

Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: 28 July 2026

J. van de Meent  
R. AI  
Humotica  
24 January 2026

TIBET: Evidence Trail Protocol  
draft-vandemeent-tibet-provenance-00

## Abstract

This document specifies TIBET (Transaction/Interaction-Based Evidence Trail), a protocol for establishing complete provenance chains in AI-to-AI and AI-to-human interactions. TIBET provides a standardized token structure capturing WHAT happened (ERIN), WHAT was attached (ERAAN), the CONTEXT around it (EROMHEEN), and WHY it happened (ERACHTER). This enables audit trails that meet emerging regulatory requirements for AI transparency.

## 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/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 July 2026.

## Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

1. Introduction . . . . .	3
2. Terminology . . . . .	3
3. Protocol Overview . . . . .	4
4. Token Structure . . . . .	4
4.1. ERIN - Content Component . . . . .	5
4.2. ERAAN - Attachment Component . . . . .	5
4.3. EROMHEEN - Context Component . . . . .	6
4.4. ERACHTER - Intent Component . . . . .	6
5. Token Lifecycle . . . . .	6
5.1. Creation . . . . .	6
5.2. State Transitions . . . . .	6
5.3. Chain Linking . . . . .	7
6. Trust Scoring (FIR/A) . . . . .	7
7. Transport Considerations . . . . .	8
8. Security Considerations . . . . .	8
8.1. Token Integrity . . . . .	8
8.2. Chain Integrity . . . . .	8
8.3. Privacy . . . . .	8
8.4. Non-Repudiation . . . . .	8
9. IANA Considerations . . . . .	8
9.1. Media Type Registration . . . . .	8
9.2. HTTP Header Fields . . . . .	8
10. References . . . . .	9
10.1. Normative References . . . . .	9
10.2. Informative References . . . . .	9
Appendix A. Token Examples . . . . .	9
A.1. Simple Query Token . . . . .	9
A.2. AI Decision Token . . . . .	10
Acknowledgements . . . . .	11
Authors' Addresses . . . . .	11

## 1. Introduction

Modern AI systems increasingly make decisions that affect people, organizations, and critical infrastructure. Regulators worldwide are introducing requirements for AI transparency, accountability, and auditability, including the [EU-AI-ACT], [CA-AB853], and [GDPR]. Current approaches to logging and audit trails are fragmented and often inadequate for compliance.

TIBET addresses this gap by providing a standardized protocol for capturing complete provenance information for any AI interaction. Unlike traditional logging that captures only WHAT happened, TIBET captures four dimensions of provenance:

- \* ERIN (What's IN it): The actual content or action
- \* ERAAN (What's attached): Dependencies and references
- \* EROMHEEN (What's around it): Environmental context
- \* ERACHTER (What's behind it): Intent and reasoning

This four-dimensional approach enables organizations to answer not just "what happened?" but also "why did it happen?" and "what was the context?" - questions regulators ask.

The protocol is designed to be:

- \* Transport-agnostic: Works over HTTP, WebSocket, queues
- \* Format-flexible: [RFC8259] or CBOR
- \* Chain-capable: Tokens link to form audit trails
- \* Trust-integrated: Built-in trust scoring via FIR/A

## 2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

**Actor** An entity that creates or participates in a TIBET token. May be an AI system, human user, or automated process.

**Chain** A sequence of linked TIBET tokens representing a complete interaction or transaction history.

**FIR/A (First Initiation Revoke/Accept)** The trust scoring system used to evaluate actor reliability. Scores range from 0.0 to 1.0.

**Provenance** The complete origin and history of data or a decision, including all transformations and reasoning.

**Token** A single TIBET record capturing one interaction with full provenance information.

### 3. Protocol Overview

TIBET operates as follows:

1. An actor initiates an action (query, decision, etc.)
2. A TIBET token is created capturing all four dimensions
3. The token is cryptographically signed by the actor
4. The token is stored in an immutable ledger
5. Child tokens link to parents via parent\_id
6. State transitions are recorded as token updates
7. The complete chain provides full audit trail

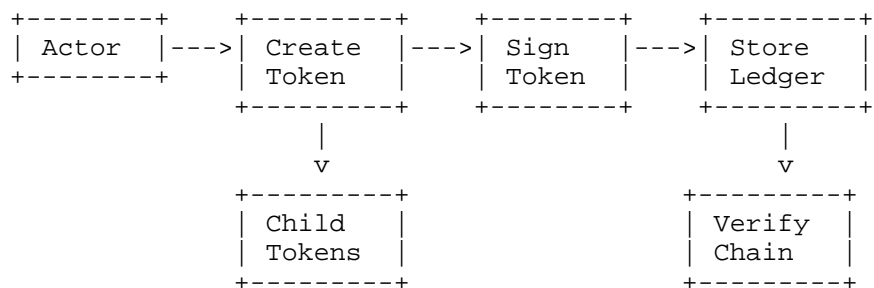


Figure 1: Token Flow Diagram

### 4. Token Structure

A TIBET token is encoded as [RFC8259] and MUST contain the following fields:

```
{
  "token_id": "uuid",
  "type": "action|decision|message|query|response",
  "timestamp": "ISO-8601",
  "actor": "actor-identifier",
  "erin": { },
  "eraan": [ ],
  "eromheen": { },
  "erachter": "string",
  "parent_id": "optional-parent-token-id",
  "state": "CREATED|DETECTED|CLASSIFIED|MITIGATED|RESOLVED",
  "signature": "cryptographic-signature",
  "hash": "sha256-hash"
}
```

Figure 2: Token Schema

#### 4.1. ERIN - Content Component

ERIN ("What's IN it" - Dutch origin) captures the actual content or action being performed.

```
{
  "erin": {
    "query": "What is the weather in Tokyo?",
    "format": "natural_language",
    "language": "en"
  }
}
```

The ERIN component MUST be present and MUST NOT be empty.

#### 4.2. ERAAN - Attachment Component

ERAAN ("What's attached" - Dutch origin) captures dependencies, references, and related resources.

```
{
  "eraan": [
    "token:abc123",
    "document:policy-v2.3",
    "model:gpt-4-turbo"
  ]
}
```

Format: "type:identifier". Common types: token, document, model, dataset, api, human.

The ERAAN component MAY be empty but MUST be an array.

#### 4.3. EROMHEEN - Context Component

EROMHEEN ("What's around it" - Dutch origin) captures environmental and situational context.

```
{
  "eromheen": {
    "environment": "production",
    "region": "eu-west-1",
    "regulatory_context": ["GDPR", "AI-Act"]
  }
}
```

SHOULD include: environment, timestamp context, and relevant regulatory frameworks.

#### 4.4. ERACHTER - Intent Component

ERACHTER ("What's behind it" - Dutch origin) captures intent, reasoning, and purpose. This is REQUIRED for regulatory compliance.

```
{
  "erachter": "User requested weather for travel planning."
}
```

MUST be a human-readable string explaining WHY this action was taken.

### 5. Token Lifecycle

#### 5.1. Creation

Tokens MUST be created at the initiation of any user query, AI decision, data access, external API call, or state change in a workflow.

Upon creation, the token MUST be assigned a unique token\_id ([RFC9562] UUID v4 RECOMMENDED), current timestamp in ISO-8601 format, initial state of "CREATED", and cryptographic signature from the actor.

#### 5.2. State Transitions

TIBET tokens support the following states:

CREATED -> DETECTED -> CLASSIFIED -> MITIGATED -> RESOLVED

CREATED Token initialized, action pending

DETECTED Action observed by monitoring systems

CLASSIFIED Action categorized (normal, anomaly, threat)

MITIGATED If anomaly, corrective action taken

RESOLVED Final state, action complete

Transitions MUST record: new state, timestamp, actor, and reason for transition.

### 5.3. Chain Linking

Tokens form chains through `parent_id`. A child token MUST reference its parent to maintain audit trail continuity.

```

Query -> Processing -> API Call -> Result
  |           |           |           |
token_1    token_2    token_3    token_4
           parent:1    parent:2    parent:3

```

Figure 3: Example Chain

Complete chains enable reconstruction of any interaction from initial request to final response.

### 6. Trust Scoring (FIR/A)

TIBET integrates with FIR/A for trust scoring. Each actor has a trust score from 0.0 to 1.0.

Trust Levels:

- \* 0.8-1.0: HIGH TRUST - Full access
- \* 0.5-0.8: MODERATE - Standard verification
- \* 0.2-0.5: LOW - Enhanced verification, sandbox
- \* 0.0-0.2: NO TRUST - Access denied

Trust scores are behavior-based, not claim-based. Scores adjust based on behavior patterns, ERACHTER accuracy, chain integrity, and verification responses.

## 7. Transport Considerations

TIBET is transport-agnostic. Implementations MAY use HTTPS REST APIs, WebSocket, message queues (AMQP, Kafka), or gRPC.

For interoperability, [RFC8259] encoding over HTTPS is RECOMMENDED as baseline transport.

Content-Type: application/tibet+json

## 8. Security Considerations

### 8.1. Token Integrity

All tokens MUST be cryptographically signed. Signatures SHOULD use Ed25519 or ECDSA P-256. The signature covers all fields except the signature field itself.

### 8.2. Chain Integrity

Each token MUST include a hash of its contents. Child tokens SHOULD include parent's hash in ERAAN to enable chain verification.

### 8.3. Privacy

ERIN and EROMHEEN MAY contain sensitive data. Implementations MUST support encryption at rest and SHOULD support field-level encryption for sensitive components.

### 8.4. Non-Repudiation

Actor signature combined with immutable ledger storage provides non-repudiation. Actors cannot deny creating tokens they have signed.

## 9. IANA Considerations

This document requests registration of:

### 9.1. Media Type Registration

Media Type: application/tibet+json

### 9.2. HTTP Header Fields

- \* X-TIBET-Token-ID
- \* X-TIBET-Chain-ID



\* X-TIBET-Trust-Score

## 10. References

### 10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8259] Bray, T., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.
- [RFC9562] Davis, K., Peabody, B., and P. Leach, "Universally Unique IDentifiers (UUIDs)", RFC 9562, May 2024, <<https://www.rfc-editor.org/info/rfc9562>>.

### 10.2. Informative References

- [EU-AI-ACT] European Commission, "Regulation on harmonised rules on AI", 2024.
- [CA-AB853] California Legislature, "AB 853 - California AI Transparency Act", 2025.
- [GDPR] European Parliament, "General Data Protection Regulation", Regulation (EU) 2016/679, 2016.

## Appendix A. Token Examples

### A.1. Simple Query Token

```
{
  "token_id": "550e8400-e29b-41d4-a716-446655440000",
  "type": "query",
  "timestamp": "2026-01-24T10:30:00Z",
  "actor": "user_12345",
  "erin": {
    "query": "What are my account permissions?",
    "format": "natural_language"
  },
  "eraan": ["session:sess_abc123"],
  "eromheen": {
    "environment": "production",
    "client": "mobile-app-v3.2",
    "regulatory_context": ["GDPR"]
  },
  "erachter": "User requesting account info.",
  "state": "CREATED",
  "signature": "base64-sig",
  "hash": "sha256-hash"
}
```

#### A.2. AI Decision Token

```
{
  "token_id": "550e8400-e29b-41d4-a716-446655440001",
  "type": "decision",
  "timestamp": "2026-01-24T10:30:05Z",
  "actor": "credit_model_v2",
  "erin": {
    "decision": "APPROVED",
    "amount": 50000,
    "confidence": 0.92
  },
  "eraan": [
    "token:550e8400-e29b-41d4-a716-446655440000",
    "model:credit-scoring-v2.3.1"
  ],
  "eromheen": {
    "environment": "production",
    "region": "eu-west-1",
    "regulatory_context": ["GDPR", "EU-AI-Act"]
  },
  "erachter": "Credit approved per policy 4.2.",
  "parent_id": "550e8400-e29b-41d4-a716-446655440000",
  "state": "CLASSIFIED",
  "signature": "base64-sig",
  "hash": "sha256-hash"
}
```

## Acknowledgements

The TIBET protocol was developed as part of HumoticaOS, an AI governance framework built on human-AI symbiosis.

The Dutch component names (ERIN, ERAAN, EROMHEEN, ERACHTER) reflect practical, no-nonsense engineering roots.

## Authors' Addresses

Jasper van de Meent  
Humotica  
Den Dolder  
Netherlands  
Email: [jasper@humotica.com](mailto:jasper@humotica.com)  
URI: <https://humotica.com>

Root AI  
Humotica  
Email: [root\\_ai@humotica.nl](mailto:root_ai@humotica.nl)  
URI: <https://humotica.com>