



QILAHK

SECURITY YOU CAN PROVE

ADVANCED RESEARCH WHITE PAPER

QTA-R02 / 02 OF 05

Temporal Trust Continuity

Research into material change time boundaries and reproducible reassessment

Version 1.1 | October 12, 2026

QILAHK Trust Authority

LOCK / KEY / PROOF

qilahk.com | blog.qilahk.com

Table of contents

Temporal Trust Continuity

Abstract	3
Time semantics and the research problem	4
A decision has several clocks	4
Two historical questions	4
Existing foundations	4
A proposed transition and materiality model	5
Separate appraisal from lifecycle action	5
Define materiality as a contract	5
An illustrative delayed evidence scenario	6
From initial state to reassessment	6
Bound what continuous means	6
Threat model and resilience	6
Evaluation of selective reassessment	7
Compare against full reassessment	7
Hypotheses and counterexamples	7
Prototype gates and references	8
A staged research implementation	8
Open questions	8
References	8

Abstract

Temporal Trust Continuity

Trust continuity is a problem of justified reliance through time. A valid observation can become stale, a policy can change, and events can arrive after an earlier decision has been recorded. This paper proposes a QILAHK research architecture that retains observation time and ingestion time separately, pins the basis of each determination, and evaluates material changes against explicit claim dependencies.

The proposed mechanism combines an event history, versioned materiality rules, bounded freshness, and an authority-controlled transition model. It aims to reconstruct what was known at a decision time while also supporting a corrected view when delayed evidence appears. A materiality result initiates review; it does not by itself authorize a certificate, registry, or TrustMark action.

The research contribution to test is whether selective reassessment can reduce redundant reviews while preserving detection of changes that invalidate required conditions. The evaluation must include collector outages, clock uncertainty, out-of-order events, policy revisions, and concurrent authority actions. No uninterrupted trust guarantee or measured latency improvement is asserted.

Intended audience

Distributed systems researchers, security engineers, and policy designers

Research scope

This white paper presents proposed QILAHK research mechanisms, illustrative scenarios, and evaluation protocols. Its hypotheses are intended for prototype testing and independent review.

QTA-R02 | Version 1.1 | October 12, 2026

Time semantics and the research problem

A decision has several clocks

An observation can occur at one time, reach the service later, and support a decision later still. Treating these times as interchangeable makes historical reconstruction unreliable. The proposed event contract records observation time, ingestion time, processing time, and decision effective time, together with clock source and uncertainty where available.

An observed state is evidence about the observation interval. It is not proof that the state remained unchanged until ingestion. The relying policy must specify acceptable evidence age and any events that terminate reliance earlier. A scheduled collection job is not equivalent to a successful, qualified observation.

Two historical questions

An operational review may ask what the authority knew when it decided. A retrospective investigation may ask what the evidence now indicates about that earlier period. Keep both views. The first uses events available at the decision cutoff; the second can incorporate delayed evidence but must label the correction and retain the original determination.

The decision contract pins the subject, claim, evidence set, policy version, authority context, and effective interval. Materiality checks compare qualified observations against the dependencies in that contract. A device identity change can be material even when a measured value is identical, because the binding that justified the original decision has changed.

Existing foundations

NIST SP 800-137 describes an information security continuous monitoring strategy and program [1]. This proposal narrows the research problem to decision continuity and temporal reconstruction. It does not equate a materiality engine with a complete monitoring program. Collector coverage, escalation, and response remain separate operational requirements.

A proposed transition and materiality model

Separate appraisal from lifecycle action

Maintain a current determination and a review state as distinct records. Opening a review identifies uncertainty or an affected dependency; it need not rewrite the earlier determination. A relying party receives the current determination together with any unresolved condition relevant to its policy. The policy specifies whether pending review limits reliance.

Each lifecycle transition is an event with an expected prior version, authorized actor, rationale, evidence references, and effective time. Use a concurrency check to prevent two reviewers from silently overwriting each other. An idempotency key makes a repeated submission distinguishable from a new decision.

Define materiality as a contract

For each dependency, a rule declares the property compared, permitted measurement method, comparison semantics, evidence-age limit, uncertainty treatment, and resulting review requirement. Numeric tolerances need units and context. Categorical changes need an allowed-state set or an explicit transition rule. Missing values need their own treatment.

Selective reassessment follows affected dependency edges from a changed observation to claims and determinations. This can save work only when the dependency map is adequate. The system must retain a conservative fallback for changes whose affected claims cannot be bounded. Unknown relevance should initiate review rather than quietly preserve an established status.

Illustrative event handling rules

Event	Required treatment
Relevant value changes	Compare against the pinned materiality rule and open review if required
Observation becomes stale	Record the affected claim and apply the relying policy
Policy version changes	Assess whether existing determinations require migration or reassessment
Late event arrives	Preserve the prior decision and add a corrected retrospective view
Concurrent decision	Reject stale version submission or require explicit reconciliation

These rules are proposed prototype behavior. An authorized disposition can continue, condition, supersede, or end a determination according to policy. Artifact issuance and status propagation should occur through explicit workflow actions linked to that disposition, with failure recovery and audit records.

An illustrative delayed evidence scenario

From initial state to reassessment

At 09:00 a collector observes an approved configuration on device D17. At 09:02 the evidence reaches the service; at 09:05 an authority records a scoped determination. At 09:10 the device changes configuration, but the collector is disconnected. At 09:20 a relying party requests current status. Whether reliance is permitted depends on the declared freshness and outage policy, not on the mere existence of the 09:05 determination.

At 09:25 the delayed 09:10 observation arrives. The system records both times, opens a review, and identifies the claims dependent on the changed configuration. A retrospective view can now show that the property was affected at 09:10. It should not falsify history by pretending the authority possessed that evidence at 09:05 or 09:20.

Bound what continuous means

Continuity is bounded by observability. A collector may miss a transient change between observations. The architecture cannot establish uninterrupted validity from periodic snapshots alone. A claim about continuous state needs a collection and enforcement method capable of supporting that interval, plus explicit coverage and failure assumptions.

A change detector also needs to address measurement uncertainty. Compare intervals or classified states when a scalar difference is not reliable. Hysteresis can reduce oscillating review triggers, but it creates a tradeoff between noise tolerance and detection delay that must be documented and tested.

Threat model and resilience

The threat model includes replay, event suppression, collector compromise, clock manipulation, queue delay, and duplicate or reordered delivery. Signed timestamps do not make a compromised clock accurate. Source sequence numbers and freshness challenges can help under a declared collector model but cannot establish facts outside that model.

Retain raw events and corrections; reject unauthorized lifecycle transitions; reconcile materiality checks to pinned rules. During an outage, expose the unavailable dependency instead of converting lack of observations into a positive health signal. The response may reduce availability and must be chosen by the relying policy owner.

Evaluation of selective reassessment

Compare against full reassessment

Build a simulator with known state transitions, multiple collector schedules, controllable delays, clock offsets, and policy revisions. Compare full reassessment on every event, fixed-interval reassessment, and dependency-based selective reassessment. Keep the same event streams and authority rules across all variants.

Define a material event in the ground truth as one that changes whether a mandatory claim meets the specified criteria. Separately track events that are relevant but preserve the condition, such as a permitted change. This prevents the evaluation from labeling every observed difference a security failure.

Hypotheses and counterexamples

H1 proposes that selective reassessment reduces review count at matched recall for material events. H2 proposes that the event model reconstructs both the as-known and corrected histories deterministically despite reordering. Counterexamples include a missing dependency that suppresses a required review or a late event that changes the reconstructed as-known decision.

Run fault injection for dropped messages, stale caches, repeated submissions, and simultaneous decisions. Include an ablation that removes clock uncertainty handling. Report results by outage length and event rate; aggregate averages can conceal long periods of unjustified reliance.

Proposed continuity measures

Measure	Definition
Material-event recall	Detected ground-truth material events divided by all such events
Exposure interval	Time from invalidating event to enforcement of the governing reliance rule
Review reduction	Difference in review count at a matched detection operating point
Replay consistency	Replays reproducing the pinned as-known outcome divided by replay trials
Uncertainty duration	Elapsed time with insufficient qualified evidence for the current claim

Choose acceptable bounds before production evaluation. A lower review count is not a benefit if invalidation recall falls or exposure intervals grow beyond the policy limit. Report the assumptions that make each bound possible, including observation coverage and propagation delay.

Prototype gates and references

A staged research implementation

Begin with an append-only event interface and a replayable evaluator. Add subject identity versioning and a pinned rule contract. Next implement dependency traversal and explicitly versioned decision transitions. Connect a test relying party only after the simulator verifies delayed-event and concurrency behavior.

Promotion gates should require deterministic historical replay, explicit treatment of collector outages, rejection of stale decision submissions, and reconciliation of published statuses to authority events. A failed downstream registry update should appear as a propagation failure rather than a completed lifecycle transition.

Open questions

How much observation coverage is sufficient for a given interval claim? Which dependency models remain usable as asset counts grow? How should a policy revision affect earlier valid determinations? What is the safest reliance behavior when a materiality result is known but the authority review is delayed?

Research should also test whether reviewers understand the difference between pending review, stale evidence, and revoked status. The architecture succeeds only when its time semantics remain visible to the people and systems relying on the result.

References

[1] NIST SP 800 137 Information Security Continuous Monitoring

Monitoring program foundation.

<https://csrc.nist.gov/pubs/sp/800/137/final>

[2] IETF RFC 9334 Remote ATtestation procedureS Architecture

Attestation freshness and appraisal context.

<https://www.rfc-editor.org/rfc/rfc9334>

[3] QILAHK Journal When a Change Becomes Material to Trust

Materiality and authority workflow framing.

<https://blog.qilahk.com/material-change-and-trust-continuity.html>

Reference editions and source pages checked October 12, 2026. The proposed mechanisms and evaluation plans in this paper are research designs, not reported experimental findings.