r = (0,0,0), the retained return is twelve micrometres whether the display reads twelve or has been set to zero. For g = (0,212,12) and r = (0,12,12), the calculated return is zero. That second case can represent a common sensor shift or a coupled reference following a real simulated residual. The arrays alone do not distinguish the histories.

The six-micrometre allowance can be applied only under the reference and acquisition conditions assumed by the procedure. The two-micrometre manual target offset on January 6 is a different observation and cannot be substituted into this return test.

A.6 Offset correction and a hidden excursion

The supplied linear clock relation through reference/local points (0,7) and (1000,1017) is local = 7 + 1.01 reference, in milliseconds. Its inverse maps local 512 to reference 500. Subtracting only seven would give 505 and leave the linear rate error uncorrected. [M-LUM-003]

Now add a local twenty-millisecond excursion at the midpoint without changing either endpoint. The local midpoint becomes 532. Applying the same inverse gives approximately 519.801980, though the authored reference time remains 500. A midpoint witness can expose that discrepancy in the supplied case. It does not establish a bound on all excursions between witnesses. [M-LUM-004]

The arithmetic uses a specified clock path. It estimates neither real oscillator stability nor communication latency. A receiver prototype would require characterization and an explicit rule for what interval the observed witnesses support.

A.7 Atomic consumption and a lost reply

Suppose two callers both read challenge C as unused before either writes consumed. If admission follows that read, both can succeed even though the final stored value is consumed. Serializing a combined check-and-consume makes only the first transition succeed under the model's assumptions. [M-LUM-010]

There remains an interval after the consumed decision is committed and before a reply is delivered. A cut in that interval leaves no successful reply at the caller, yet the committed event has occurred. The correct recovery operation is to locate that decision by its request and payload binding. It must not restore the challenge to unused merely because the reply was lost. [M-LUM-011]

This argument assumes a valid original commit and sufficient surviving evidence to recover it. Where the decision or witness record is incomplete, reconciliation remains a hold. Atomicity does not create current-state authenticity or physical storage durability.

A.8 Authentic answers can be old answers

The Cycle 013 reply binds realm, nonce, request, payload and epoch. Verification checks its tag and active values before the retained recovery rules. Modifying a signed epoch fails the local authentication check; signing a different nonce with the correct key passes authenticity but fails active binding. [M-LUM-013]

After a verifier reset, reuse of the old active nonce and the same request values makes the recorded reply applicable again to that local comparison. The response bytes have not become less authentic. The missing property is retained evidence that the question must not be asked in the same supposedly fresh form.

Durable non-reuse, protected current-epoch custody and authenticated request binding therefore address different parts of the problem. The public fixture key and fixed strings test control flow only. They should never be copied into an operational service as credentials or a nonce-generation design.

A.9 A residual that still contains the command

Under y(c,s) = 3c + 2s + 0.5cs, the command-only response y(c,0) is 3c and source-only response y(0,s) is 2s. Their sum misses 0.5cs in the combined condition. Subtracting command-only from combined yields 2s + 0.5cs, which remains command dependent. [M-ALE-012]

At c = s = 1 the solo sum is five and combined output is five and a half. Reversing one input changes the interaction sign. With c = 0 the interaction vanishes even though the same response law applies. A zero residual in one control therefore does not show that interaction is absent everywhere.

These dimensionless calculations motivate multiple amplitudes, reversals and combined conditions. They do not establish an electrical transfer function, biological contribution or a permissible subtraction method for acquired data.

A.10 Order can reproduce the interaction contrast

For binary conditions define C = y11 - y10 - y01 + y00. In a static additive law C is zero; adding kcs gives C = k. Under the forward order 00,10,11,01 at times 0,1,2,3, a linear drift dt contributes d(2 - 1 - 3 + 0) = -2d. At d = 0.25 this is -0.5. Reversing the order gives +0.5. [M-ALE-013]

In the forward sequence followed by its reverse, each condition's two acquisition times average 3.5. The linear term cancels in the contrast. Their mean squared times are 24.5 for 00, 18.5 for 10, 14.5 for 11 and 12.5 for 01. The quadratic contribution is q(14.5 - 18.5 - 12.5 + 24.5) = 8q. With q = 0.125 it is one.

The paired order therefore yields one without interaction and one and a half with k = 0.5 under the supplied quadratic drift. Balanced mean time is not balanced mean squared time. Cancellation is a property of a specified design and nuisance model, not a universal consequence of reversal.

B. Proposed unloaded metrology protocol

B.1 Purpose and entry conditions

This protocol is a proposal for the next geometry review and unloaded comparison. It records no new session and authorizes no loaded strip test. Entry requires identified equipment actually available for the session, named custody, a retained mounting drawing, usable range, relevant calibration evidence and environmental observations. A renewed comparator loan requires a new agreement; NQ-17 is not presumed present.

The primary question is which relative motions are observed by the proposed reference arrangement. The initial output is a channel-to-motion map. The map should identify a target movement, a head or stand movement, and a support movement that might be shared. Where the apparatus cannot separately excite those motions safely and measurably, the limitation should be explicit. A diagram alone is not evidence that the separation exists.

B.2 Retained records

Retain equipment identifiers, mounting orientation, contact locations, cable routing, transport and setup changes, operator, custodian and observation times. Record raw starting displays before any tare. Every later tare must have a time, reason and preserved pre-tare values. Record environmental quantities only where instruments and their resolution are identified; an unmeasured room condition must not acquire an invented numerical bound.

The session table should contain both channels, the common target condition, any commanded or manual action, and explicit missing readings. Display increments and nominal divisions should be distinguished from calibrated displacement. Preserve original units. A second table may present changes from each start, but it must retain the link to the raw values.

B.3 Sequence and challenges

Begin with the retained unloaded configuration and a no-command observation interval. The purpose is to characterize what the arrangement reports over that session, not to establish a universal warm-up time. If a controlled movement is introduced, identify which object moved and whether the input is calibrated. A hand movement can challenge an inference but cannot supply an absolute calibration step.

Where supported by safe apparatus limits and an independently justified reference, compare outward and return positions at the retained grid. Do not infer full-grid coverage from a single step. Record approach direction and repeated positions. Stop and retain the incomplete sequence if usable travel, contact, reference support or record retention fails.

B.4 Interpretation and exit

A stable channel difference supports only the defined comparison under the observed geometry. If a known shared movement preserves the difference, report that directly. If the reference itself moves with the intended target, no corrected specimen return can be claimed from the pair alone. Missing calibration or mounting terms keep acceptance incomplete rather than producing a smaller budget.

The review should distinguish a usable intake record, a supported unloaded comparison, completed calibration coverage and loading acceptance. These are separate exits. The present twenty-five-micrometre budget, six-micrometre return allowance and thousand-micrometre grid remain unchanged unless a separately reviewed research decision explicitly revises the requirement. No stress or strength output follows without additional force, geometry and material evidence.

C. Proposed receiver recovery protocol

C.1 Purpose and state boundaries

This proposed implementation test concerns preservation and recovery of a request-bound decision. It is not an operational lunar service. Before testing, specify the issuer, witness, verifier, key holder and current-epoch authority, and identify the storage and reset boundaries each can share. A service labeled outside the receiver is insufficient if restoring the receiver can restore that service's relevant history too.

Document the expected behavior when current authority is unavailable. A hold must preserve the retained request and payload rather than silently issuing a substitute decision. Key provisioning and protection require their own implementation review; the published fixture key is demonstration data only.

C.2 Required observations

For each attempt retain active realm, nonce, request and payload binding; supplied issuer and witness state; authenticated reply fields; verification result; session closure; consumption count; commit and reply completion evidence; and recovery action. A digest may be used in a later implementation, but the exact bytes and canonicalization rules must be specified rather than assumed from a field name.

Keep a distinction between the system's inputs and the test controller's complete intervention history. If an old snapshot is deliberately restored, the reviewer should not receive an extra rollback flag unless the proposed service actually observes such a flag. The retained controller history then lets the test report the blind spot honestly.

C.3 Challenge set

Test a genuine first use, duplicate caller, split check-and-consume, cut before commit, cut after commit before reply, repeated recovery read, changed request, changed payload, missing witness and withdrawn authority. Test an authenticated reply to a different active question, altered authenticated content, verifier reset with nonce reuse, receiver-only rollback, issuer/witness rollback and joint rollback of all authorities.

Each cut needs observable evidence of which transition completed. A software flag named flushed does not establish a physical flush. Where the test runs only in a model, report the result at that level. Recovery reads must retain zero new consumptions and must not fill missing time anchors.

C.4 Review criteria

The output should distinguish new admission from retrieval of an existing decision, pending reconciliation from unavailable authority, and payload survival from timing permission. An old valid history and a fully rolled-back history with identical supplied state must remain an explicit indistinguishable pair unless the implementation adds a surviving discriminating observation.

The review should explain which failures are detected, which are prevented under specified assumptions, and which remain undetectable. Availability costs should be recorded alongside safety holds. A prototype that withholds every request has not demonstrated useful recovery, even though it avoids favorable false classifications.

D. Proposed electrical source and storage protocols

D.1 Two experiments with separate conclusions

The electrical surrogate experiment concerns source response. The storage experiment concerns retained completion. They should use separate result headings and acceptance criteria even if a later apparatus combines them. Neither proposal involves a participant, implant, neural recording or stimulation.

For the electrical surrogate, identify the source, command path, acquisition input, range, resolution, timebase, bandwidth and any filtering. Retain the exact generated input and observed raw output. Units must be defined before a dimensionless software coefficient is compared with an electrical amplitude. State which parts of the acquisition path can saturate, clip or share a ground or supply.

D.2 Conditions and order

Retain baseline, command-only, source-only and combined conditions. Use repeated acquisitions, controlled amplitude changes and input reversals where supported by the apparatus. Preserve the actual order and times. A planned order and the performed order are separate records if interruption changes the sequence.

Analyze the combined response against the sum of isolated responses, but preserve the discrepancy as a model residual rather than naming its source. Compare alternative lag ranges and record zero-energy conditions as undefined for cosine measures. The finite search interval should appear beside any selected lag. No numerical flag should imply that an unresolved signal is neural or externally originated.

Reversed or balanced order can challenge a linear drift explanation. It cannot establish the absence of nonlinear drift. Repeated baseline observations and additional order conditions should test whether the response changes with time or preceding input. Nuisance-factor control should be specified before viewing a favored residual, with deviations recorded afterward.

D.3 Storage completion under interruption

Bind one request to one retained value. Define local completion, witness completion, readback and acknowledgement as observable transitions. Place failure cuts before local completion, between local and witness completion, after witness completion, and before acknowledgement. Include loss of each supply and any shared supply path. Record the original value and every attempted recovery without drawing a replacement.

A queued write should remain distinct from a completed witness record. A readback must bind request and value, not merely state that some write occurred. If neither readable retained copy survives, recovery should preserve the unresolved request and its missing evidence. It should not create a new value under the old identity.

D.4 Review and limits

Report the storage behavior separately from the attribution result. Preserving a value establishes no independent origin. Recognizing a command echo establishes no storage durability. A response that remains after one cancellation method is only a remainder under that method until further evidence supports its source.

The next review should identify the smallest additional experiment that could contradict each leading explanation. If a proposed condition cannot change the distinction among those explanations, classify its role as repeatability, range checking or redundancy rather than discrimination. The goal is an acquisition path whose ordinary contributions can be inspected before the program asks it to support an extraordinary claim.

E. Selected software cases and online supplement

Eight representative entries connect the argument to the retained software constructions. The complete 337-entry register is available in the online supplement, grouped by cycle with expandable data views. Constructed histories are test definitions, not additional observations. Authored construction and reviewer observations remain distinct. No key here is an operational credential.

meridianrc.org/working-papers/001/supplement/

Cycle 002 / shared calibration error

A shared calibration error survives agreement.

extension um: 400

slip um: 200

gauge bias um: 0

crosshead bias um: 0

calibration: shared_unverified

shared scale percent: 10

gauge um: 440

crosshead um: 660

residual um: 220

difference bound um: 40

naive discrepancy flag: True

decision: withhold_calibration_unverified

Full entry c002-005 in the supplement

Cycle 002 / distinct labels shared upstream

Different labels can conceal a shared upstream origin.

inputs: [{"channel": "input", "origin_tag": "a", "acquisition_context": "s1", "tag_context": "s1", "upstream_root": "s1/root-a"}, {"channel": "feedback", "origin_tag": "b", "acquisition_context": "s1", "tag_context": "s1", "upstream_root": "s1/root-a"}]

nominal channels: 2

unique nonmissing labels: 2

verified origins: 1

decision: shared_origin

authored independent origins: 1

withheld independent pair: False

Full entry c002-019 in the supplement

Cycle 009 / cloned issuer state

An issuer restored with its witness cannot expose their common rollback.

reviewer input: {"sample": {"session": "S3", "tick": 50, "challenge": "C41"}, "journal": [{"session": "S3", "sequence": 1, "tick": 0, "challenge": "C41"}, {"session": "S3", "sequence": 2, "tick": 100, "challenge": "C41"}], "issuer": {"available": true, "validated": true, "state": "issued_unused", "session": "S3", "challenge": "C41"}}

result: conditional_model_bracket

Authored construction (not a reviewer observation): All three inputs are cloned together from an earlier state; issuer independence is false but unavailable to the reviewer.

Full entry c009-031 in the supplement

Cycle 013 / current

A valid current response recovers the retained decision.

inputs: {"reply": {"body": {"realm": "lumen-recovery", "nonce": "fresh-session-013", "request": "Q9", "payload": "block-A", "epoch": 8}, "tag": "9d1a89f09325fe8ae597e35361f66a75572ae3cb3f30ee83c759e3b7e41f9290"}, "active": {"realm": "lumen-recovery", "nonce": "fresh-session-013", "request": "Q9", "payload": "block-A", "closed": false}, "issuer": {"epoch": 8, "challenge": "C42", "state": "consumed", "request": "Q9", "payload": "block-A"}, "witness": {"epoch": 8, "challenge": "C42", "state": "consumed", "request": "Q9", "payload": "block-A"}, "authority": "permitted"}

decision: recover_bound_current_model_decision

session closed: True

new consumptions: 0

Full entry c013-001 in the supplement

Cycle 013 / all domains rollback

Authenticity alone does not establish presentness.

inputs: {"reply": {"body": {"realm": "lumen-recovery", "nonce": "fresh-session-013", "request": "Q9", "payload": "block-A", "epoch": 7}, "tag": "c2fb232bf0e40cc5809302fd62aba4189884b64e24d3f6dac2029d021e49edd3"}, "active": {"realm": "lumen-recovery", "nonce": "fresh-session-013", "request": "Q9", "payload": "block-A", "closed": false}, "issuer": {"epoch": 7, "challenge": "C42", "state": "consumed", "request": "Q9", "payload": "block-A"}, "witness": {"epoch": 7, "challenge": "C42", "state": "consumed", "request": "Q9", "payload": "block-A"}, "authority": "permitted"}

decision: recover_bound_current_model_decision

session closed: True

new consumptions: 0

Full entry c013-011 in the supplement

Cycle 013 / reset reuses nonce

A reused exchange value leaves a recovery boundary open.

inputs: {"reply": {"body": {"realm": "lumen-recovery", "nonce": "fresh-session-013", "request": "Q9", "payload": "block-A", "epoch": 8}, "tag": "9d1a89f09325fe8ae597e35361f66a75572ae3cb3f30ee83c759e3b7e41f9290"}, "active": {"realm": "lumen-recovery", "nonce": "fresh-session-013", "request": "Q9", "payload": "block-A", "closed": false}, "issuer": {"epoch": 8, "challenge": "C42", "state": "consumed", "request": "Q9", "payload": "block-A"}, "witness": {"epoch": 8, "challenge": "C42", "state": "consumed", "request": "Q9", "payload": "block-A"}, "authority": "permitted"}

decision: recover_bound_current_model_decision

session closed: True

new consumptions: 0

Full entry c013-015 in the supplement

Cycle 013 / linear forward

Acquisition order can imitate an interaction.

interaction: 0

drift per step: 0.25

quadratic coefficient: 0

rows: [[0, 0, 0, 0.0], [1, 1, 0, 3.25], [2, 1, 1, 5.5], [3, 0, 1, 2.75]]

contrast: -0.5

Full entry c013-018 in the supplement

Cycle 013 / balanced quadratic

Balancing linear drift does not remove every time-dependent alternative.

interaction: 0

drift per step: 0

quadratic coefficient: 0.125

rows: [[0, 0, 0, 0.0], [1, 1, 0, 3.125], [2, 1, 1, 5.5], [3, 0, 1, 3.125], [4, 0, 1, 4.0], [5, 1, 1, 8.125], [6, 1, 0, 7.5], [7, 0, 0, 6.125]]

contrast: 1.0

Full entry c013-022 in the supplement

F. Sources and document catalogue

External technical sources

E1 / Joint Committee for Guides in Metrology. JCGM 100:2008. Evaluation of measurement data - Guide to the expression of uncertainty in measurement. Section 5.2: correlated input quantities. DOI: 10.59161/JCGM100-2008E.

https://www.bipm.org/documents/20126/2071204/JCGM_100_2008_E.pdf

E2 / H. Krawczyk, M. Bellare and R. Canetti. HMAC: Keyed-Hashing for Message Authentication. RFC 2104, February 1997. Background for the authentication concept only; the Meridian toy exchange is not a security validation.

https://www.rfc-editor.org/rfc/rfc2104.html

E3 / NIST/SEMATECH e-Handbook of Statistical Methods. Section 5.3.3.2: Randomized block designs. Background for nuisance-factor control; the ordered Meridian examples remain deterministic calculations.

https://www.itl.nist.gov/div898/handbook/pri/section3/pri332.htm

External sources provide background for the identified methods. They do not validate Meridian equipment or findings. Historical note references to other sources remain part of those notes; this draft does not claim a new reproduction of the referenced external experiments.

Meridian document catalogue

M-SEI-001 / Where the displacement comes from

2026-10-07 / Mara Voss, Jonas Reed

https://meridianrc.org/research/cycle-001/space-elevator/

M-LUM-001 / Keeping the events a correction discards

2026-10-07 / Elias Chen, Leila Haddad

https://meridianrc.org/research/cycle-001/lumen/

M-ALE-001 / Two channels, one source

2026-10-07 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-001/aletheia/

M-SEI-002 / When a difference is not yet a displacement finding

2026-10-12 / Mara Voss, Jonas Reed

https://meridianrc.org/research/cycle-002/space-elevator/

M-LUM-002 / The event neither archive can return

2026-10-12 / Elias Chen, Leila Haddad

https://meridianrc.org/research/cycle-002/lumen/

M-ALE-002 / The cost of withholding corroboration

2026-10-12 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-002/aletheia/

M-SEI-003 / Which part of the strip moved?

2026-10-19 / Mara Voss, Jonas Reed

https://meridianrc.org/research/cycle-003/space-elevator/

M-LUM-003 / Two copies of event five

2026-10-19 / Elias Chen, Leila Haddad

https://meridianrc.org/research/cycle-003/lumen/

M-ALE-003 / The circle that came back twice

2026-10-19 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-003/aletheia/

M-SEI-004 / The bench that is still on paper

2026-10-27 / Mara Voss, Jonas Reed

https://meridianrc.org/research/cycle-004/space-elevator/

M-LUM-004 / Before the receiver makes its second copy

2026-10-27 / Elias Chen, Leila Haddad

https://meridianrc.org/research/cycle-004/lumen/

M-ALE-004 / A valid signature and an unfinished ledger

2026-10-27 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-004/aletheia/

M-SEI-005 / Choosing the instrument before choosing the strip

2026-11-05 / Mara Voss, Jonas Reed

https://meridianrc.org/research/cycle-005/space-elevator/

M-LUM-005 / The silence between clock witnesses

2026-11-05 / Elias Chen, Leila Haddad

https://meridianrc.org/research/cycle-005/lumen/

M-ALE-005 / When the withdrawal reaches the ledger

2026-11-05 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-005/aletheia/

M-SEI-006 / The step between zero and one millimetre

2026-11-13 / Mara Voss, Jonas Reed

https://meridianrc.org/research/cycle-006/space-elevator/

M-LUM-006 / The timestamp that the receiver must leave open

2026-11-13 / Elias Chen, Leila Haddad

https://meridianrc.org/research/cycle-006/lumen/

M-ALE-006 / A third account of the draw

2026-11-13 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-006/aletheia/

M-SEI-007 / The zero after unloading

2026-11-21 / Mara Voss, Jonas Reed

https://meridianrc.org/research/cycle-007/space-elevator/

M-LUM-007 / The receiver starts counting again

2026-11-21 / Elias Chen, Leila Haddad

https://meridianrc.org/research/cycle-007/lumen/

M-ALE-007 / The request that must not draw twice

2026-11-21 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-007/aletheia/

M-SEI-008 / The travel the head cannot use

2026-11-28 / Mara Voss, Jonas Reed

https://meridianrc.org/research/cycle-008/space-elevator/

M-LUM-008 / The last witness record

2026-11-28 / Elias Chen, Leila Haddad

https://meridianrc.org/research/cycle-008/lumen/

M-ALE-008 / A copy beyond the missing tail

2026-11-28 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-008/aletheia/

M-SEI-009 / The bracket on the certificate

2026-12-07 / Jonas Reed, Mara Voss

https://meridianrc.org/research/cycle-009/space-elevator/

M-LUM-009 / A challenge outside the receiver

2026-12-07 / Leila Haddad, Elias Chen

https://meridianrc.org/research/cycle-009/lumen/

M-ALE-009 / Before the acknowledgement leaves

2026-12-07 / Tomás Ibarra, Nadia Okafor

https://meridianrc.org/research/cycle-009/aletheia/

M-SEI-010 / Borrowing the whole instrument

2026-12-16 / Jonas Reed, Mara Voss

https://meridianrc.org/research/cycle-010/space-elevator/

M-LUM-010 / The challenge that stays used

2026-12-16 / Leila Haddad, Elias Chen

https://meridianrc.org/research/cycle-010/lumen/

M-ALE-010 / A receipt cannot identify a signal

2026-12-16 / Tomás Ibarra, Nadia Okafor

https://meridianrc.org/research/cycle-010/aletheia/

M-SEI-011 / The first hour on the table

2026-12-24 / Jonas Reed, Mara Voss

https://meridianrc.org/research/cycle-011/space-elevator/

M-LUM-011 / Recovering the reply, not the challenge

2026-12-24 / Leila Haddad, Elias Chen

https://meridianrc.org/research/cycle-011/lumen/

M-ALE-011 / The echo one sample later

2026-12-24 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-011/aletheia/

M-SEI-012 / Two readings of the same stillness

2027-01-01 / Jonas Reed, Mara Voss

https://meridianrc.org/research/cycle-012/space-elevator/

M-LUM-012 / Who keeps the current epoch?

2027-01-01 / Leila Haddad, Elias Chen

https://meridianrc.org/research/cycle-012/lumen/

M-ALE-012 / When the inputs stop adding

2027-01-01 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-012/aletheia/

M-SEI-013 / The last hour of the loan

2027-01-08 / Jonas Reed, Mara Voss

https://meridianrc.org/research/cycle-013/space-elevator/

M-LUM-013 / A reply to this question

2027-01-08 / Leila Haddad, Elias Chen

https://meridianrc.org/research/cycle-013/lumen/

M-ALE-013 / The order in the signal

2027-01-08 / Nadia Okafor, Tomás Ibarra

https://meridianrc.org/research/cycle-013/aletheia/

P-SEI / Space Elevator: exogenous environment and interaction

https://meridianrc.org/projects/space-elevator/#anomalous

P-LUM / Lumen: exogenous interaction and contact research

https://meridianrc.org/projects/lumen/#exogenous

P-ALE / Aletheia: exogenous cognition

https://meridianrc.org/projects/aletheia/#anomalous

LUM-ACK-1031 / Uncommanded acknowledgement

https://meridianrc.org/archive/uncommanded-acknowledgement/

G. Terms used in this report

ET/ED preparedness

Preparedness and integration of exogenous hypotheses into research design, instrumentation and governance. Lumen retains level 4, Aletheia level 5; the Space Elevator numerical level is unspecified. A level is not a probability or a finding about an event.

Retained exception

A record preserved with its unresolved provenance and competing explanations. LUM-ACK-1031 is a receiver-log transcription without the original capture or certified timing witness; the reported shutdown-field discrepancy remains part of that unverified account.

Acquisition

The event or process that produces a retained input. A later copy is not a new independent acquisition.

Active binding

The fields connecting an authority response to the question currently being verified, including realm, nonce, request and payload in the toy exchange.

Admission

The decision that first uses a challenge under the issuer rule. A recovery read retrieves an existing decision and must not create a second admission.

Agreement

A specified relationship between observed records, such as equal payloads or a small channel difference. Its meaning depends on the comparison and failure model.

Allowance

A selected design or acceptance limit. The 25 um displacement budget and 6 um return allowance are not estimated standard uncertainties.

Archive arrival

Receipt at an archive, distinct from the time of source acquisition and affected by transport history.

Atomic consumption

A combined check-and-state transition preventing two callers from both admitting against the same unused challenge under the model assumptions.

Authentication

Verification of message origin or integrity under a specified mechanism and key assumptions. It does not establish the truth or current applicability of the assertion.

Calibration coverage

The range, configuration, conditions and evidence to which a calibration account applies. Head-only paperwork does not automatically cover a new mounting.

Common mode

A contribution shared by channels. It may cancel in their difference while remaining important to an absolute claim.

Conditional result

A result valid under explicitly supplied assumptions, such as separate reference custody or a current authority state.

Confounding

A situation in which alternative factors can produce the same observed contrast or response. The ordered drift examples are explicit constructed cases.

Corroboration

Additional support from a justified distinct origin. Multiple paths or labels alone do not establish that distinction.

Covariance

A joint-variation term in a statistical input model. Its sign affects uncertainty propagation for a difference.

Custody

Responsibility for possession, retention and permitted use of a record or object. Different custody may support some separations without proving all relevant independence.

Decision history

The retained sequence of decisions, evidence and reasons. A later reconsideration appends a result rather than deleting the earlier account.

Digest

A computed representation used to compare retained content. A digest recomputed after the same edit cannot independently witness the earlier record.

Disputed event

An event with incompatible retained accounts. It differs from a missing event and from a surviving event without its required witness.

Draw

One authored source acquisition in the symbol models. Recovery of its retained value is not permission to issue another draw.

Epoch

A scoped state identifier used in the receiver recovery models. Its number does not itself prove that a supplied record is current.

Failure domain

The boundary within which one interruption, restoration or change can affect several supposed witnesses together.

First-record witness

A retained observation or digest originating before the processing change it is intended to expose. Its actual capture stage and custody matter.

Flush

A modeled or implemented completion boundary for a write. The software fixtures assume its survival properties rather than measuring physical storage.

Gauge extension

Change in distance between specimen gauge marks, distinct from crosshead travel or grip movement.

Hidden history

A constructed case fact not supplied to the reviewer. It is used to assess the reviewer limit and must not leak into its decision input.

Indistinguishable pair

Two candidate histories that give the reviewer identical observations while differing in a property relevant to the desired claim.

Interaction

A joint input contribution not predicted by adding isolated responses under the selected model. Time-dependent drift can imitate a contrast attributed to interaction.

Lag

A shift in samples in the finite synthetic source screen. Without a sampling rate it has no assigned physical duration.

Measurand

The quantity intended to be measured. A difference between channels and an absolute target displacement are different quantities.

Missing event

An expected event whose retained payload is unavailable. The reason for absence may itself remain unknown.

Nonce

A value intended to distinguish an active exchange. A fixed label or a reused value after reset does not demonstrate durable non-reuse.

Planning bound

An assumed allowance used to examine a requirements envelope. It is not an empirical instrument performance estimate.

Raw return

The return computed from retained initial and final gauge and reference readings, preserving the effect hidden by a display tare.

Recovery read

Retrieval of an already bound decision or acquired value. It should retain its original identity and avoid new consumption or replacement acquisition.

Reference geometry

The relation among contacts, targets, supports and frames that determines which motions a comparison actually observes.

Residual

The difference left by a stated comparison or fitted model. A residual has no automatic source identity.

Rollback

Restoration of an earlier state. If every available authority shares the restoration, their agreement can fail to reveal the later history.

Source attribution

An inference about where an acquired signal originated. Preserving, authenticating or replaying the signal does not alone establish that origin.

Standard uncertainty

An uncertainty expressed as a standard deviation within the specified measurement model, distinct from a worst-case allowance.

Tare

An adjustment of display origin. Retaining the raw readings permits a later review of what the adjustment would otherwise hide.

Timing permission

The supported use of a retained event for a timing claim. Payload service can remain useful while this permission is withheld.

Unwitnessed event

A surviving payload without its required first-record support. It is not automatically missing or known to be false.

Usable range

Travel actually available in the offered instrument configuration, which can be smaller than its nominal range.

Working draft

A document prepared for technical review. Proposed methods and synthesis can change without erasing the identified historical records.

H. Open review register

These are unresolved technical questions for review, not completed actions or newly assigned institutional approvals.

Metrology geometry

Which relative target, support and head motions can the next arrangement distinguish? Identify a control capable of breaking channel agreement under the competing explanation.

Calibration scope

Which range and configuration are actually covered? Retain missing mounting, transport, alignment and environmental terms.

Return acceptance

Can the initial and returned raw values be recovered across every tare and restart? Do not substitute the manual target offset for a strip return.

Time bracketing

Which clock paths remain possible between witnesses? Identify the evidence supporting any claimed bound.

Freshness custody

What later observation survives the rollback under test? Identify key, epoch and non-reuse failure domains separately.

Commit boundary

What observable event establishes local and witness completion? Retain the lost-reply interval without a second admission or draw.

Signal attribution

Which command, lag, amplitude, interaction and drift alternatives survive the proposed electrical controls? Name residuals by their models.

ET/ED record discrimination

For LUM-ACK-1031, can the original capture, independent timing and originating shutdown exporter be recovered? Which records could discriminate replay, stale-buffer reuse, transcription fault and an unfamiliar source without granting the content control authority?

Reconsideration

Can later evidence both support and undermine a withheld claim while leaving the earlier decision intact?

Subject index

Page references locate the terms in the body, documentary record and appendices. Ranges include contiguous pages; the index does not claim that every mention is a distinct finding.

Acquisition order

16, 30, 117-118, 147

Agreement

9-11, 21, 24, 45, 54, 58, 61-62, 71-73, 76, 80, 86-88, 96, 101-102, 107-109, 111-113, 136, 139, 145, 155, 159-160

Aletheia

7, 9, 11, 13, 16, 28, 31, 34, 36, 92-93, 95-100, 102-118, 130, 132, 149-155

Authentication

16, 25, 29, 36, 81-82, 86-91, 94, 96-99, 101-107, 137-138, 141, 149, 156

Calibration

11, 15, 19, 24, 33, 35, 40-43, 46-48, 50, 52-55, 60-61, 63, 67-68, 116, 122, 127, 135-136, 139-140, 145, 156, 160

Challenge consumption

25, 36, 82-88, 90, 137, 141, 146-147, 156, 158

Common target

11, 18, 139

Covariance

19, 56, 135, 156

ET/ED preparedness

7, 9, 12-13, 21, 31, 37, 119, 124, 130, 133-134, 155, 161

Exogenous hypotheses

7, 9, 13, 21, 95, 104-107, 120-126, 129-131, 134, 154-155

LUM-ACK-1031

7, 12-14, 26-27, 34, 37, 133-134, 154-155, 161

Uncommanded acknowledgement

7, 9, 26, 133-134, 154

Drift

16, 30, 43, 50, 58-59, 68, 70, 117-118, 127, 138, 143, 147-148, 156, 158, 160

Epoch

6, 15-16, 24-25, 33-34, 36, 78-79, 86-91, 137-138, 141, 146-147, 153, 157, 160

First-record witness

11, 23-24, 71-74, 157, 159

Grip motion

18, 42-43, 48, 50-52, 157

HMAC

25, 90, 100-104, 149

Indistinguishable histories

15-16, 36, 48, 71, 73, 76, 78, 84, 86, 88, 98-99, 103, 109, 136, 142, 157

Interaction

7, 13, 21, 30, 115-118, 120-123, 125-126, 128, 130, 138, 147-148, 154, 158, 160

KIS-61

10, 15, 18, 20, 36, 52-56, 58, 60-63

Lumen

6-7, 9, 11, 13, 15, 23, 26, 34, 36, 65-73, 75-91, 125-127, 129, 133, 146-147, 149-155

Missing records

7, 15-16, 19, 23-26, 28-29, 33-34, 36-37, 47-49, 51, 67, 69-76, 80-89, 91, 93-94, 96-100, 102-107, 109, 111-112, 135, 137, 139-141, 144, 146, 152, 157-160

NQ-17

10, 15, 18, 20, 36, 60-63, 139

Nonce reuse

25, 29, 90-91, 96, 137-138, 141, 146-147, 155, 158

North Quay

10-11, 35, 60, 62

Power interruption

15, 29-30, 32-34, 36, 51, 58-61, 63, 73, 75, 77, 105, 109-110, 113-117, 120, 125, 127, 129, 131

Raw return

136, 158

Recovery

8-9, 11, 16, 23-25, 27, 29, 33-34, 36, 67-69, 74-75, 78, 80-84, 86-90, 105-111, 137, 141-144, 146-147, 155, 157-158

Reference geometry

11, 18, 21, 34-36, 47, 60, 120-121, 123, 125, 135, 139-140, 158, 160

Residual

19, 36, 40, 48-51, 62, 71, 73-78, 115, 135-136, 138, 143, 145, 158, 160

Rollback

15, 36, 84-86, 88-91, 141, 146-147, 159-160

Source attribution

9, 11, 14, 21, 28-29, 31, 37, 96, 144, 159-160

Storage completion

29, 32, 34, 36, 109-110, 113, 115, 117-118, 141, 143, 157, 160

Tare

20, 32, 50-51, 58-60, 62, 136, 139, 158-160

Timing permission

12-14, 16, 24, 26-27, 30, 32, 34, 37, 41, 44, 67-69, 71-86, 89-90, 99, 103, 111, 117, 123, 125-126, 128-129, 133-134, 136, 142, 155, 159, 161