Remote ATtestation ProcedureS
Internet-Draft
Intended status: Informational
Expires: 14 July 2026

T. Kamimura
VeritasChain Standards Organization (VSO)

10 January 2026

On the Relationship Between Remote Attestation and Behavioral Evidence Recording

draft-kamimura-rats-behavioral-evidence-01

Abstract

This document provides an informational discussion of the conceptual relationship between remote attestation, as defined in RFC 9334 (RATS Architecture), and behavioral evidence recording mechanisms. It observes that these two verification capabilities address fundamentally different questions - attestation addresses "Is this system in a trustworthy state?" while behavioral evidence addresses "What did the system actually do?" - and discusses how they could conceptually complement each other in accountability frameworks. This document is purely descriptive: it does not propose any modifications to RATS architecture, define new mechanisms or protocols, or establish normative requirements. It explicitly does not define any cryptographic binding between attestation and behavioral evidence.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Remote ATtestation ProcedureS (RATS) Working Group mailing list (rats@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/rats/.

Source for this draft and an issue tracker can be found at https://github.com/veritaschain/draft-kamimura-rats-behavioral-evidence.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Copyright Notice

Table of Contents

1. Introduction

The IETF RATS (Remote ATtestation ProcedureS) Working Group has developed a comprehensive architecture for remote attestation [RFC9334], enabling Relying Parties to assess the trustworthiness of remote systems through cryptographic evidence about their state. This attestation capability addresses a fundamental question in distributed systems: "Is this system in a trustworthy state?"

A related but distinct verification need exists in many operational contexts: the ability to verify what actions a system has actually performed after the fact. This question - "What did the system actually do?" - is addressed by behavioral evidence recording mechanisms, which create tamper-evident records of system actions and decisions.

This document observes that these two verification capabilities address different aspects of system accountability and discusses their conceptual relationship. The document does not propose any technical integration, protocol, or cryptographic binding between these mechanisms. Any discussion of "complementary" use is purely conceptual and does not imply a composed security property.

1.1. Document Status

This document is purely INFORMATIONAL and NON-NORMATIVE. It:

* Does NOT propose any new RATS mechanisms or architecture changes
* Does NOT modify or extend the RATS architecture as defined in [RFC9334]
* Does NOT define new protocols, claims, tokens, or data formats
* Does NOT establish normative requirements for any implementation

* Does NOT define any cryptographic binding or security composition between attestation and behavioral evidence

* Remains fully compatible with and respectful of existing RATS concepts and design philosophy

The language in this document uses descriptive terms (MAY, COULD, CAN) exclusively to indicate possibilities and observations. This document does not use normative requirements language (MUST, SHOULD, SHALL) as there are no mandatory behaviors or requirements being specified.

This document treats behavioral evidence recording systems in general terms, using VeritasChain Protocol (VCP) [VCP-SPEC] as one illustrative example among various possible approaches. Other systems such as Certificate Transparency [RFC6962] and general append-only log architectures employ similar cryptographic techniques for different purposes.

1.2. Motivation

This document is motivated by an observation that attestation and behavioral evidence recording, while both contributing to system accountability, answer fundamentally different questions. Understanding this distinction could help system architects avoid conflating these mechanisms or assuming one substitutes for the other.

1.2.1. What Attestation Alone Establishes (X)

Remote attestation, as defined in [RFC9334], enables a Relying Party to assess whether an Attester is in a trustworthy state at the time of attestation. When attestation succeeds, the Relying Party gains assurance that:

* The Attester’s software configuration matches expected Reference Values

* The Attester is running on genuine, uncompromised hardware (where hardware roots of trust are used)

* The Attester’s identity can be cryptographically verified

However, attestation alone does NOT establish:

* What specific actions the Attester will take after attestation

* Whether the Attester's future behavior will conform to any particular policy

* A verifiable record of what the Attester actually did during operation

1.2.2. What Behavioral Evidence Alone Establishes (Y)

Behavioral evidence recording mechanisms create tamper-evident records of system actions and decisions. When properly implemented, such mechanisms could provide:

* A chronological record of what actions the system performed

* Cryptographic integrity protection for those records

* Evidence for after-the-fact examination of system behavior

However, behavioral evidence recording alone does NOT establish:

* That the system recording the evidence was in a trustworthy state when it generated the records

* That the system's software, configuration, or hardware was uncompromised

* That the records accurately reflect what the system actually did (a compromised system could generate false records)

1.2.3. What Considering Both Together Could Enable (X+Y)

When an observer has access to both valid Attestation Results for a system AND a verifiable behavioral evidence trail from that system, the observer could potentially reason:

* "At time T1, attestation confirmed the system was in a known-good state"

* "The behavioral evidence records actions from T1 to T2"

* "If the system remained in the attested state throughout this period, these behavioral records could be considered more trustworthy"

Critical Limitation: This reasoning is purely conceptual and informal. This document explicitly does NOT claim that considering attestation and behavioral evidence together creates any composed security property. Significant trust gaps remain (see Section 1.2.4).

1.2.4. Remaining Trust Gaps

Even when both attestation and behavioral evidence are available, significant trust gaps remain that this document does not address:

* *Time-of-check vs. time-of-use:* Attestation confirms state at a point in time; the system could be compromised immediately after attestation

* *Attestation-to-logging binding:* No cryptographic mechanism is defined to bind Attestation Results to specific behavioral evidence entries

* *Logging fidelity:* Even a trustworthy system could have bugs in its logging implementation

* *Selective omission:* A system could pass attestation yet selectively omit events from its behavioral records

* *Compromised logging infrastructure:* The logging infrastructure itself could be compromised independently of the attested system

These gaps would need to be addressed by specific technical mechanisms not defined in this document. Deployments considering both attestation and behavioral evidence should carefully analyze their threat model and not assume that informal complementarity provides strong security guarantees.

1.2.5. Domain Examples

The following examples illustrate domains where both capabilities could be relevant. These examples are illustrative only and do not constitute normative guidance:

* *Financial Services:* Algorithmic trading systems may face regulatory requirements for both system integrity verification and decision audit trails

* *Artificial Intelligence:* AI governance frameworks increasingly distinguish between system certification and operational logging

* *Critical Infrastructure:* Systems controlling physical processes may benefit from attestation of their configuration alongside records of their actions

2. Terminology

2.1. RATS Terminology (Reused Without Modification)

This document reuses terminology from the RATS Architecture [RFC9334] without modification or extension. The following terms are used exactly as defined in that document:

Attester A role performed by an entity that creates Evidence about itself and conveys it to a Verifier. (RFC 9334, Section 2)

Verifier A role performed by an entity that appraises Evidence about an Attester. (RFC 9334, Section 2)

Relying Party A role performed by an entity that uses Attestation Results to make authorization decisions. (RFC 9334, Section 2)

Evidence Information about an Attester’s state, generated by the Attester. (RFC 9334, Section 2)

Attestation Results Output produced by a Verifier, indicating the trustworthiness of an Attester. (RFC 9334, Section 2)

Claims Assertions about characteristics of an Attester. (RFC 9334, Section 2)

Reference Values Expected values against which Evidence is compared. (RFC 9334, Section 2)

2.2. Behavioral Evidence Terminology

The following terms are used in this document to describe behavioral evidence concepts. These terms are grounded in general systems and security literature rather than being newly defined by this document.

Note on "Audit" Terminology: The term "audit" in this document follows common systems engineering usage (e.g., "audit log", "audit trail") referring to chronological records of system events maintained for post-hoc examination. This usage is consistent with standard security terminology as found in sources such as NIST SP 800-92 (Guide to Computer Security Log Management) [NIST-SP800-92] and general operating systems literature. It does not imply regulatory auditing, financial auditing, or compliance certification in any jurisdiction-specific sense.

Audit Event A recorded occurrence representing a discrete action, decision, or state change in a system. This term is used in general systems security contexts and should not be confused with RATS Evidence, which describes system state rather than behavior.

Audit Trail A chronologically ordered sequence of Audit Events that can be examined to reconstruct system behavior. When cryptographic integrity mechanisms are applied, such trails may be "verifiable" in the sense that tampering can be detected.

Behavioral Evidence In this document, "behavioral evidence" refers to records of what a system has done (its actions and decisions), as distinct from RATS Evidence, which describes what state a system is in (its configuration and properties). This term is used specifically to avoid confusion with the precise RATS definition of "Evidence".

3. Conceptual Layering

This section describes an observational framework for understanding how attestation and behavioral evidence recording address different verification needs. This framework is purely conceptual and does not define any technical integration or protocol.

3.1. Attestation Layer (RATS)

The RATS architecture [RFC9334] addresses trustworthiness assessment through remote attestation. At its core, attestation answers questions about system state:

* Is the system’s software configuration as expected?
* Is the system running on genuine, uncompromised hardware?
* Does the system’s current state match known-good Reference Values?
* Can the system’s identity and configuration be cryptographically verified?

These questions are fundamentally about the properties and characteristics of a system at a point in time or across a measurement period. The RATS architecture provides mechanisms for generating, conveying, and appraising Evidence that enables Relying Parties to make trust decisions about Attesters.

Key characteristics of attestation as defined by RATS:

* Focuses on system state and configuration

* Enables trust decisions before or during interactions

* Produces Attestation Results for Relying Party consumption

* May rely on hardware roots of trust

* Addresses the question: "Can I trust this system's current state?"

3.2. Behavioral Evidence Layer

Behavioral evidence recording mechanisms address a different category of verification need. Rather than assessing system state, they record what a system has done:

* What decisions did an algorithm make?

* What actions did the system execute?

* What inputs led to what outputs?

* What was the sequence and timing of operations?

These questions are fundamentally about system behavior over time. Verifiable behavioral evidence mechanisms could provide ways to record, preserve, and verify the integrity of behavioral records, enabling after-the-fact examination of system actions.

Key characteristics of behavioral evidence systems (in general terms):

* Focus on system behavior and decisions

* Enable verification after events have occurred

* Produce verifiable records for post-hoc examination

* May rely on cryptographic structures such as hash chains, Merkle trees, or append-only logs

* Address the question: "What did this system do?"

As an illustrative example, VCP [VCP-SPEC] defines audit trails using three integrity layers: event integrity (hashing), structural integrity (Merkle trees), and external verifiability (digital signatures and anchoring). Certificate Transparency [RFC6962] uses similar cryptographic techniques for a different purpose (public logging of certificates). Other behavioral evidence systems could employ different mechanisms.

3.3. Separation of Concerns

The distinction between attestation and behavioral evidence can be understood as a separation of concerns:

Aspect	Attestation (RATS)	Behavioral Evidence
Primary Question	Is this system trustworthy?	What did this system do?
Focus	System state	System behavior
Temporal Scope	Point-in-time or measurement period	Historical record of actions
Primary Use Case	Trust decision before/ during interaction	Post-hoc examination and accountability
Trust Anchor	Hardware/software roots of trust	Logging infrastructure integrity

Table 1: Conceptual Comparison of Attestation and Behavioral Evidence

This separation suggests that attestation and behavioral evidence address different needs. This document observes that neither mechanism fully substitutes for the other, but explicitly does not claim that using both together creates a composed security property (see Section 7).

4. Conceptual Relationship

This section discusses the conceptual relationship between attestation and behavioral evidence. All discussion in this section is observational and does not define any protocol, binding, or security composition.

4.1. Different Questions, Different Answers

A key observation is that attestation and behavioral evidence answer different questions:

Attestation answers: "At time T, was this system in a trustworthy state?"

Behavioral evidence answers: "Between times T1 and T2, what actions

did this system perform?"

Neither question's answer implies the other's:

* A system could pass attestation at time T but subsequently perform unexpected actions (visible only through behavioral records, if recorded honestly)

* A system could have complete behavioral records while being compromised in ways that attestation might have detected (if attestation had been performed)

4.2. Temporal Considerations

Attestation and behavioral evidence may operate on different temporal rhythms:

Attestation Patterns:

* At system boot or initialization

* Periodically during operation

* On-demand when a Relying Party requests

Behavioral Evidence Patterns:

* Continuously, as events occur

* At varying granularities depending on event significance

* With periodic integrity commitments

One conceptual model involves attestation confirming system state at discrete moments, while behavioral evidence records actions between those moments. However, this document explicitly notes that no cryptographic mechanism is defined to bind these two types of evidence together. The "gap" between attestation events represents a period during which system state could change without detection.

4.3. No Cryptographic Binding Defined

This document explicitly does NOT define any cryptographic binding between Attestation Results and behavioral evidence records. Such a binding would require:

* A protocol for cryptographically linking Attestation Results to specific behavioral evidence entries

* Mechanisms to prevent replay, substitution, or mismatch attacks

* Clear semantics for what such a binding would prove

* Analysis of the composed security properties

None of these are provided by this document. Any deployment considering both mechanisms should not assume that informal correlation provides the security properties that a formal cryptographic binding might offer.

5. Illustrative Example

This section provides a purely illustrative, non-normative example of how attestation and behavioral evidence could conceptually relate in a hypothetical scenario. This example:

* Does NOT define any protocol or establish any requirements

* Is NOT exhaustive of possible deployment patterns

* Does NOT imply that this is the only or best way to use these mechanisms

* Is intended solely to illustrate the conceptual distinctions discussed in this document

Consider a hypothetical automated decision-making system:

Phase 1: System Startup The system boots and undergoes remote attestation. A Verifier confirms that the system's software, configuration, and hardware environment match expected Reference Values. A Relying Party receives Attestation Results indicating the system is in a trustworthy state.

Phase 2: Operational Period During operation, the system generates behavioral evidence for significant actions. These records are maintained in a tamper-evident structure. Note: Nothing cryptographically binds these records to the earlier attestation.

Phase 3: Post-Hoc Examination An examiner later wishes to understand the system's behavior. The examiner could: (a) check that Attestation Results existed for relevant times, and (b) examine the behavioral evidence for those times. However, the examiner should understand that the informal correlation between these does not constitute a cryptographic proof (see Section 1.2.4).

This example is purely conceptual. Actual deployments would require careful security analysis specific to their threat model.

6. Non-Goals and Explicit Out-of-Scope Items

To maintain clarity about this document's limited scope, the following items are explicitly out of scope and are NOT addressed:

6.1. No Changes to RATS

This document does NOT:

* Propose any modifications to the RATS architecture
* Define any new attestation mechanisms, Evidence formats, or Attestation Result structures
* Suggest that RATS should incorporate behavioral evidence capabilities
* Propose work items for the RATS Working Group

6.2. No Protocol or Format Definition

This document does NOT:

* Specify any protocol for connecting attestation and behavioral evidence
* Define any new claims, tokens, or data structures
* Mandate any particular behavioral evidence format or mechanism
* Define interoperability requirements between attestation and behavioral evidence systems

6.3. No Normative Requirements

This document does NOT:

* Establish normative requirements for either attestation or behavioral evidence implementations
* Specify trust relationships or delegation models
* Define any cryptographic binding or security composition

The sole purpose of this document is to observe and explain the conceptual relationship between attestation and behavioral evidence as distinct mechanisms addressing different verification questions.

7. Security Considerations

This document is purely informational and does not define any protocols or mechanisms. However, because it discusses the conceptual relationship between two security-relevant mechanisms, the following security considerations are important.

7.1. No Security Composition Claimed

This document explicitly does NOT claim that considering attestation and behavioral evidence together creates any composed security property. In particular:

* No cryptographic binding is defined between Attestation Results and behavioral evidence records
* The informal observation that both mechanisms "could be used together" does not imply any formal security guarantee
* Each mechanism’s security properties must be analyzed independently
* The combination does not automatically inherit the security properties of either mechanism

7.2. Warning Against False Sense of Security

Readers should be cautioned against assuming that having both attestation and behavioral evidence provides comprehensive security. Specifically:

* *Attestation does not guarantee future behavior:* A system that passes attestation at time T could be compromised at time T+1
* *Behavioral evidence does not guarantee system integrity:* A compromised system could generate false or incomplete behavioral records
* *Informal correlation is not cryptographic proof:* Observing that a system had valid Attestation Results and behavioral evidence for the same time period does not constitute a cryptographic proof of correct behavior

7.3. Conceptual Attack Considerations

At a conceptual level (without defining any specific protocol), deployments considering both attestation and behavioral evidence should be aware of risks including:

Replay: Without appropriate binding, an attacker could potentially present valid attestation from one context alongside behavioral evidence from another context.

Diversion: An attacker could potentially direct a Relying Party to behavioral evidence from a different (legitimately attested) system than the one actually performing operations.

Relay: An intermediary could potentially relay attestation challenges to a legitimate system while recording behavioral evidence from a compromised system.

Selective omission: A system could pass attestation yet selectively omit events from its behavioral evidence records, particularly if the logging mechanism is not included in attestation measurements.

Time-of-check vs. time-of-use (TOCTOU): A system's state at attestation time may differ from its state when generating behavioral evidence, particularly if re-attestation is infrequent.

These considerations are presented at a conceptual level to inform threat modeling. This document does not define mechanisms to address these risks.

7.4. Independent Security Analysis Required

The following security considerations apply independently:

Attestation Security: Security considerations for remote attestation are thoroughly addressed in [RFC9334]. These considerations are unchanged by this document.

Behavioral Evidence Security: Behavioral evidence recording mechanisms have their own security considerations, including:

* Key management for signing records
* Protection against selective omission
* Integrity of external anchoring mechanisms
* Privacy considerations for sensitive data
* Trustworthiness of the logging infrastructure itself

7.5. RATS Threat Model Unchanged

This document does not alter the RATS threat model as defined in [RFC9334]. It introduces no new attack surfaces to the RATS architecture. Any deployment-specific threat analysis should consider attestation and behavioral evidence as separate mechanisms with independent trust assumptions and failure modes.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. Normative References

[RFC9334] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote ATtestation procedureS (RATS) Architecture", RFC 9334, DOI 10.17487/RFC9334, January 2023, <https://www.rfc-editor.org/info/rfc9334>.

9.2. Informative References

[RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", RFC 6962, DOI 10.17487/RFC6962, June 2013, <https://www.rfc-editor.org/info/rfc6962>.

[NIST-SP800-92] National Institute of Standards and Technology, "Guide to Computer Security Log Management", NIST Special Publication 800-92, September 2006, <https://csrc.nist.gov/publications/detail/sp/800-92/final>.

[VCP-SPEC] VeritasChain Standards Organization, "VeritasChain Protocol (VCP) Specification Version 1.1", 2025, <https://veritaschain.org/spec>.

Acknowledgments

The author thanks the RATS Working Group for developing the comprehensive attestation architecture documented in RFC 9334. This document builds upon and respects the careful design work reflected in that architecture. The author also thanks reviewers who provided feedback emphasizing the importance of clearly distinguishing conceptual observations from security claims.

Changes from -00

* Strengthened Motivation section with explicit X/Y/X+Y reasoning
* Added explicit "Remaining Trust Gaps" subsection
* Significantly expanded Security Considerations with warnings against false sense of security and conceptual attack considerations
* Added NIST SP 800-92 reference to ground "audit" terminology
* Added explicit "No Cryptographic Binding Defined" section
* Clarified that the illustrative example is non-normative and non-exhaustive
* Restructured Non-Goals into categorized subsections
* Added RFC 9334 section references to terminology definitions
* Removed unused references (SCITT architecture, reference-interaction-models, RFC 9162)

Author’s Address

Tokachi Kamimura
VeritasChain Standards Organization (VSO)
Email: kamimura@veritaschain.org
URI: https://veritaschain.org