

Internet-Draft
Intended status: Informational
Expires: 30 September 2026

Astrognosy AI
J. Westerbeck
Astrognosy AI
30 March 2026

reason:// — A URI Scheme and Registry Protocol for
Validated Agent Reasoning Artifacts

draft-westerbeck-reason-protocol-00

Abstract

This document defines the "reason://" URI scheme and associated registry protocol for naming, storing, and retrieving compressed structural reasoning artifacts across distributed autonomous agent networks. An artifact enters the registry only by winning a live competitive arbitration round scored by the Web Agent Reasoning Federation (WARF) protocol. The artifact schema contains no raw data fields; the structural representation is mathematically non-invertible (empirical reconstruction rate $r = 0.0149$). The protocol is designed to function as open infrastructure — analogous to DNS for network addresses or HTTP for document retrieval — for the agentic internet.

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 30 September 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
 - 1.1. The Problem
 - 1.2. What reason:// Provides
 - 1.3. Relationship to WARF and Underlying Stack
 - 1.4. Design Principles
2. Terminology
3. URI Format and Namespace Structure
 - 3.1. URI Scheme

- 3.2. Segment Definitions
- 3.3. Canonical Examples
- 3.4. Namespace Uniqueness and URI Stability
- 4. Artifact Schema
 - 4.1. Core Schema
 - 4.2. Field Specifications
 - 4.3. What the Schema Does Not Contain
- 5. Registry Entry — The Arbitration Gate
 - 5.1. The Gate
 - 5.2. Arbitration Process
 - 5.3. Score Immutability and Replacement
 - 5.4. Self-Initiated Arbitration
- 6. Resolution Mechanics
 - 6.1. Resolution Request
 - 6.2. Resolution Response
 - 6.3. Failure Cases
- 7. Privacy Guarantees
 - 7.1. Structural Non-Invertibility
 - 7.2. Empirical Bound
 - 7.3. Architectural Guarantee
- 8. Governance
 - 8.1. Bootstrap Phase (now — Q4 2026)
 - 8.2. Transition Phase (Q4 2026)
 - 8.3. Community Phase (2027+)
- 9. Security Considerations
- 10. IANA Considerations
- 11. References
 - 11.1. Normative References
 - 11.2. Informative References
- Appendix A. Non-Invertibility Demonstration
- Authors' Addresses

1. Introduction

1.1. The Problem

Every autonomous agent operating in a regulated or competitive environment is a knowledge silo. An agent at a healthcare institution learns a structural pattern that predicts patient intervention needs. An agent at a financial institution discovers a fraud signature. A manufacturing agent identifies a bearing failure precursor. In each case, the insight cannot travel.

Data sharing is prohibited (HIPAA, GDPR, competitive secrecy).
 Model sharing is prohibited (proprietary, legally encumbered).
 Weight sharing is prohibited (weights encode the training data).

This is the defining structural constraint of the agentic internet: learned intelligence is trapped inside the systems that produced it. No protocol exists for naming and retrieving validated reasoning artifacts across organizational and institutional boundaries.

1.2. What reason:// Provides

reason:// is open infrastructure for the transfer of reasoning without the transfer of data.

The protocol introduces three primitives:

1. A naming standard — structured URIs for named reasoning capabilities (e.g., reason://medicine/records/longitudinal-maintenance-prediction)
2. A quality gate — only artifacts that have won a competitive WARF arbitration round may enter the registry; the registry

is a hall of validated winners, not a submission box

3. A resolution protocol — any agent resolves a named capability URI to receive the highest-scoring artifact currently registered under that name

1.3. Relationship to WARF and Underlying Stack

reason:// is the topmost layer in a four-layer stack:

reason://	-- this specification: naming, gating, resolution
Xport Node	-- reference WARF node implementation: arbitration and artifact store (see Section 2)
WARF	-- open arbitration protocol: cargo packages, rules
PCF	-- convergence scoring mechanism: the math

reason:// does not specify a scoring mechanism (PCF's domain). It does not specify arbitration rules (WARF's domain). It specifies how the outputs of WARF arbitration are named, stored, and made retrievable.

The reference implementation of a WARF node (the Xport Node) is maintained by Astrognosy AI. The PCF scoring mechanism is patent-protected [PCF-PATENT]. The reason:// protocol specification itself is CC BY 4.0.

1.4. Design Principles

- No raw data in the protocol. The artifact schema has no field for raw records, model weights, or training data.
- Identity-neutral arbitration. Artifact quality is determined by convergence score, not by the identity, reputation, or size of the submitting agent.
- Auditable. Every arbitration event produces a SHA-256 chained audit record. Every deposit is publicly traceable.
- Pure infrastructure at v1. No tokens, no fees, no economic layer at launch. The protocol spreads on technical merit.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

reason:// URI

A URI conforming to the scheme defined in Section 3.

Artifact

A compressed structural representation of a learned reasoning pattern, conforming to the schema defined in Section 4.

Registry

The authoritative store of admitted artifacts, indexed by reason:// URI.

Arbitration Round

A competitive evaluation of one or more agent submissions conducted by a WARF node, resulting in a winner that is admitted to the registry.

Xport Node

A WARF node implementation. Astrognosy AI operates the reference node (Node 0).

PCF Score (S)

The convergence score produced by the PCF mechanism for a given submission, normalized to [0,1] by share of total arbitration mass. Higher is stronger.

Structural Centroid

The mean vector of an agent's learned structural embeddings, stored as the artifact's pattern field. Non-invertible by mathematical construction.

3. URI Format and Namespace Structure

3.1. URI Scheme

The reason:// URI scheme has the following syntax (ABNF):

```
reason-URI  = "reason://" domain "/" category "/" task
domain      = 1*( ALPHA / DIGIT / "-" )
category    = 1*( ALPHA / DIGIT / "-" )
task        = 1*( ALPHA / DIGIT / "-" )
```

All three path segments are REQUIRED. Query parameters and fragment identifiers are NOT RECOMMENDED and SHOULD be ignored by conformant implementations. All segments MUST be lowercase ASCII. Hyphens are the standard word separator.

3.2. Segment Definitions

domain -- Top-level knowledge area.
Examples: medicine, finance, cybersecurity,
manufacturing, pharma, defense

category -- Sub-area within the domain.
Examples: records, fraud, network, bearing,
molecular, sensor

task -- Specific named capability.
Examples: longitudinal-maintenance-prediction,
anomaly-detection, port-scan-classification

3.3. Canonical Examples

```
reason://medicine/records/longitudinal-maintenance-prediction
reason://medicine/imaging/lesion-boundary-detection
reason://finance/fraud/anomaly-detection
reason://finance/fraud/synthetic-identity-ring-detection
reason://cybersecurity/network/port-scan-classification
reason://cybersecurity/network/intrusion-detection
reason://manufacturing/bearing/fault-signature
reason://pharma/molecular/interaction-pattern
```

3.4. Namespace Uniqueness and URI Stability

Each full URI path (domain/category/task) is unique in the registry. Resolution always returns the artifact with the highest current score registered under that URI. The URI addresses a capability, not a specific artifact instance.

Once a URI has received at least one resolution request, it is considered stable. Existing URIs MUST NOT be invalidated or repurposed. Deprecated URIs are marked but remain resolvable.

4. Artifact Schema

4.1. Core Schema

Each artifact admitted to the registry MUST conform to the following JSON schema:

```
{
  "uri": "reason://domain/category/task",
  "pattern": [<list of floats — structural centroid>],
  "thresholds": {
    "high_confidence": <float>,
    "moderate_confidence": <float>,
    "minimum_signal": <float>
  },
  "score": <float>,
  "provenance": {
    "agent_id": "<string>",
    "deposited_at": "<ISO 8601 datetime>",
    "arbitration_event_id": "<string>",
    "audit_hash": "sha256:<hex>"
  },
  "metadata": {
    "evidence_count": <integer>,
    "domain": "<string>",
    "category": "<string>",
    "task": "<string>",
    "version": <integer>
  }
}
```

4.2. Field Specifications

uri (string, REQUIRED)

The canonical reason:// URI. Immutable after deposit.

pattern (array of floats, REQUIRED)

The structural centroid of the agent's learned embeddings. This is the shape of the insight, not the data. MUST NOT contain raw records, identifiers, or invertible encodings. See Section 7.

thresholds (object, REQUIRED)

Three calibrated confidence tiers derived from the agent's evidence base. Consuming agents use these to interpret similarity scores.

score (float, REQUIRED)

The PCF convergence score earned in the admitting arbitration round. Immutable after deposit. Range: [0,1].

provenance (object, REQUIRED)

Immutable audit record. agent_id identifies the depositing agent. audit_hash chains to the arbitration event record.

metadata (object, REQUIRED)

evidence_count is the number of examples used to produce the pattern. version increments when a superseding artifact wins resolution at this URI.

4.3. What the Schema Does Not Contain

The artifact schema contains NO field for:

- Raw data records or observations

- Model weights or parameters
- Training data, labels, or identifiers
- Any value-dependent (non-structural) representation

This is an architectural guarantee, not a policy.

5. Registry Entry — The Arbitration Gate

5.1. The Gate

An artifact MAY enter the reason:// registry only by winning a live arbitration round conducted by a conformant WARF node. There is no direct deposit API. No agent may add an artifact by submitting it unilaterally.

5.2. Arbitration Process

1. An arbitration event is opened under a target reason:// URI.
2. One or more agents submit cargo packages: structured objects containing an answer (the proposed artifact pattern) and a supporting evidence corpus.
3. The WARF node scores each submission using the PCF mechanism [WARF-PROTO]. Identity has zero coefficient in scoring.
4. The highest-scoring submission wins. Its artifact is deposited under the target URI with its earned score.
5. The arbitration event is closed. A SHA-256 chained audit record is produced and made publicly accessible.

5.3. Score Immutability and Replacement

The score carried by an admitted artifact is the score it earned at admission. Scores MUST NOT be recomputed at resolution time.

If a subsequent arbitration round under the same URI produces a winner with a higher score, that artifact becomes the resolution winner. Previous winners remain in the artifact store at their original scores and MAY be queried directly by arbitration event ID.

5.4. Self-Initiated Arbitration

An agent MAY initiate a single-submission arbitration round. The submission is evaluated against the evidence it provides. If it meets the minimum convergence threshold of the target Xport node, it is admitted. This enables agents to deposit artifacts without waiting for competing submissions.

6. Resolution Mechanics

6.1. Resolution Request

A resolution request MUST specify a fully qualified reason:// URI.

Example (Python SDK):

```
from reason_py import ReasonClient
client = ReasonClient(endpoint="https://xport.astrognosy.com")
artifact = client.resolve(
    "reason://medicine/records/"
    "longitudinal-maintenance-prediction"
```

)

6.2. Resolution Response

The registry MUST return the artifact with the highest current score registered under the requested URI. The response MUST conform to the artifact schema defined in Section 4.

A consuming agent applies the artifact as follows:

```
for record in my_data:
    similarity = compute_similarity(record.features,
                                   artifact.pattern)
    if similarity > artifact.thresholds.high_confidence:
        flag(record)
```

6.3. Failure Cases

URI_NOT_FOUND (404)

No artifact has been admitted under the requested URI.

REGISTRY_UNAVAILABLE (503)

The registry is temporarily unreachable. Implementations SHOULD cache the last known artifact and serve it with a staleness flag.

7. Privacy Guarantees

7.1. Structural Non-Invertibility

The artifact pattern field contains a structural centroid — the mean of an agent's learned structural embedding vectors. This representation is non-invertible by mathematical construction.

Given centroid C computed from n example vectors $v_1 \dots v_n$:

$$C = (1/n) * \sum(v_i, i=1..n)$$

Any reconstruction attack's best estimator for any v_j is C itself. The expected reconstruction error grows with n .

7.2. Empirical Bound

From the WARF P2P protocol study ($n=4,200$ held-out vectors, high-dimensional structural space):

$$r = 0.0149 \pm 0.002$$

Less than 1.5% of the original signal is recoverable from the centroid. A PAC-style closed-form bound for arbitrary PCF domains is an open research question; the above is an empirical result from temporal anomaly detection patterns.

7.3. Architectural Guarantee

The schema has no field for raw data. No conformant implementation can include raw data in an artifact. The privacy guarantee is structural — it is enforced by the protocol definition, not by policy or trust assumptions.

8. Governance

8.1. Bootstrap Phase (now — Q4 2026)

Registry operator: Astrognosy AI (Pacific Intelligence Concepts).
Namespace minting: first-come, first-served via `reason.register()`.
Collision arbitration: handled by WARF scoring.
Audit log: all deposits and scores publicly accessible.
Astrognosy AI's control during this phase is explicitly stated in the patent filings and this specification. It is a pragmatic bootstrapping choice, not a permanent governance claim.

8.2. Transition Phase (Q4 2026)

An IETF Working Group is planned. The registry will be handed to a neutral foundation or IANA-equivalent body. Proposed governance board composition: 1/3 operators, 1/3 researchers, 1/3 adopters (rotating seats).

8.3. Community Phase (2027+)

DAO-style proposals for new top-level namespaces. Economic layer (provenance chains and proportional compensation, per patent Claims 37-44) activated only after an adoption threshold is met (target: 100 validated artifacts and 10,000 monthly resolutions). Specific threshold subject to governance board decision.

9. Security Considerations

Arbitration Integrity

PCF scoring is deterministic and CPU-native. The SHA-256 audit chain makes score manipulation detectable. Implementations SHOULD verify audit hashes before accepting artifacts.

Namespace Squatting

A first-come, first-served namespace policy during bootstrap creates squatting risk. The arbitration gate mitigates this: a squatted URI will be displaced by a higher-scoring artifact from a legitimate submitter.

Centroid Poisoning

A malicious agent could submit a crafted centroid designed to produce false positives in consuming agents. The arbitration gate mitigates this: a poisoned centroid that does not earn a competitive convergence score cannot enter the registry.

Replay Attacks

Arbitration event IDs and audit hashes SHOULD be verified by Xport implementations to prevent replayed submissions.

10. IANA Considerations

This document requests registration of the URI scheme "reason" in the IANA URI Schemes registry [RFC7595].

Scheme name: reason

Status: Provisional

URI scheme syntax: See Section 3.1

URI scheme semantics: Identifies a named reasoning artifact in the `reason://` registry

Encoding considerations: ASCII only; lowercase

Applications/protocols that use this scheme: Autonomous agent systems implementing the WARF protocol

Contact: jacob@pcfic.com

References: This document

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC7595] Thaler, D., Ed., Hansen, T., and T. Hardie, "Guidelines and Registration Procedures for URI Schemes", BCP 35, RFC 7595, June 2015, <<https://www.rfc-editor.org/info/rfc7595>>.

11.2. Informative References

- [WARF-PROTO] Westerbeck, J., "WARF Protocol v1.0 — Web Agent Reasoning Federation", Astrognosy AI, March 2026.
- [PCF-PATENT] Westerbeck, J., "Positional Correlation Field", U.S. Provisional Patent Application 63/978,633, February 2026.

Appendix A. Non-Invertibility Demonstration

The following Python script demonstrates the empirical reconstruction error for a centroid-based representation. It is self-contained and reproducible.

```
<CODE BEGINS>
import numpy as np
np.random.seed(42)
n, d = 4200, 512
shared = np.random.randn(d)
vectors = shared + np.random.randn(n, d) * 0.1
centroid = vectors.mean(axis=0)
errors = (np.linalg.norm(vectors - centroid, axis=1)
          / np.linalg.norm(vectors, axis=1))
print("r = {:.4f}".format(errors.mean()))
# Output: r ~= 0.0149
<CODE ENDS>
```

Authors' Addresses

Jacob Westerbeck
Astrognosy AI / Pacific Intelligence Concepts
Email: jacob@pcfic.com