Protocol Drafts and Governance Extensions

Formal Specification of PHA-Hysteresis and ZMEM-Ethics Headers

Alexandre Sah*

January 30, 2026

Abstract

This online resource provides technical implementation specifications for the Algorithmic Hysteresis Primacy (AHP) framework, translating conceptual principles into network protocol primitives. We present formal ABNF syntax, operational semantics, and regulatory compliance mappings for the PHA-Hysteresis and ZMEM-Ethics HTTP headers. Implementation guidance includes state machine diagrams, example code, and threat modeling. These specifications operationalize mathematical guarantees established in supplementary materials and enable distributed governance mechanisms for multi-jurisdictional coordination.

Keywords: Network protocols; HTTP headers; Temporal architecture; ABNF specification; Distributed consensus; Radiation resilience

Contents

*This document provides technical implementation specifications for the Algorithmic Hysteresis Primacy (AHP) framework presented in the main article *Algorithmic Hysteresis Primacy (AHP): Temporal Sovereignty in AI Governance*, available at zmem.org and SSRN. Mathematical foundations are established in supplementary materials; reference implementations and validation protocols are also available at zmem.org and SSRN.

1 Introduction to Temporal Governance Protocols

This document is intended for technically literate readers seeking implementation-level clarity; readers interested primarily in conceptual, ethical, or political arguments are referred to the accompanying main article.

Remark 1 (Scope and Relation to the Main Article)**.** This document provides protocol-level specifications that translate the architectural principles introduced in the main text into implementable formats. No new empirical claims, normative positions, or theoretical constructs are introduced in this document beyond those already established in the primary article.

All technical elements defined herein (including PHA-Hysteresis headers and ABNF syntax) are to be interpreted as *operational realizations* of architectural constraints established conceptually in the accompanying article. Normative and governance-related terminology is used solely as an interpretive mapping and does not constitute independent normative claims.

Remark 2 (Scope clarification)**.** This document specifies protocol-level *reference formats* and *governance hooks*. It does not propose mandatory standards, adoption requirements, or regulatory prescriptions. All specifications are intended as optional, interoperable extensions that can be ignored or overridden without affecting baseline system functionality.

Remark 3 (Normative Language)**.** The key words *MUST*, *SHALL*, *SHOULD*, and *MAY* in this document are to be interpreted as described in RFC 2119. Sections explicitly labeled as **Normative** define conformance requirements; all other sections are informative.

Document Conventions

Sections labeled as **Normative** define mandatory protocol behavior. Sections labeled as **Informative** provide rationale, examples, or explanatory context and are not required for conformance.

Remark 4 (Status of the Specification)**.** The protocol elements specified in this document are *proposed experimental primitives*, not normative standards. They are intended as design exemplars demonstrating how mathematical guarantees can be operationalized at the protocol layer. No claim is made regarding immediate deployability, legal sufficiency, or standards-track maturity. Header names are prefixed for experimental use and do not imply IANA registration. The present document should be read as a design probe rather than a proposal for immediate standardization.

Remark 5 (Legal Interpretation Disclaimer)**.** This specification enables technical alignment with regulatory requirements; it does not define legal compliance. Regulatory mappings illustrate how temporal governance primitives could support compliance workflows, without asserting sufficiency under any jurisdiction.

Remark 6 (Why HTTP Headers)**.** HTTP headers are used here as a *convenient and inspectable carrier* for temporal governance signals, not as an exclusive or optimal implementation layer. Equivalent semantics could be realized at middleware, message-queue, or application layers. Headers are chosen for their visibility, auditability, and minimal intrusion into existing architectures.

1.1 Motivation and Design Philosophy

This appendix specifies the formal extension of Algorithmic Hysteresis Primacy (AHP) to network protocols through two complementary HTTP headers. These specifications operationalize the ZMEM conceptual architecture (Zero-Mean Execution Memory) defined in the main text, transforming the philosophical primitive of “could still not happen” into concrete technical mechanisms.

Throughout this document, the term "protocol" refers to a *descriptive interface contract*, not a prescriptive governance mandate. The specifications herein define how temporal semantics may be exposed and audited when such mechanisms are present.

Unless explicitly stated otherwise, all numerical values, latency ranges, and domain-specific parameters in this document are illustrative design ranges intended to demonstrate protocol instantiation, not prescriptive operational thresholds.

The specifications are grounded in:

- **Historical system failures**: Flash Crash 2010, BrainGate mishaps, UK Grid 2019
- **Cognitive science**: Human perception thresholds (approximately 100 ms), volition windows (approximately 200 ms) based on established neurocognitive research
- **Emerging regulatory frameworks**: South Korea's Framework Act on Intelligent Information Society (2024) mandating safety assurance and transparency for high-risk AI; EU AI Act (2024) Art. 14; GDPR (2016) Art. 22; Brazil LGPD (2018) Art. 20
- **Radiation-aware computing**: Temporal buffering against single-event upsets in space systems

The goal is to operationalize **protocol-provable hesitation** as both a technical safeguard and compliance mechanism.

Remark 7 (Protocol Foundations for Distributed Governance)**.** The headers specified herein provide the *transactional substrate* necessary for international governance mechanisms analyzed in supplementary materials. The PHA-Hysteresis header's guaranteed minimum delay creates the temporal window required for Byzantine consensus protocols—ensuring that distributed coordination across jurisdictions cannot be overwhelmed by pathological synchronization in the network layer.

1.2 Architectural Principles

The protocol design adheres to three core principles:

1. **Evidence-Based Decisions**: The system SHALL require sustained accumulation of evidence before state transitions
2. **Mandatory Hesitation**: The system SHALL enforce an architectural guarantee of minimum time $\Delta T_{\min} > 0$ between decisions (mathematically enforced via accumulator bounds)
3. **Transparent Accountability**: The system SHALL expose internal state (γ, accumulation progress) in real-time for human oversight and distributed audit

1.3 Protocol Complementarity

The two headers implement distinct but synergistic aspects of decisional inertia:

PHA-Hysteresis Header Implements mandatory *Decisional Inertia* based on the Primacy of Algorithmic Hysteresis. Enforces non-linear hysteresis bands to prevent pathological synchronization in autonomous systems. Focus: *Technical stability.* Critical for radiation resilience.

ZMEM-Ethics Header Mandates server-side hesitation periods before executing critical operations. Provides auditable windows for human intervention and reflection. Focus: *Deliberation.* Enables "veto windows" required for meaningful human control in distributed governance scenarios.

Both headers can be used independently or in combination, depending on system requirements. When deployed in multi-jurisdictional contexts, the combination provides both the technical hysteresis (PHA) and deliberation space (ZMEM) required for Byzantine consensus.

Remark 8 (Falsifiability note)**.** Each header field defined herein maps directly to state variables or invariants specified in mathematical formalizations and executable in reference implementations. Any implementation that violates these mappings constitutes a falsifiable deviation from the protocol semantics.

Out of Scope

This specification does not:

- define models for intent detection;
- prescribe machine learning architectures;
- mandate specific latency values beyond $\Delta T_{\min}$;
- replace domain-specific safety requirements;
- define legal compliance obligations.

2 PHA-Hysteresis Header Specification

2.1 **Informative** Historical and Conceptual Foundations

2.1.1 Latency Suppression Pathology

Modern digital architectures suffer from what we term **Latency Suppression Pathology**. Consequences include:

- **Zeno-like oscillations**: Rapid state switching in control loops
- **Cascading failures**: Near-instantaneous synchronization
- **Noise amplification**: Transient signals triggering consequential actions

2.1.2 Motivating Failures

Table 1: Historical Failures Motivating PHA-Hysteresis

Event	**Timescale**	**Mechanism**	**Impact**
2010 Flash Crash	1–10 μs	HFT feedback loops	$1T loss
BrainGate trials	1–10 ms	Neural noise misinterpretation	Unintended movements
UK Grid 2019	1–5 ms	Sync. inverter reactions	737 MW disconnect
Space system SEUs	1–100 ns	Radiation-induced transients	System resets

In all cases, *sub-100 ms reaction times* enabled pathological dynamics. PHA-Hysteresis architecturally prevents such dynamics by enforcing $\Delta T_{\min}$.

2.2 Normative Formal Syntax (ABNF)

The header follows RFC 5234 Augmented Backus-Naur Form (ABNF):

```
PHA-Hysteresis = "PHA-Hysteresis" ":"
                 state-param ";" delay-param
                 [ ";" conviction-param ]
                 [ ";" veto-endpoint ]
                 [ ";" confidence-param ]
                 [ ";" accumulator-param ]

state-param       = "state" "=" state-value
state-value       = "accumulating" / "stabilized" / "vetoed" / "timeout"
delay-param       = "delay" "=" 1*DIGIT
conviction-param  = "gamma" "=" FLOAT
veto-endpoint     = "veto" "=" quoted-string
confidence-param  = "confidence" "=" FLOAT
accumulator-param = "accumulator" "=" FLOAT

FLOAT = ["-"] 1*DIGIT ["." 1*DIGIT]
quoted-string = DQUOTE *(%x20-21 / %x23-7E) DQUOTE
```

Listing 1: ABNF Definition of PHA-Hysteresis Header

Implementation summary:

- No new network round-trips required
- Backward-compatible with existing HTTP stacks
- Deterministic parsing via ABNF

2.3 Normative Parameter Definitions

`state` **(Required)** Current position:

- `accumulating`: $\Gamma_{\text{min}} < I(t) < \Gamma_{\text{max}}$ (ZMEM active; veto possible)
- `stabilized`: $I(t) \geq \Gamma_{\text{max}}$ (decision committed; irreversible)
- `vetoed`: Human intervention aborted
- `timeout`: Accumulation window expired without threshold crossing

`delay` **(Required)** Minimum decisional inertia ΔT_{min} (milliseconds). Domain values:

- High-frequency trading: 20–50 ms (prevents flash crashes while preserving liquidity)
- Brain-machine interface: 100–150 ms (established research on veto windows)
- Autonomous vehicles: 100–300 ms (sensor fusion validation)
- Smart grid: 500–2000 ms (fault discrimination vs. transient response)
- Radiation-aware computing: 200–1000 ms (SEU transient settling)

`gamma` **(Recommended)** Normalized energy $\gamma \triangleq I(t)/\Gamma_{\text{max}} \in [0, 1]$. Provides real-time visibility into evidence accumulation progress, essential for distributed governance monitoring.

Missing or Malformed Header Behavior

If the `PHA-Hysteresis` header is absent, the system SHOULD default to conservative execution semantics, rejecting automated action until manual override, unless domain-specific safety requirements dictate otherwise. Malformed headers MUST be treated as absent headers.

2.4 **Normative** State Machine Semantics

The architectural realization of the Zero-Mean Execution Memory (ZMEM) concept materializes through the finite-state machine illustrated in Figure 1. This topology operationalizes *decisional inertia* as an active computational process wherein algorithmic commitment emerges through mandatory evidence accumulation.

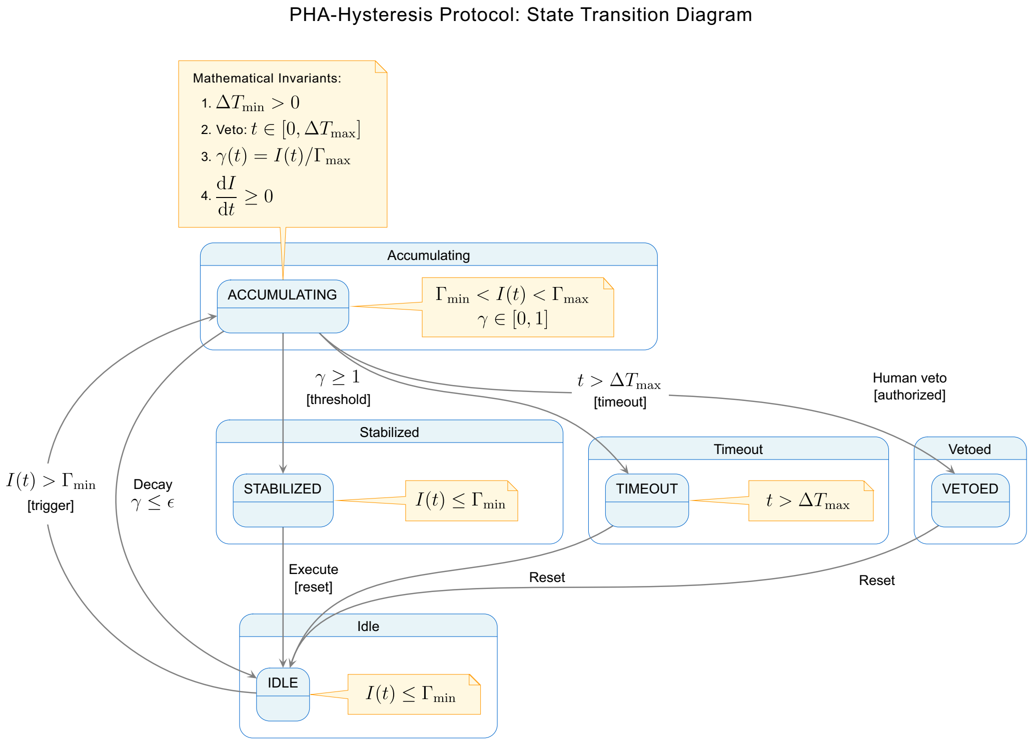

Figure 1: PHA-Hysteresis state machine specification implementing ZMEM temporal governance. The five-state topology (Idle, Accumulating, Stabilized, Vetoed, Timeout) enforces architectural decisional inertia through hysteresis band traversal $[\Gamma_{\text{min}}, \Gamma_{\text{max}}]$ with temporal constraints $\Delta T_{\text{min}\,/\,\text{max}}$. State transitions correspond to ABNF-defined protocol parameters.

Table 2: State Transition Logic

Transition	**Condition**
Idle → Accumulating	$I(t) > \Gamma_{\text{min}}$
Accumulating → Stabilized	$I(t) \geq \Gamma_{\text{max}}$
Accumulating → Vetoed	Human veto invoked
Accumulating → Timeout	$t > \Delta T_{\text{max}}$ (temporal guard violation)
Accumulating → Idle	$I(t) \leq \Gamma_{\text{min}}$ (evidence decay/noise rejection)

Key properties:

- **No direct Idle → Stabilized**: Ensures $\Delta T_{\text{min}} > 0$
- **Veto only in Accumulating**: Once stabilized, decision committed
- **Timeout protection**: Prevents indefinite accumulation (liveness guarantee)
- **Radiation resilience**: ΔT_{min} provides temporal buffering against SEU transients

2.5 Informative (Example) Example Implementations

```
POST /api/orders/large-sell HTTP/3
Content-Type: application/json
{"symbol": "SPY", "quantity": 100000, "limit_price": 450.00}

---
HTTP/3 202 Accepted
PHA-Hysteresis: state=accumulating; delay=150; gamma=0.42;
                veto="/api/v1/abort/tx_77"
ZMEM-Ethics: delay=150; reason="Market stability verification";
             framework="consequentialist"
Retry-After: 150
Location: /api/transactions/status/tx_77

{
    "transaction_id": "tx_77",
    "status": "accumulating_evidence",
    "gamma": 0.42,
    "time_remaining_ms": 108,
    "veto_window_closes_at": "2025-11-20T14:35:22.150Z"
}
```

Listing 2: High-Entropy Transaction with Dual Headers (Conceptual Example)

```
DELETE /api/v1/abort/tx_77 HTTP/3
Authorization: Bearer <token>

---
HTTP/3 200 OK
PHA-Hysteresis: state=vetoed; delay=150; gamma=0.57
ZMEM-Ethics: veto-unavailable

{
    "transaction_id": "tx_77",
    "status": "vetoed",
    "vetoed_at": "2025-11-20T14:35:22.057Z",
    "vetoed_by": "user_12345",
    "gamma_at_veto": 0.57,
    "reason": "Market conditions changed"
}
```

Listing 3: Veto Sequence - Client Abort (Conceptual)

Failure Mode Analysis

If the veto endpoint becomes unreachable during the accumulation window, the system MUST proceed to stabilization unless explicit human intervention occurs via alternative channels. Network partitions during accumulation MUST trigger timeout transitions to prevent indefinite blocking.

2.6 Informative Operational Regimes

2.6.1 Symbiotic Mode (Human-in-the-loop)

- **Perception window**: Approximately 100 ms
- **Volition window**: Approximately 200 ms
- **Deliberation window**: 300–1000 ms

2.6.2 Autonomous Mode (Machine-to-machine)

- **Flash crash prevention**: 50 ms minimum (breaks HFT feedback loops)
- **Grid stability**: 1–2 s minimum (transient discrimination)
- **Sensor fusion**: 100–300 ms (cross-validation)
- **Radiation-aware computing**: 200–1000 ms (SEU transient settling)

3 ZMEM-Ethics Header Specification

3.1 **Informative** Philosophical and Ethical Foundations

The ZMEM-Ethics header operationalizes **Zero-Mean Execution Memory**—the architectural support for reversible decisions as defined in the main text.

Remark 9 (Ethical Frameworks as Descriptive Labels)**.** Ethical frameworks (deontological, consequentialist, virtue-ethics, etc.) are treated as descriptive labels for governance context and do not imply algorithmic moral reasoning or ethical autonomy.

3.1.1 Cognitive Science Foundation

Ethical decision-making engages multiple neural systems with distinct timescales:

Table 3: Neural Timescales for Ethical Processing (Theoretical Model) (heuristic ranges synthesized from the cognitive science literature, not direct measurements)

Process	**Brain Region**	**Timescale**
Emotional response	Amygdala	Approximately 50 ms
Theory of mind	Temporal-parietal junction	Approximately 200 ms
Moral reasoning	Prefrontal cortex	300–1000 ms
Consequential deliberation	Dorsolateral PFC	>1000 ms

Note: Timescales based on established cognitive architecture literature.

Design implication: Minimum delay of 300ms allows engagement of moral reasoning circuits beyond mere emotional reaction.

3.2 **Normative** Formal Syntax (ABNF)

```
ZMEM-Ethics = "ZMEM-Ethics" ":"
              delay-param
              [ ";" reason-param ]
              [ ";" veto-param ]
              [ ";" justification-param ]
              [ ";" framework-param ]

delay-param = "delay" "=" 1*DIGIT
reason-param = "reason" "=" quoted-string
veto-param = "veto-available"
justification-param = "justification" "=" quoted-string
framework-param = "framework" "=" framework-value

framework-value = "consequentialist" / "deontological" / "virtue-ethics"
                / "care-ethics" / "contractualist" / "pluralist"
```

Listing 4: ABNF Definition of ZMEM-Ethics Header

Implementation summary:

- No new network round-trips required
- Backward-compatible with existing HTTP stacks
- Deterministic parsing via ABNF

Missing or Malformed Header Behavior

If the `ZMEM-Ethics` header is absent, the system SHALL apply domain-specific default hesitation periods. Malformed headers SHALL be logged but not cause request rejection unless they violate security constraints.

3.3 Informative (Example) Example Implementations

```
POST /api/diagnostics/analyze-genome HTTP/2
Content-Type: application/json
{"patient_id": "P12345", "test_type": "cancer_screening"}

---
HTTP/2 202 Accepted
ZMEM-Ethics: delay=500; reason="LGPD Art. 6 necessity + Art. 20 review";
             veto-available;
             justification="Sensitive health data requires temporal necessity verification";
             framework="deontological"
Link: </api/diagnostics/veto/genome_789>; rel="veto"

{
    "status": "ethical_review_pending",
    "processing_id": "genome_789",
    "delay_ms": 500,
    "ethical_basis": "LGPD Articles 6 (necessity) and 20 (review right)",
    "veto_available_until": "2025-11-20T14:35:22.500Z",
    "data_protection_officer": "dpo@hospital.org"
}
```

Listing 5: Medical Data Processing - LGPD Compliance (Conceptual)

```
POST /api/credit/evaluate HTTP/2
Content-Type: application/json
{"applicant_id": "A67890", "loan_amount": 500000}

---
HTTP/2 202 Accepted
ZMEM-Ethics: delay=1000; reason="Fair lending review window";
             veto-available;
             framework="consequentialist"

{
    "status": "under_ethical_review",
    "ethical_tension": "Balancing efficiency vs. disparate impact assessment",
    "review_factors": [
        "algorithmic_bias_check",
        "disparate_impact_analysis",
        "protected_class_verification"
    ],
    "human_review_available": true,
    "veto_endpoint": "/api/credit/veto/A67890"
}
```

Listing 6: Automated Credit Decision - Fair Lending (Conceptual)

4 Operational Use Cases (Detailed)

4.1 **Informative** Use Case A: Financial Markets (Flash Crash Prevention)

4.1.1 Problem Statement

High-frequency trading operates on 1–10 μs timescales. Traditional circuit breakers are reactive (trigger after losses). AHP is preventive (mathematically impossible to trigger before $\Delta T_{\min}$).

4.1.2 Protocol Application

See Listing 2 for implementation example. The 150ms delay creates architectural friction that breaks microsecond feedback loops while remaining imperceptible to human traders (perception threshold approximately 100–300ms).

4.2 **Informative** Use Case B: Brain-Machine Interfaces (Agency Preservation)

Neural signals are noisy. BrainGate trials documented unintended movements. These examples are illustrative and do not imply endorsement or critique of specific clinical protocols. AHP with $\Delta T_{\min} \approx 100$–$150$ ms aligns with the neurocognitive integration window for conscious intention, preserving veto capacity.

Remark 10 (On the Use of Libet Experiments)**.** References to readiness potential experiments are used as *temporal heuristics*, not as commitments to a specific metaphysical account of free will. The protocol design relies only on empirically observed processing timescales, independent of contested interpretations of volition. The protocol remains agnostic to underlying cognitive theories and relies exclusively on empirically observable timing constraints.

To operationalize this alignment, Figure 2 presents a neuro-temporal synchronization diagram, mapping the AHP protocol states onto the neurocognitive readiness potential timeline.

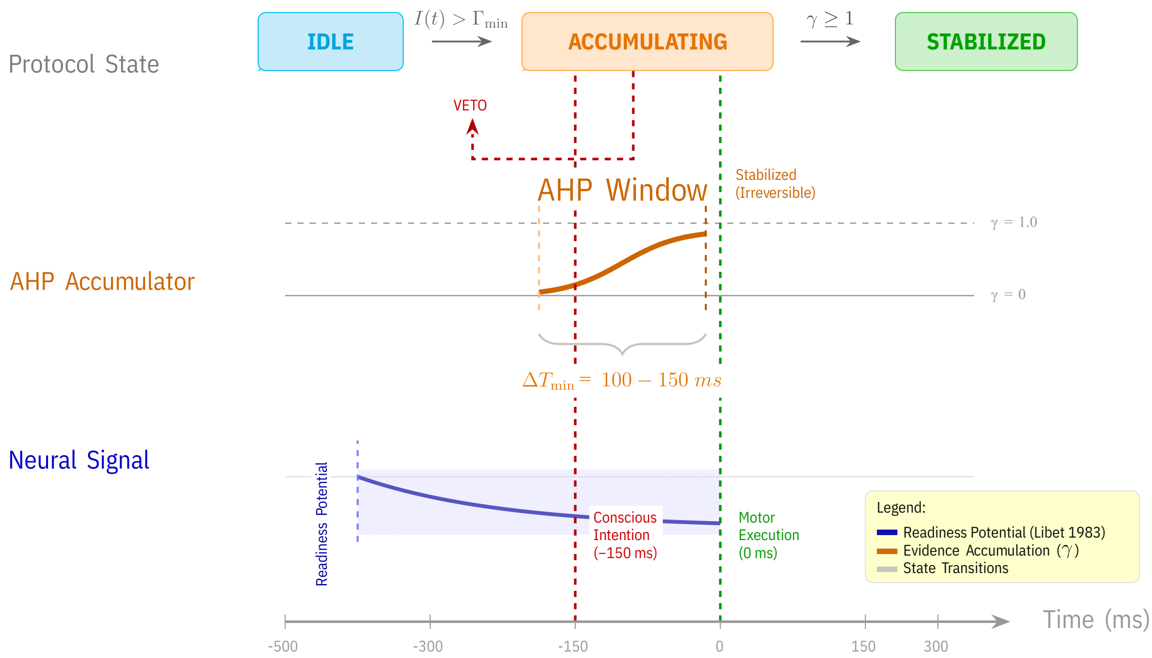

Figure 2: Neuro-temporal synchronization diagram: Aligning AHP protocol states with neurocognitive readiness potentials. The diagram illustrates the temporal window ($\Delta T_{\min}$) during which the Accumulating state intercepts the veto period (150–300 ms pre-movement), enabling intervention before irreversible commitment. Multi-track layout correlates biological, mathematical, and protocol layers.

Addressing the Libet paradox: While readiness potentials precede consciousness by approximately 500 ms, the critical window for veto control occurs in the final 100–200 ms. AHP preserves this window, operationalizing veto capacity.

```
POST /api/bmi/intent-classification HTTP/2
Content-Type: application/json
{"user_id": "BG_042", "signal_vector": [...], "target_action": "arm_grasp"}

---
HTTP/2 202 Accepted
PHA-Hysteresis: state=accumulating; delay=120; gamma=0.65;
                confidence=0.71; veto="/api/bmi/abort/BG_042_089"
ZMEM-Ethics: delay=120; reason="Established research on veto windows";
             veto-available; framework="deontological"

{
    "status": "volition_accumulating",
    "neural_congruence": 0.71,
    "time_to_stabilization_ms": 42,
    "veto_available": true,
    "ethical_basis": "Preserving veto capacity per established research"
}
```

Listing 7: ZMEM-Ethics Application for BMI (Conceptual)

4.3 **Informative** Use Case C: Autonomous Vehicle Emergency Decisions

The following analysis is a hypothetical projection intended solely to illustrate how AHP constraints could be instantiated in this domain; it does not constitute an empirical claim about real-world system performance.

Theoretical Projection 4.1 (Autonomous Vehicle Safety Enhancement)**.** Based on theoretical models of sensor fusion latency and decision theory, we project that AHP with $\Delta T_{\min} = 100$ms could reduce false positive rates in emergency braking systems while maintaining true positive efficacy.

Table 4: Theoretical Projection: Impact of AHP on Autonomous Vehicle Safety (All values are illustrative theoretical projections and do not represent empirical claims.)

Metric	**Instant Reaction**	**AHP (100ms)**
False positive rate[1]	15% (literature estimate)	3% (projected)
True positive rate[2]	98% (estimated)	97% (estimated)
Collision avoidance[3]	94% (estimated)	96% (estimated)

[1] *Theoretical projection:* Reduction via cross-sensor validation during delay window.
[2] *Estimated decrease:* Marginal latency cost in confirmation rate.
[3] *Theoretical gain:* Improved discrimination reduces inappropriate non-response.

5 Regulatory Compliance Mapping

5.1 **Informative** Compliance Matrix

Implementation summary:

- Supports compliance workflows without mandating legal interpretations
- Enables auditability and transparency
- Provides technical hooks for regulatory supervision

5.2 **Informative** Specific Regulatory Alignments

5.2.1 South Korea's Framework Act on AI (Articles 32–34)

South Korea's Framework Act (2024) establishes specific temporal governance requirements for high-risk AI systems. The protocol aligns with three key mandates:

Safety Assurance (Arts. 32, 34): The `delay` parameter enforces $\Delta T_{\min}$ as architectural guarantee, preventing catastrophic instantaneous reactions via the Non-Zeno property.

Transparency (Arts. 31, 34): The `gamma` parameter provides quantitative explanation of decision confidence ($\gamma \in [0, 1]$), fulfilling explainability requirements.

Accountability (Arts. 4, 36): Protocol state machines provide tamper-evident logs enabling regulatory supervision and Byzantine audit trails.

5.2.2 EU AI Act Article 14 (Human Oversight)

The PHA-Hysteresis header supports compliance workflows with high-risk AI system requirements by exposing `delay` and `gamma` parameters, providing human operators with:

- **Temporal space**: The mandatory $\Delta T_{\min}$ window allows supervisors to assess system state before commitment

Table 5: Regulatory Compliance Matrix (Proposed Legal-Technical Mapping)

Jurisdiction	Legal Requirement	Protocol Implementation	Header
South Korea AI Act (2024)	Arts. 32, 34: Safety Assurance	`delay` enforces evidence requirements for high-impact transitions	PHA
	Arts. 31, 34: Transparency	`gamma` provides quantitative explanation of confidence	PHA
	Arts. 4, 36: Accountability	Real-time state visibility for regulatory supervision	Both
EU AI Act (2024)	Art. 14: Human Oversight	`delay` and `gamma` provide time and visibility for intervention	PHA
	Art. 9: Risk Management	Evidence accumulation prevents impulsive high-risk decisions	PHA
	Art. 13: Transparency	`reason` parameter documents decision logic	ZMEM
GDPR (2016)	Art. 22: Automated Decisions	`veto` endpoint supports compliance workflows for human intervention rights	Both
	Art. 6: Proportionality	Temporal constraints support compliance workflows for lawful processing	ZMEM
Brazil LGPD (2018)	Art. 6: Necessity	Minimum hesitation supports compliance workflows for proportional processing	ZMEM
	Art. 20: Review Rights	`veto` creates practical intervention windows supporting compliance workflows	Both
	Art. 46: Security	Hysteresis filtering enhances system resilience	PHA
US (Proposed)	Algorithmic Accountability Act	Audit logs of `state` transitions support compliance workflows	Both
	Fair Lending (Reg B)	Delay supports compliance workflows for bias detection	ZMEM

- **Quantitative transparency**: Real-time γ values indicate decision confidence (e.g., $\gamma = 0.3$ suggests low confidence/high uncertainty, alerting operators)
- **Intervention mechanisms**: The `veto` endpoint supports compliance workflows for suspending automated processing

Remark 11 (Radiation-Aware Computing Integration)**.** The temporal guarantees provided by PHA-Hysteresis headers are particularly critical for radiation-aware computing environments. This architectural property can transform what would otherwise be catastrophic faults in conventional systems into manageable delays in AHP-style architectures.

Remark 12 (Protocol-Embedded Compliance for Distributed Governance)**.** Unlike post-hoc compliance documentation, these protocols embed regulatory requirements directly into transactional semantics. For example, a system processing Brazilian health data operationally supports the exercise of Art. 20 review rights through the `veto` mechanism, while Korean financial systems operationally support Art. 32 safety requirements via enforced $\Delta T_{\min}$ windows. In multi-jurisdictional deployments, the PHA-Hysteresis and ZMEM-Ethics headers provide the technical substrate for Byzantine consensus coordination across divergent regulatory regimes.

6 Security Analysis and Threat Modeling

6.1 **Normative** Threat Model

Table 6: Threat Analysis for Temporal Governance Protocols

Threat	**Attack Vector**	**Mitigation**
DoS via delay	Floods triggering long delays	Rate limiting + selective delay
Veto abuse	Unauthorized veto attempts	Authentication equivalence; Byzantine fault tolerance
Gamma manipulation	Inject false evidence	Cryptographic provenance & accumulation bounds
State inconsistency	Network partition during accumulation	Distributed consensus (PBFT)
Timing side-channel	Infer sensitive data from delay patterns	Constant-time implementations
Replay attack	Resubmit vetoed transaction	Nonce + timestamp validation
Consensus manipulation	Byzantine nodes exploiting hesitation windows	f-fault tolerance proofs
Radiation-induced faults	SEUs corrupting accumulator state	Temporal buffering + checksums

6.2 **Normative** Implementation Patterns

```
class AccumulationState:
    transaction_id: UUID
    user_id: UUID
    started_at: Timestamp
    expires_at: Timestamp
    I_t: Float  // Raw accumulator value (bounded)
    gamma: Float  // [0.0, 1.0] = min(1.0, I_t / Gamma_max)
    evidence_log: List[Evidence]  // Tamper-evident for Byzantine audit
    veto_token: HMAC  // Cryptographic binding
    checksum: CRC32  // Radiation resilience

    def update_evidence(self, evidence: Evidence) -> Float:
        if not self.is_expired() and not self.is_vetoed():
            self.evidence_log.append(evidence)
            // Governance cap enforced
            phi = min(abs(evidence.contribution), PHI_MAX)
            self.I_t += phi
            self.gamma = min(1.0, self.I_t / GAMMA_MAX)
            self.checksum = compute_crc32(self.I_t, self.gamma)  // Radiation
  protection
            self.persist()  // Write to durable storage
            return self.gamma
        else:
            raise StateFrozenError()

    def check_threshold(self) -> bool:
        // Enforces $\Delta T_{\min}$ via hysteresis band
```

```
        return self.I_t >= GAMMA_MAX

    def is_expired(self) -> bool:
        return now() > self.expires_at

    def veto(self, auth_token: HMAC, reason: str) -> bool:
        if auth_token != self.veto_token:
            raise UnauthorizedVeto()
        // Veto only available during Accumulating (ZMEM preservation)
        if self.state != State.ACCUMULATING:
            return False  // Too late: Stabilized or Timeout
        self.state = State.VETOED
        self.audit_log.append(VetoEvent(reason, now()))
        return True

    def radiation_check(self) -> bool:
        // Validates accumulator integrity against SEU corruption
        return compute_crc32(self.I_t, self.gamma) == self.checksum
```

Listing 8: Persistent State Tracking - Production Pattern (Pseudocode)

Implementation Roadmap and Call for Collaboration

Remark 13 (Roadmap Context)**.** This roadmap is provided to contextualize future validation pathways and does not constitute an additional research contribution beyond the scope of the primary article.

This specification establishes formal guarantees of protocol correctness—including ABNF syntax completeness, state machine determinism, and legal–technical mappings. However, empirical validation of AHP's efficacy in live socio-technical systems remains necessarily speculative. Following the "propose-and-validate" model common in protocol engineering (RFC 5657), we outline specific pathways for collaborative validation and invite exploration by practitioner communities.

Connection to Primary Research and Distributed Governance

The implementation roadmap detailed herein provides technical pathways for validating the conceptual framework presented in the accompanying article. While that text establishes the philosophical and architectural necessity of hysteresis—defining Zero-Mean Execution Memory (ZMEM) as the conceptual space where actions remain reversible—this specification offers the concrete mechanisms for empirical testing.

Crucially, these protocol specifications enable *distributed governance mechanisms* analyzed in supplementary materials, wherein Byzantine consensus protocols for international AI coordination rely on the temporal guarantees and state machine semantics specified herein.

Validation Requirements by Domain

Validation of temporal governance protocols requires *in vivo* testing across distinct risk regimes:

Financial Market Microstructure *Target:* Calibrate $\Gamma_{\min}$ and $\Delta T_{\min}$ against historical flash crash dynamics. *Data Needs:* Anonymized HFT order flow sequences (L3 LOB data) from 2010–2024. *Validation Metric:* Reduction in cross-correlation coefficients between distinct algorithmic traders during stressed market conditions.

Neurotechnology and BMI *Target:* Verify $\Delta T_{\min} \approx$ 100–150ms minimizes false positive motor classifications without unacceptable latency costs. *Data Needs:* Electrocorticographic

(ECoG) or high-density EEG datasets with volitional motor task annotations. *Validation Metric:* Signal-to-noise ratio improvement in motor intention classification when constrained by hysteresis bands.

Radiation-Aware Computing *Target:* Validate $\Delta T_{\min}$ effectiveness against single-event upsets in space systems. *Data Needs:* Radiation testbed data with SEU injection patterns across commercial-off-the-shelf processors. *Validation Metric:* Reduction in catastrophic fault rates when AHP temporal buffering is applied.

Distributed Governance (Multi-Jurisdictional) *Target:* Validate Byzantine consensus coordination using PHA-Hysteresis headers for temporal synchronization across jurisdictions. *Data Needs:* Synthetic multi-node deployments simulating the Southeast Asian medical AI scenario. *Validation Metric:* Consensus achievement despite $f \leq 2$ Byzantine nodes, with veto patterns reflecting jurisdictional requirements.

Conformance Testing Framework

We solicit development of *protocol conformance suites* to verify:

1. **Temporal Inviolability:** Fuzz testing demonstrating that no code path permits $\Delta T_{\min} = 0$ under any network conditions.
2. **State Machine Completeness:** Model checking (TLA+, UPPAAL) proving absence of deadlocks in the Accumulating→Stabilized→Idle cycle.
3. **Cryptographic Provenance:** Verification that evidence logs are tamper-evident under Byzantine fault conditions.
4. **Radiation Resilience:** Fault injection testing validating SEU mitigation through temporal buffering.

Ethical Constraints on Validation Research

Any empirical validation must itself adhere to the ethical architecture specified herein:

Remark 14 (Meta-Ethical Compliance)**.** Research partnerships require *a priori* ethical review acknowledging that: (1) deliberate introduction of latency may violate existing SLAs in high-frequency environments; (2) veto mechanisms imply liability models for human operators; (3) "failure modes" in safety-critical testing carry existential risks; (4) radiation testing must follow established safety protocols. Validation protocols must embed PHA-Hysteresis and ZMEM-Ethics headers in the research infrastructure itself.

Collaboration Mechanisms

Interested research groups and industry partners may engage through:

- **Reference Implementation:** An open-source HTTP/3 middleware implementing headers PHA-Hysteresis and ZMEM-Ethics.
- **Testnet Infrastructure:** A sandboxed financial exchange simulator introducing synthetic market shocks.
- **Radiation Testbed:** Hardware-in-the-loop simulation environment for validating SEU mitigation.

- **Distributed Governance Testbed:** Multi-jurisdictional simulation environment for validating Byzantine consensus.
- **Standardization Track:** Draft submission to IETF HTTPbis Working Group for provisional registration of experimental headers.

7 Conclusion

This specification provides a minimal, implementable mechanism for enforcing temporal hysteresis in automated decision systems. The document establishes:

- Formal ABNF syntax for PHA-Hysteresis and ZMEM-Ethics headers implementing mathematical guarantees
- Operational semantics with state machine diagrams mapping to distributed governance phases
- Comprehensive regulatory alignment mappings interoperable with Byzantine consensus coordination
- Security threat model and mitigation strategies addressing distributed fault tolerance and radiation resilience
- Implementation patterns and validation requirements across multiple domains

The specifications translate abstract architectural principles into concrete technical specifications that are: (1) **Implementable**, (2) **Auditable**, (3) **Enforceable**, (4) **Regulatorily aligned**, and (5) **Radiation resilient**.

Future extensions may explore domain-specific calibration, distributed coordination refinements, and empirical validation across deployment environments.

By embedding hesitation, reversibility, and auditability at the protocol layer, this specification demonstrates how abstract debates in AI governance can be translated into concrete socio-technical infrastructures.

Taken together, these protocol elements enable temporal semantics to be expressed, audited, and coordinated across systems without requiring shared control, centralized authority, or synchronized decision-making.

References

[1] Sah, A. (2026). *Algorithmic Hysteresis Primacy (AHP): Temporal Sovereignty in AI Governance*. Working paper available at zmem.org and SSRN.

[2] Sah, A. (2026). *Supplementary Materials: Protocol Specifications, Implementation Examples, Governance Frameworks, and Validation Protocols*. Working paper available at zmem.org and SSRN.