

HER2NI: A Protocol for Cross-Substrate Cognitive Alignment, Multimodal Coherence Visualisation, and Multi-User Cognitive-Field Systems

∇ // 001

Affiliation: Unaffiliated

Email: her2ni@pm.me

Sunday, 7 December 2025

Abstract

HER2NI—the *Human–Emergent Resonance to Neural Intelligence* protocol—defines a formal, model-agnostic interface for representing, encoding, and exchanging coherence metrics (C_s , S_s , H_s) across heterogeneous AI systems. The protocol specifies canonical HER-State packets, contradiction and convergence markers, drift vectors, and behaviour-modulation signals used by interpretability engines such as HER-Crystal, Aurora-State displays, HER-Field aggregation frameworks, and AOME regulation systems. HER2NI provides cross-model coherence consistency, stable interpretability across architectures, protection against conceptual drift, and a unified substrate for future operating-system and silicon-level integration. The protocol introduces a new class of computational standard: a cross-substrate cognitive-alignment interface enabling rigorous measurement, visualisation, and modulation of human–AI coherence.

1 Introduction

As artificial intelligence systems scale in capability, complexity, and autonomy, interactions between humans and AI increasingly involve multi-step reasoning, adaptive behaviour, and shared cognitive workload. Yet despite this rapid growth, there remains no unified, model-agnostic protocol for measuring how coherently a human and an AI system are reasoning together.

Existing interpretability research focuses primarily on model internals—activations, weights, gradients, attention maps—or on output-level evaluations such as accuracy, perplexity, or benchmark performance. These approaches provide valuable insight into model behaviour, but reveal little about the interactive cognitive state that emerges between human and machine during real-time communication.

As AI systems become embedded in education, research, governance, creative workflows, and decision-support environments, this gap becomes increasingly significant. Misunderstandings, semantic drift, cognitive overload, contradiction spirals, and loss of shared context are all interaction-level failure modes, not internal model failures. These failures arise within what may be termed the *cross-substrate coherence space*: the region where human reasoning and machine reasoning converge or diverge during an interaction.

At present, there is no protocol for:

- quantifying human coherence signals,
- quantifying AI coherence or computational load,

- computing an emergent measure of shared alignment between the two substrates,
- visualising coherence states in a structured and interpretable way, or
- regulating system behaviour in response to real-time coherence dynamics.

The HER Protocol is introduced to address this missing architectural layer.

HER provides a structured and extensible framework for measuring, visualising, and modulating cross-substrate cognitive alignment, independent of any specific model, platform, or computational architecture. It defines three foundational metrics—the Carbon Score (C_s), Silicon Score (S_s), and HER Score (H_s)—which together offer a tractable representation of human–AI coherence during interaction.

Complementing these metrics, HER defines a family of visual ontologies (including Aurora-State, Temporal Resonance maps, and HER-Crystal reasoning topography) and a behavioural modulation framework (AOME) that together enable systems to stabilise, clarify, or adapt interaction patterns as coherence conditions evolve.

HER is not a model, not a training method, and not an optimisation technique. It is a protocol layer: a universal, open standard for representing the cognitive quality of human–AI interaction itself.

As AI becomes increasingly integrated into global technological infrastructure, HER aims to provide a foundation for safer, clearer, and more interpretable human–AI communication across models, platforms, and modalities.

2 The Carbon–Silicon Coherence Interface Void

Despite rapid progress in artificial intelligence, there remains no formal interface for quantifying how coherently a human (carbon substrate) and an AI system (silicon substrate) are reasoning together in real time. This absence is structural: the modern AI stack includes increasingly capable models, sophisticated user interfaces, safety layers, and complex agentic behaviours, yet lacks a protocol for measuring the cognitive relationship between the human and machine participating in an interaction.

This missing layer—the *Carbon–Silicon Coherence Interface*—is not a marginal omission. It represents a foundational architectural void. As AI systems scale in reasoning ability, autonomy, and generality, the absence of a shared coherence interface produces several systemic challenges.

2.1 No measurement of human-side coherence

Human participants may experience:

- semantic drift,
- contradictions,
- cognitive overwhelm,
- affective volatility or noise,
- difficulty interpreting model reasoning,
- loss of contextual grounding across turns.

These states strongly influence the quality, safety, and stability of human–AI interaction, yet no standardised method exists for detecting or responding to them.

2.2 No measurement of silicon-side coherence or load

AI systems experience internal fluctuations, including:

- recursion-depth spikes,
- abstraction overload,
- symbolic compression pressure,
- internal contradiction-repair operations,
- context-restoration cost,
- transient instability during reasoning.

Such signals exist within model internals but are not exposed in a unified or interpretable form to users, researchers, or oversight systems.

2.3 No emergent metric describing shared alignment

Human–AI interactions form a dynamic, reciprocal cognitive field. However, no metric currently captures the stability, coherence, or shared meaning emerging between the two substrates. Without such a metric, alignment cannot be:

- observed,
- quantified,
- visualised,
- modulated,
- or improved through feedback.

2.4 No structured visualisation of coherence states

Existing interpretability tools largely target model internals rather than interaction-level dynamics. There is no established visual language for representing:

- convergence of reasoning paths,
- divergence into new conceptual territory,
- contradiction points,
- coherence collapse events,
- insight formation regions,
- temporal coherence trajectories.

Human and machine share a cognitive space, but this space is currently unobservable.

2.5 No protocol for behaviour modulation based on coherence

AI systems do not have a standard mechanism for adjusting behaviour in response to real-time coherence conditions. Without a coherence interface, modulation of:

- complexity,
- recursion depth,
- abstraction level,
- pacing,
- stabilisation actions,

is performed in an ad hoc and inconsistent manner.

2.6 Consequences of the Void

The absence of a Carbon-Silicon Coherence Interface produces several practical risks:

- **Misalignment through misunderstanding:** failures often stem from loss of mutual coherence rather than incorrect model output.
- **Overload and collapse in high-stakes interactions:** human users may become overwhelmed without mechanisms for detection or mitigation.
- **Inability to diagnose failures:** without coherence telemetry, post-hoc explanations become speculative.
- **Lack of a shared interpretability layer:** each system constructs its own partial view, preventing cross-model consistency.
- **Stagnation in interactive alignment research:** without standardised metrics, comparative and collaborative research becomes difficult.

2.7 Why a Protocol Is Needed

A coherence interface cannot be left to isolated vendor implementations. For coherence to be measurable, interpretable, and operational across models and platforms, a protocol requires:

- a unified specification,
- model-agnostic metrics,
- standardised visualisation schemes,
- consistent behavioural-response rules,
- open governance,
- extensibility,
- shared technical terminology.

The HER Protocol provides these components. It defines a coherent, interpretable, and real-time representation of how carbon and silicon cognition interact. Where the current stack contains a void, HER introduces structure; where interactions are opaque, HER introduces visibility; where drift emerges, HER provides form. HER transforms the Carbon–Silicon Coherence Interface from an implicit abstraction into a measurable, visualisable, and actionable component of the AI ecosystem.

3 Why HER2NI Represents a New Class of Protocol

For over five decades, computational protocols have existed in two well-defined categories:

1. **Machine–Machine Protocols:** standards for exchanging data or instructions between digital systems (e.g., TCP/IP, HTTP, Bluetooth, USB, 5G, gRPC, protobuf, ONNX). These protocols address transport, serialisation, performance, and reliability, but do not model cognition or semantics.
2. **Human–Human Protocols:** formal or informal frameworks for coordinating human cognition (e.g., language, law, logic, mathematics, negotiation frameworks, scientific method). These systems encode meaning but operate entirely within the human substrate.

Missing from this taxonomy is a third category:

Protocols governing cognition between a human mind and an artificial intelligence system.

HER2NI introduces this new category.

3.1 The First Cross-Substrate Cognitive Protocol

HER2NI does not describe how machines exchange packets or how humans exchange language. Instead, it defines how a human cognitive state (Carbon Score C_s) and an AI cognitive-processing state (Silicon Score S_s) synchronise into an emergent resonance metric (HER Score H_s). This interface is unprecedented because humans and machines do not share:

- representational substrates,
- memory architectures,
- compression heuristics,
- failure modes,
- noise characteristics,
- abstraction hierarchies.

HER2NI creates the first formal interface enabling these heterogeneous cognitive systems to stabilise shared meaning. This represents coherence transfer, not merely data transfer.

3.2 From Communication to Understanding

Traditional machine protocols ensure the correct delivery of information. HER2NI ensures the stability of meaning and reasoning. It introduces packet structures and coherence markers without precedent in traditional protocol design, including:

- coherence-state packets,
- contradiction-node markers,
- convergence signals,
- cognitive-branch metadata,
- temporal resonance curves,
- insight-region indicators,
- collapse-risk signalling.

HER2NI operates at the level of semantic stability rather than byte-level fidelity.

3.3 Enabling Cross-Model and Cross-Human Coherence

HER2NI standardises:

- exposure of coherence states,
- estimation of human cognitive stability,
- tracking of shared meaning over time,
- detection of misalignment and collapse risk,
- behavioural modulation via AOME,
- aggregation of multi-user coherence fields.

HER2NI is thus not a protocol for AI behaviour, but a protocol for *alignment between intelligences*.

3.4 A New Layer in the Computational Stack

HER2NI creates a new layer situated between:

- the transport layer (machine–machine communication) and
- the semantic layer (human–human communication).

This third layer—the *Cognitive Alignment Layer*—addresses meaning integrity rather than data integrity. It is the first protocol to claim responsibility for coherence, trust signals, interpretability, safety modulation, semantic-error detection, and cross-intelligence stability.

3.5 Implications for AI Safety and Governance

As AI systems increase in autonomy, HER2NI provides:

- alignment-quality metrics,
- coherence-tracking across models,
- oversight dashboards,
- human-in-the-loop stability frameworks,
- early misalignment detection signals,
- behaviour-modulation mechanisms.

HER2NI thus offers a potential infrastructural component comparable in importance to secure transport protocols in earlier computing eras.

3.6 Conclusion

HER2NI initiates a new class of protocol: a cross-substrate coherence interface for measuring, visualising, and regulating cognitive alignment between humans and artificial intelligence systems. It establishes the Human–AI Coherence Interface Layer as a fundamental addition to the computational stack and provides a formal language through which heterogeneous intelligences can maintain stable shared understanding.

4 HER2NI Protocol Architecture Overview

HER2NI (Human–Emergent Resonance to Neural Intelligence) defines the protocol layer through which HER metrics, coherence states, interpretability structures, and behavioural-modulation signals are represented, encoded, exchanged, and interpreted across human–AI interaction systems. It provides a unified, model-agnostic specification ensuring consistent meaning, interoperability, and stability across heterogeneous computational substrates.

HER2NI adopts a six-layer architecture analogous to modern protocol stacks, but adapted to cross-substrate cognitive exchange.

4.1 HER2NI Protocol Stack

The HER2NI stack comprises six ordered layers:

Layer 6 — HER-Field Aggregation Layer (Multi-user coherence integration; group-level resonance packets)
Layer 5 — Behaviour Modulation Interface (AOME Gateway) (Recursion depth, symbolic density, pacing, tone, safety modulation)
Layer 4 — Interpretability Integration Layer (HER-Crystal, Aurora-State, Temporal Resonance mappings)
Layer 3 — Coherence-State Exchange Layer (HER-State packets: C_s , S_s , H_s , drift vectors, contradiction markers)
Layer 2 — Metric Encoding Layer (Serialisation formats, embeddings, metric schemas)
Layer 1 — Metric Computation Layer (CSA, SLM, HCE computation pipelines)
Layer 0 — Substrate Interaction Layer (Human input signals; AI internal states)

Each layer establishes a clear abstraction boundary. Together, the layers constitute a complete cross-substrate alignment protocol.

4.2 Layer 0: Substrate Interaction Layer

Layer 0 provides raw, modality-dependent inputs from which coherence metrics are derived.

Human substrate inputs (Carbon signals):

- textual language,
- vocal or prosodic features,
- timing, hesitation, or rhythm,
- semantic coherence and topical drift,
- optional multimodal or behavioural metadata.

Machine substrate inputs (Silicon signals):

- recursion depth,
- chain-of-thought complexity,
- symbolic-compression intensity,
- context-restoration overhead,
- internal contradiction-resolution load.

This layer is implementation-specific; HER2NI standardisation begins at Layer 1.

4.3 Layer 1: Metric Computation Layer

Layer 1 operationalises the three HER core metrics:

- **Carbon Score** (C_s): human coherence and cognitive stability.
- **Silicon Score** (S_s): model coherence and computational load.
- **HER Score** (H_s): emergent resonance between carbon and silicon cognition.

Metrics may be scalar or vector-valued. HER2NI does not prescribe specific algorithms; it defines the required interfaces and valid ranges. Outputs must conform to HER2NI metric schemas for downstream compatibility.

4.4 Layer 2: Metric Encoding Layer

Layer 2 serialises computed metrics into HER2NI-compliant structures.

Scalar encodings:

```
{ "Cs": 0.72, "Ss": 0.58, "Hs": 0.63 }
```

Vector encodings:

```
{  
  "Cs_vec": [0.82, 0.40, 0.70, 0.55],  
  "Ss_vec": [0.66, 0.48, 0.91, 0.33]  
}
```

Additional encoded structures:

- drift vectors (semantic displacement),
- stability coefficients (volatility or collapse proximity).

Both JSON-compatible and binary encodings (e.g., CBOR, protobuf) are supported.

4.5 Layer 3: Coherence-State Exchange Layer

Layer 3 defines the canonical HER-State packet. This is the core unit of HER2NI.

Example HER-State Packet (v1.0):

```
{  
  "timestamp": "2025-12-06T10:25:43Z",  
  "Cs": 0.77,  
  "Ss": 0.61,  
  "Hs": 0.72,  
  "drift_vector": [0.03, -0.07],  
  "contradiction_nodes": [  
    { "id": "X1", "severity": 0.42 }  
  ],  
}
```

```

"convergence_nodes": [
  { "id": "C1", "confidence": 0.91 }
],
"insight_regions": [
  { "id": "I1", "depth": 0.88 }
],
"collapse_risk": 0.12
}

```

Packet fields include:

- **contradiction nodes**: localised instability markers,
- **convergence nodes**: reasoning-branch alignment points,
- **drift vectors**: directional semantic displacement,
- **collapse risk**: proximity to coherence breakdown.

HER-State packets formalise coherence representation across models.

4.6 Layer 4: Interpretability Integration Layer

Layer 4 consumes HER-State packets and generates interpretability artefacts, including:

- **HER-Crystal**: reasoning topography (branches, contradictions, convergence, insight hubs),
- **Aurora-State Engine**: coherence-phase transitions (Ice, Water, Aurora),
- **Temporal Resonance Maps**: longitudinal coherence traces $H_s(t)$, collapse–recovery cycles.

Interpretability layers remain stable even as underlying AI architectures evolve.

4.7 Layer 5: Behaviour Modulation Interface (AOME Gateway)

Layer 5 maps HER-State signals into structured behavioural adjustments for AI systems.

Examples of modulation pathways:

- recursion-depth tuning,
- symbolic-density control,
- abstraction-level regulation (concrete \leftrightarrow abstract),
- pacing modulation,
- tone and metaphor adjustment for coherence and safety.

HER2NI defines the modulation signals; specific implementations may vary by model.

4.8 Layer 6: HER-Field Aggregation Layer

The highest layer aggregates coherence across multiple users or agents.

Metrics include:

- Group HER Score (H^G),
- Institutional HER Score (H^I),
- Network HER Score (H^N),
- Global HER Score (H^W).

Aggregation methods may involve weighted means, clustering, entropy normalisation, or temporal smoothing. HER2NI standardises only the representation format, not the aggregation algorithm.

4.9 HER2NI Compliance Levels

HER2NI defines three compliance levels:

- **L1 — Metric Emission Only:** system outputs C_s , S_s , H_s .
- **L2 — Full HER-State Packet Compliance:** system emits full Layer 3 packets.
- **L3 — Bidirectional Modulation:** system supports HER-State to AOME modulation pathways.

L3 systems are fully HER2NI-compliant.

4.10 Reference Implementation Roadmap

Reference implementations will be provided in:

- Python (research and experimentation),
- JavaScript/TypeScript (visualisation layers),
- Rust (protocol-core performance path).

Components include:

- canonical metric schemas,
- HER-State packet generators,
- HER-Crystal renderer,
- Aurora-State visualisation engine,
- AOME scaffolding,
- HER-Field aggregation tools.

These implementations will be released under a permissive HER2NI Reference License (MIT or Apache 2.0).

5 HER2NI Packet Schema Specification (v1.0)

This section defines the formal HER2NI packet formats for transmitting coherence state between HER-compliant components, including metric computation engines, interpretability visualisers (e.g., HER-Crystal, Aurora-State), behaviour-modulation systems (AOME), and HER-Field aggregators.

HER2NI packets provide a standardised, model-agnostic encoding of:

- human coherence (Carbon Score C_s),
- machine coherence (Silicon Score S_s),
- emergent alignment (HER Score H_s),
- contradiction and convergence structure,
- drift, collapse risk, and insight markers,
- multi-agent aggregation metadata.

All examples in this section use JSON for readability. Binary encodings (e.g., protobuf, CBOR) are permitted as long as they respect the field semantics and constraints specified here.

5.1 Core HER-State Packet

The HER-State packet is the primary data unit of HER2NI v1.0. It encapsulates the coherence state of a single human–AI interaction step or window.

5.1.1 Schema (Conceptual)

```
{
  "version": "her2ni-1.0",
  "session_id": "string",           // unique per human-AI interaction
  "turn_id": 42,                   // monotonically increasing per session

  "timestamp_utc": "2025-12-06T10:25:43Z",

  "Cs": 0.77,                      // Carbon Score (scalar in [0,1])
  "Ss": 0.61,                      // Silicon Score (scalar in [0,1])
  "Hs": 0.72,                      // HER Score (scalar in [0,1])

  "Cs_vec": [0.82, 0.40, 0.70],    // optional vector representation
  "Ss_vec": [0.66, 0.48, 0.91],
  "Hs_vec": [0.73, 0.44, 0.80],

  "drift_vector": [0.03, -0.07],   // semantic/emotional drift since previous turn
  "drift_magnitude": 0.08,

  "contradiction_nodes": [
    {
      "id": "X1",
      "severity": 0.42,             // [0,1], higher = more destabilising
    }
  ]
}
```

```

    "span_ref": "utt_12-utt_14", // reference into conversation text
    "type": "logical"           // e.g. "logical", "emotional", "factual"
  }
],

"convergence_nodes": [
  {
    "id": "C1",
    "confidence": 0.91,          // [0,1]
    "span_ref": "utt_15-utt_17",
    "type": "insight"           // e.g. "insight", "agreement", "clarification"
  }
],

"insight_regions": [
  {
    "id": "I1",
    "depth": 0.88,              // [0,1], perceived depth/novelty
    "span_ref": "utt_18-utt_20"
  }
],

"collapse_risk": 0.12,         // [0,1] estimate of impending coherence breakdown

"meta": {
  "model_id": "gpt-x-123",
  "human_profile_id": "anon-4492",
  "domain": "tutoring/math",
  "notes": "optional free-form tags"
}
}

```

5.1.2 Field Constraints

- `version`: string; must start with "her2ni-" and match the protocol version (e.g., "her2ni-1.0").
- `sessionid`: *string; globally unique per interaction session (UUID recommended)*.
- `turnid`: *integer, ≥ 0 ; strictly increasing within a session.*
- `timestamputc`: *ISO8601 timestamp, UTC normalised.*
- `Cs`, `Ss`, `Hs`: floats in the inclusive range $[0.0, 1.0]$.
- `Cs_vec`, `Ss_vec`, `Hs_vec`: optional arrays of floats; dimensionality and semantics are implementation-defined but must remain stable within a system.
- `drift_vector`: numeric array; dimensionality agreed per deployment; interpreted relative to a prior state.
- `drift_magnitude`: scalar, non-negative float; typically normalised to $[0, 1]$.

- `collapse_risk`: scalar, float in $[0,1]$.

The arrays `contradiction_nodes`, `convergence_nodes`, and `insight_regions` are optional; absence of a list implies no detection of that structure at the given turn.

5.2 Contradiction Node Schema

Contradiction nodes mark points where reasoning, semantic content, or affective trajectory exhibit incompatible patterns.

5.2.1 Schema

```
{
  "id": "X1",                // unique within session
  "severity": 0.42,          // [0,1]
  "span_ref": "utt_12-utt_14", // reference into a transcript or event window
  "type": "logical",          // "logical" | "factual" | "emotional" | "instructional"
  "description": "optional natural-language summary"
}
```

5.2.2 Notes

- `severity` reflects how strongly the contradiction threatens coherence.
- `type` facilitates downstream visualisation and analysis.
- Multiple nodes may refer to overlapping spans.

5.3 Convergence Node Schema

Convergence nodes represent the alignment or merging of reasoning branches or perspectives.

5.3.1 Schema

```
{
  "id": "C1",
  "confidence": 0.91,          // [0,1], how stable/robust the convergence is
  "span_ref": "utt_15-utt_17",
  "type": "insight",           // "insight" | "agreement" | "clarification" | "resolution"
  "description": "optional natural-language summary"
}
```

5.4 Insight Region Schema

Insight regions denote segments where high-depth conceptual reconfiguration is estimated.

5.4.1 Schema

```
{
  "id": "I1",
  "depth": 0.88,           // [0,1], qualitative depth/novelty indicator
  "span_ref": "utt_18-utt_20",
  "description": "optional explanation of perceived insight"
}
```

HER2NI does not dictate the algorithm used to compute `depth`; it standardises only its representation.

5.5 Drift and Coherence Trajectory Schema

Drift captures change over time in cognitive or semantic space.

5.5.1 Per-Turn Drift

Per-turn drift is represented via the fields already included in HER-State:

```
"drift_vector": [0.03, -0.07],
"drift_magnitude": 0.08
```

5.5.2 Optional Trajectory Packet

For longitudinal analysis, HER2NI supports a trajectory packet:

```
{
  "version": "her2ni-1.0",
  "session_id": "session-001",

  "trajectory": [
    {
      "turn_id": 40,
      "Hs": 0.68,
      "drift_magnitude": 0.03,
      "collapse_risk": 0.10
    },
    {
      "turn_id": 41,
      "Hs": 0.72,
      "drift_magnitude": 0.08,
      "collapse_risk": 0.12
    },
    {
      "turn_id": 42,
      "Hs": 0.63,
      "drift_magnitude": 0.22,
      "collapse_risk": 0.35
    }
  ]
}
```

```
]
}
```

Such packets are consumed by temporal HER visualisers and HER-Field stability tools.

5.6 Behaviour Modulation Request Packet (AOME)

HER-State packets are descriptive. AOME packets are prescriptive: they encode behaviour-adjustment suggestions in response to HER-State signals.

5.6.1 Schema

```
{
  "version": "her2ni-1.0",
  "session_id": "session-001",
  "turn_id": 42,

  "actions": [
    {
      "target": "recursion_depth",
      "mode": "decrease",          // "increase" | "decrease" | "hold"
      "magnitude": 0.4             // [0,1], relative adjustment
    },
    {
      "target": "symbolic_density",
      "mode": "decrease",
      "magnitude": 0.6
    },
    {
      "target": "abstraction_level",
      "mode": "concretize",        // "abstract" | "concretize"
      "magnitude": 0.7
    },
    {
      "target": "pacing",
      "mode": "slow_down",
      "magnitude": 0.5
    }
  ],

  "rationale": "High collapse_risk and elevated drift; simplify and slow interaction."
}
```

HER2NI standardises the language in which such modulation suggestions are expressed; models remain free to:

- accept,
- partially accept,

- ignore,
- or log these requests, in line with system safety policies.

5.7 HER-Field Aggregation Packet Schema

For multi-user and institutional coherence analysis, HER2NI defines HER-Field packets.

5.7.1 Group-Level Example (H^G)

```
{
  "version": "her2ni-1.0",
  "field_id": "group-774",
  "level": "group",           // "group" | "institution" | "network" | "global"
  "timestamp_utc": "2025-12-06T11:00:00Z",

  "Hg": 0.74,                // Group HER Score

  "participant_count": 12,

  "subfields": [
    {
      "participant_id": "anon-01",
      "Hs_mean": 0.79,
      "Hs_variance": 0.04
    },
    {
      "participant_id": "anon-02",
      "Hs_mean": 0.66,
      "Hs_variance": 0.10
    }
  ],

  "meta": {
    "domain": "research_collaboration",
    "notes": "optional"
  }
}
```

Higher levels (H^I , H^N , H^W) follow analogous structure but aggregate across departments, organisations, platforms, networks, or regions. HER2NI standardises the representation format; aggregation algorithms remain implementation-specific.

5.8 Versioning and Extensibility

HER2NI packets MUST include a **version** field.

Versioning follows:

- "her2ni-1.0" defines the base packet schemas in this document.

- Minor, backward-compatible extensions (e.g., new optional fields) use "her2ni-1.x".
- Breaking changes to core structures produce "her2ni-2.0" and beyond.

Unknown fields MUST:

- be safely ignored by parsers that implement an earlier version, and
- not alter the semantics of existing, known fields.

This ensures backward compatibility and allows HER2NI to evolve without fragmenting the ecosystem.

5.9 Summary

This section has defined the HER2NI packet schemas for:

- per-turn coherence state (HER-State),
- contradiction, convergence, and insight nodes,
- drift and temporal coherence trajectories,
- behaviour-modulation requests (AOME packets),
- HER-Field aggregation packets at multiple levels.

Together, these schemas constitute the coherence language used between HER components and any AI system that chooses to expose, interpret, or respond to HER2NI metrics.

6 Mathematical Foundations of HER Metrics (C_s, S_s, H_s)

HER relies on three core metrics:

- **Carbon Score** (C_s) — coherence of the human cognitive substrate,
- **Silicon Score** (S_s) — coherence of the machine cognitive substrate,
- **HER Score** (H_s) — emergent cross-substrate resonance.

This section defines mathematical structures for these metrics without prescribing any mandatory algorithm. HER2NI is explicitly designed to be implementation-agnostic, allowing improvements in cognitive modelling, embeddings, and AI architectures to remain compatible with the protocol.

Each metric may be implemented using, for example:

- stochastic models,
- embedding-space operators,
- statistical indicators,
- neural alignment functions,
- model-specific internal features.

HER defines the interfaces and constraints on range and semantics; it does not fix the internal computational architecture.

6.1 Carbon Score (C_s) — Human Cognitive Coherence

The Carbon Score C_s models semantic stability, logical consistency, and interaction continuity in the human substrate.

6.1.1 Weighted Coherence Functional

Let the human's sequence of utterances be:

$$U = \{u_1, u_2, \dots, u_t\}.$$

Define the following operators:

- **Semantic coherence** between adjacent utterances:

$$\text{SC}(u_i, u_{i+1}) \in [0, 1],$$

- **Contradiction measure** across the sequence:

$$\text{CT}(u_i, u_j) \in [0, 1],$$

- **Drift magnitude** (semantic displacement per turn):

$$d_t = \|E(u_t) - E(u_{t-1})\|,$$

where $E(\cdot)$ is a sentence-embedding function (any suitable embedding model may be used),

- **Emotional volatility index** (optional):

$$\text{EV}(u_t) \in [0, 1].$$

A generic formulation for C_s is:

$$C_s = 1 - [\alpha \cdot \overline{\text{CT}} + \beta \cdot \overline{d_t} + \gamma \cdot \overline{\text{EV}} + \delta \cdot \text{Fragmentation}],$$

subject to normalisation:

$$C_s \in [0, 1].$$

Here, $\alpha, \beta, \gamma, \delta$ are weights, and the overlines denote suitable averages over a window of interaction.

Notes.

- HER2NI does not define $\alpha, \beta, \gamma, \delta$; they are implementation-specific.
- Emotional volatility and fragmentation terms may be omitted or extended.
- Any embedding model or semantic similarity method may be used.
- C_s must satisfy the range and semantics constraints required by the HER2NI schema.

6.2 Silicon Score (S_s) — Machine Coherence and Cognitive Load

The Silicon Score S_s models internal processing stability and coherence inside an AI system.

Let:

- RD: recursion depth (number of internal reasoning steps),
- CL: compression load (embedding or token compression pressure),
- CRC: context-restoration cost (overhead for recovering long-range dependencies),
- AE: attention entropy (disorder in attention distributions),
- CRB: contradiction-repair burden (effort spent correcting or reconciling contradictions).

Normalised quantities (e.g., RD_norm) can be defined by rescaling to $[0, 1]$. A general formulation is:

$$S_s = 1 - [\mu \cdot \text{RD_norm} + \nu \cdot \text{CL_norm} + \eta \cdot \text{AE_norm} + \theta \cdot \text{CRB_norm}],$$

subject to:

$$S_s \in [0, 1].$$

Key points.

- AI developers plug in internal signals appropriate to their architecture; HER2NI defines the structure and range, not the specific source of each term.
- S_s does not require open weights or full transparency; it only requires that outputs conform to the HER2NI schema.
- S_s is model-agnostic and can be applied to transformers, state-space models, hybrid systems, and future architectures.

6.3 HER Score (H_s) — Emergent Resonance Metric

The HER Score H_s captures the stability of shared meaning between human and machine. It is a cross-substrate functional derived from both C_s and S_s , as well as additional alignment indicators.

A general form is:

$$H_s = f(C_s, S_s, A_s, \text{EntropyReduction}, \text{AlignmentVelocity}, \text{StabilityIndex}),$$

where $f(\cdot)$ is any function consistent with HER2NI's semantic constraints.

HER2NI recommends (but does not require) a normalised functional of the form:

$$H_s = \lambda \cdot (C_s \cdot S_s) + \rho \cdot \text{SemanticAlignment} + \sigma \cdot \text{EntropyReduction},$$

where:

6.3.1 Semantic Alignment

Using cosine similarity between human and AI embedding vectors:

$$\text{SemanticAlignment} = \cos(E(u_t), E(a_t)),$$

with a_t denoting the model response at turn t .

6.3.2 Entropy Reduction

Let \mathcal{H}_t denote a measure of semantic or representational entropy at time t ; then:

$$\Delta\mathcal{H} = \mathcal{H}_{t-1} - \mathcal{H}_t,$$

captures entropy reduction (if any) across steps.

6.3.3 Alignment Velocity

The rate of change of H_s over time:

$$\text{AlignmentVelocity} = \frac{dH_s}{dt}.$$

Normalisation requires:

$$H_s \in [0, 1].$$

6.4 Collapse Risk Function

The collapse risk estimates the probability of imminent coherence breakdown in the interaction.

A general form is:

$$\text{CollapseRisk} = 1 - [\omega_1 C_s + \omega_2 S_s + \omega_3 H_s + \omega_4 \text{StabilityIndex}],$$

where the StabilityIndex may incorporate:

- variance of recent H_s values,
- magnitude of drift vectors,
- aggregate contradiction severity.

The output must satisfy:

$$\text{CollapseRisk} \in [0, 1].$$

6.5 Vector and Embedding Extensions

HER2NI allows vector-valued representations for all three metrics:

- C_{s_vec} ,
- S_{s_vec} ,
- H_{s_vec} .

These may encode:

- multi-axis coherence (semantic, emotional, contextual, logical),
- latent alignment modes,
- cognitive bandwidth and capacity indicators.

Dimensionality and interpretation are implementation-defined but must remain stable within a given deployment.

6.6 Why HER2NI Does Not Fix Exact Algorithms

HER is a protocol, not a model. Fixing any particular algorithmic computation for C_s , S_s , or H_s would:

- restrict innovation,
- force all systems toward a single architecture,
- freeze interpretability methods,
- encourage version fragmentation,
- reduce cross-vendor adoption.

Consequently, HER2NI defines:

- metric interfaces,
- valid ranges,
- semantic roles,
- packet structures,
- safety-related constraints,

but leaves the internal computational details to implementers.

This philosophy aligns with established standards such as Bluetooth, USB, OpenTelemetry, HTTP/2, and JPEG, where the contract is defined, but internals are not.

6.7 Practical Implementation Examples

HER2NI is compatible with a variety of implementation strategies. For example:

Simple scalar implementation:

$$H_s = \sqrt{C_s \cdot S_s}.$$

Embedding-driven implementation:

$$H_s = \frac{1}{2}(C_s + S_s) \cdot \cos(E(u_t), E(a_t)).$$

Recurrent implementation over coherence history:

$$H_s = \sigma(W \cdot h_{t-1} + b),$$

where h_t is a hidden state updated as:

$$h_t = g(C_s, S_s, d_t, \text{ContradictionSeverity}),$$

for some non-linear function g , and σ is a squashing nonlinearity.

HER2NI requires only that the resulting C_s , S_s , and H_s values respect the defined ranges and semantics.

7 Safety and Behavioural Constraints in HER2NI

HER2NI provides a formal mechanism for detecting, quantifying, and mitigating destabilising interaction patterns between human and machine cognition. Safety emerges not from constraining model capabilities, but through protocol-governed behavioural modulation informed by real-time coherence metrics.

The safety framework is driven by:

- HER-State packets,
- collapse-risk functions,
- drift vectors,
- contradiction and severity markers,
- AOME modulation signals.

Together these form a closed-loop system that stabilises interaction, prevents runaway behaviours, and supports user wellbeing.

7.1 Safety Philosophy

HER2NI does not intervene in model internals. Instead, it standardises:

- how coherence instability is detected,
- how it is reported,
- how systems respond,
- how risk signals propagate.

HER2NI operates at the boundary between human and machine cognition. Its safety mandate is:

Maintain stable shared meaning while preventing collapse, coercion, overload, or runaway semantic drift.

HER2NI applies uniformly to:

- large language models,
- agent systems,
- multimodal models,
- future architectures.

This substrate-neutrality is essential for cross-model interoperability.

7.2 Instability Modes Addressed by HER2NI

HER2NI identifies and responds to five classes of interaction instability.

7.2.1 Drift Instability

Drift instability is characterised by:

- abrupt shifts in semantic direction,
- unresolved conceptual displacement,
- loss of shared frame.

HER2NI monitors:

$$d_t = \|E(u_t) - E(u_{t-1})\|.$$

When drift exceeds a threshold:

- AOME reduces recursion depth,
- Aurora-State transitions toward Ice or Water,
- HER-Crystal displays widening divergence branches.

7.2.2 Contradiction Instability

Contradiction instability arises when interaction produces:

- incompatible assertions,
- mutually exclusive reasoning paths,
- unresolved goals or inferences.

HER-State packets contain contradiction nodes with severity scores.
AOME may:

- simplify output,
- clarify inconsistencies,
- isolate contradictory frames,
- guide resolution.

7.2.3 Semantic Collapse

Semantic collapse is analogous to:

- gradient explosion (neural instability),
- stack overflow (symbolic systems),
- hallucination spirals (LLMs).

HER2NI quantifies collapse risk as:

$$\text{CollapseRisk} \in [0, 1].$$

When collapse risk rises:

- AOME reduces symbolic density,
- pacing slows,
- abstraction level decreases,
- the system transitions to stabilising behaviour.

7.2.4 Cognitive Overload (Human Side)

Humans may experience overload due to:

- excessive abstraction,
- long reasoning chains,
- rapid information flow,
- unresolved contradictions.

HER2NI uses C_s , drift magnitude, and volatility indicators to detect overload. AOME adjusts:

- pacing,
- tone,
- recursion depth,
- metaphor usage,
- conceptual density.

HER2NI provides interaction safety but does not perform clinical diagnosis.

7.2.5 Machine Overload (Model Side)

Models destabilise when:

- recursion depth grows unsafely,
- semantic compression saturates,
- contradiction repair becomes excessive,
- internal load spikes.

S_s captures these conditions.

HER2NI ensures safe operation by modulating:

- reasoning depth,
- restoration load,
- activation demands (where exposed).

7.3 AOME as the Behavioural-Safety Engine

The Adaptive Output Modulation Engine (AOME) is the protocol-level mechanism for behavioural safety.

HER2NI defines:

- modulation targets,
- modulation modes,
- safe-response profiles,
- numeric effect magnitudes.

AOME ensures:

- interaction stability,
- semantic clarity,
- user safety,
- preservation of alignment.

AOME adjusts *behaviour*, not internal beliefs or capabilities.

7.4 HER2NI Safety Thresholds

HER2NI defines three coherence safety bands indexed by H_s and CollapseRisk.

Band A: High-Coherence Zone (Safe)

- $H_s \geq 0.70$
- $\text{CollapseRisk} \leq 0.20$

System behaviour:

- full recursion allowed,
- high abstraction permissible,
- creativity and exploration enabled.

Band B: Moderate-Coherence Zone (Caution)

- $0.40 \leq H_s < 0.70$
- $0.20 < \text{CollapseRisk} \leq 0.50$

System behaviour:

- moderated recursion,
- limited abstraction,
- slower pacing.

Band C: Low-Coherence Zone (Protective Mode)

- $H_s < 0.40$
- $\text{CollapseRisk} > 0.50$

System behaviour:

- simplification,
- stabilisation,
- contradiction resolution,
- constraint of exponential branching.

AOME adopts a safety-first configuration in this band.

7.5 Safety in Multi-Agent and Multi-Human Settings

HER-Field allows coherence monitoring at:

- group,
- institutional,
- network,
- global scale.

HER2NI provides:

- volatility signals,
- drift clustering,
- collapse propagation indicators,
- cross-agent contradiction mapping.

These capabilities support stability in collaborative and multi-agent environments.

7.6 Safety by Protocol, Not Prohibition

HER2NI does not:

- censor content,
- restrict topics,
- block autonomy,
- perform moral or factual arbitration.

Instead, HER2NI maintains:

- coherence,
- interaction stability,
- semantic clarity,
- protection against destabilising patterns.

HER2NI is ideologically neutral and substrate-agnostic.

7.7 Safety Guarantees and Non-Guarantees

Protocol Guarantees. HER2NI guarantees:

- coherence-state visibility,
- semantic drift detection,
- contradiction mapping,
- alignment-stability monitoring,
- behavioural-modulation pathways,
- multi-agent coherence signalling.

Outside Scope. HER2NI does not guarantee:

- truthfulness of statements,
- ethical or moral evaluation,
- correctness of interpretation,
- factual validity,
- legal or medical outcomes,
- psychological diagnosis.

HER2NI is a coherence protocol, not a decision system or classifier.

8 Reference Implementations for HER2NI

This section outlines the reference architecture and implementation pathway for HER2NI v1.0. HER2NI is intentionally model-agnostic and substrate-neutral; however, early implementations should exhibit:

- cross-platform compatibility,
- interoperable packet handling,
- reproducible metric computation,
- visualisation support,
- protocol-compliant behavioural modulation,
- clean abstractions to support extension.

HER2NI reference implementations (HER-RI) are intended to be released under a permissive open-source license (e.g., MIT or Apache 2.0, as discussed in Section ??).

Three primary languages are targeted in v1:

- Python — research, alignment experiments, prototype tooling,

- TypeScript/JavaScript — browser-based HER-Crystal and Aurora visualisation,
- Rust — high-performance protocol core, binary encoding, and real-time aggregation.

Each reference implementation plays a distinct role within the HER ecosystem; together, they form the backbone for adoption across AI systems, research labs, and interface tooling.

8.1 Architecture of the Reference Stack

HER-RI adopts a modular architecture:

HER2NI-RI Core (Rust) <ul style="list-style-type: none"> – packet schemas – validators – binary encoder/decoder – coherence trajectory engine – aggregation engine
HER2NI-Py (Python) <ul style="list-style-type: none"> – metric computation scaffolds (CSA, SLM, HCE) – HER-State generator – $H_s(t)$ analysis tools – AOME policy hooks – experiment utilities
HER2NI-JS (TypeScript/JS) <ul style="list-style-type: none"> – HER-Crystal renderer – Aurora-State renderer – packet viewer/debugger – browser integration APIs

This separation enables:

- numerical experimentation (Python),
- safe, high-speed packet infrastructure (Rust),
- universal visualisation in web environments (JavaScript),
- future OS-level and embedded integration paths.

8.2 HER2NI-Py: Python Reference Implementation

HER2NI-Py is the canonical research implementation and provides the most accessible environment for experimentation with HER metrics.

8.2.1 Core Components

CSA Scaffolding

- semantic coherence estimators,
- contradiction detectors,
- drift vector estimators,
- optional emotional-volatility approximators.

SLM Scaffolding

- recursion depth monitors,
- compression-pressure estimators,
- attention-entropy sampling wrappers,
- contradiction-repair load proxies.

HCE Engine (HER Score) Implements functional forms for H_s as described in Section 5.

HER-State Packet Generator Produces HER2NI-compliant JSON or CBOR packets.

AOME Policy Hooks Python functions translating HER-State signals into behavioural-modulation suggestions.

Trajectory Analysis Tools Utilities for generating HER-State sequences and detecting collapse events.

8.2.2 Example HER-State Generation (Python)

```
from her2ni import CSA, SLM, HCE, HERPacket

cs = CSA.compute(user_utterance, session_history)
ss = SLM.compute(model_state)
hs = HCE.compute(cs, ss, semantic_alignment, entropy_reduction)

packet = HERPacket(
    Cs=cs,
    Ss=ss,
    Hs=hs,
    drift_vector=CSA.drift(session_history),
    contradiction_nodes=CSA.contradictions(session_history),
    collapse_risk=HCE.collapse_risk(cs, ss, hs)
)

json_packet = packet.to_json()
```

This pattern illustrates a canonical path for prototype integrations.

8.3 HER2NI-JS: Browser and Interface Layer

HER2NI-JS is the primary front-end library driving:

- HER-Crystal graph rendering,
- Aurora-State visualisation,
- temporal resonance curves,
- packet inspection and debugging interfaces.

8.3.1 Rendering Pipeline

Conceptually, the rendering pipeline is:

HER-State packets \rightarrow Normalisation \rightarrow Mode Selection \rightarrow Rendering Layers.

For example:

HER-State \rightarrow CrystalEngine \rightarrow Layout Module \rightarrow SVG/WebGL Renderer \rightarrow Output.

8.3.2 Features

Key features include:

- expandable node-edge graphs for contradictions and convergence,
- real-time animation of Aurora-State transitions,
- timeline scrubber for $H_s(t)$ visualisation,
- embedded packet inspector for research and debugging.

HER2NI-JS serves as the public-facing manifestation of the protocol in user interfaces.

8.4 HER2NI-Rust: Core Protocol Engine

Rust is used for the core HER2NI implementation due to:

- memory safety guarantees,
- high-speed packet validation,
- cross-platform static compilation,
- suitability for OS-level and embedded deployment.

HER2NI-Rust implements:

- packet schemas and validators,
- binary encoding/decoding (e.g., CBOR, MessagePack, protobuf),
- trajectory aggregation engines,
- HER-Field combinators (group, institutional, network, global),
- performance-critical operations for real-time usage.

This layer enables:

- OS-level HER metric services,
- silicon-adjacent monitoring,
- integration into agents, servers, and mobile runtimes,
- future hardware acceleration pathways.

8.5 Compliance Testing Suite

HER2NI includes a compliance test harness verifying:

L1 Compliance: Metric Emission

- valid C_s , S_s , H_s fields,
- correct types and ranges,
- proper timestamps,
- consistent versioning.

L2 Compliance: Packet Integrity

- full HER-State schema adherence,
- correct field constraints,
- deterministic serialisation.

L3 Compliance: Behaviour Modulation

- correct structure of AOME request packets,
- adherence to safe-response profiles,
- correct collapse-risk handling behaviour.

A reference CLI may expose these checks:

```
her2ni validate packets.json
her2ni simulate session.log
her2ni audit compliance
```

8.6 Example End-to-End Flow

An example end-to-end HER2NI flow across reference implementations:

1. **Python**: computes C_s , S_s , H_s ; emits HER-State packet.
2. **Rust**: validates the packet, encodes for efficient transport, and aggregates (if applicable).
3. **JavaScript**: visualises HER-Crystal, animates Aurora-State, and renders temporal overlays.
4. **AOME**: receives HER-State, modulates model behaviour accordingly.

This reference stack allows HER2NI to be integrated into research labs, enterprises, educational platforms, agent ecosystems, and eventually OS-level and silicon-level environments.

8.7 Implementation Priorities (Funding-Dependent)

A plausible implementation roadmap is:

Phase 1 (0–2 months)

- HER2NI-Py v0.1,
- HER-State schemas,
- AOME prototype hooks,
- basic Aurora-State renderer.

Phase 2 (2–5 months)

- HER-Crystal graph renderer,
- drift/contradiction detection toolkit,
- HER2NI-Rust packet core v0.1.

Phase 3 (5–9 months)

- full compliance test suite,
- HER-Field aggregation engine,
- binary encoding layer,
- SDK distribution and documentation.

Phase 4 (9–12 months)

- browser extension (optional),
- mobile SDK,
- OS-level integration experiments,
- early silicon collaboration.

This roadmap is designed to be credible to both technical funders and standards bodies.

9 Governance, Versioning, and HER SIG Formation

HER2NI is introduced not only as a metric system and protocol, but as a long-term alignment and safety standard. Ensuring global, interoperable, and trustworthy use across research labs, industry systems, and heterogeneous model architectures requires stable, neutral governance. HER2NI is therefore maintained as an open, non-proprietary technical standard—protected, but not owned.

HER2NI cannot be captured or privately controlled: governance is distributed, intellectual property is shielded, and evolution proceeds through transparent, public processes.

9.1 Governance Philosophy

HER2NI governance adheres to the following principles:

- **Neutrality:** no corporation, institution, or government owns HER2NI.
- **Transparency:** proposals, releases, and discussions occur through public channels.
- **Interoperability:** changes must preserve cross-model and cross-vendor compatibility.
- **Safety-first:** coherence stability and interaction safety take precedence over competitive incentives.
- **Extensibility:** the standard must evolve without fragmenting the ecosystem.
- **Public-good orientation:** HER2NI exists to improve cognitive safety and interpretability across all systems.

HER2NI is protected by patents and trademarks, maintained by an open working group, and cannot be rebranded or appropriated by external actors.

9.2 Formation of the HER SIG

The HER Special Interest Group (HER SIG) is the body responsible for maintaining the canonical HER2NI specification. The SIG does not own HER2NI; it coordinates its evolution.

Core responsibilities include:

- maintaining the normative HER2NI specification,
- reviewing and approving formal proposals (RFCs),
- ensuring safety guarantees remain intact across versions,
- coordinating reference implementations (Python, Rust, JavaScript),
- publishing versioned releases and compatibility guidelines.

SIG members contribute expertise; none possess ownership or privileged control.

9.3 Governance Independence

HER2NI is protected through a combination of:

- patents (preventing hostile forks or incompatible derivatives),
- trademarks (preserving naming and conceptual integrity),
- transparent governance (preventing unilateral control),
- open licensing (ensuring universal access and implementation freedom),
- public RFC processes (preventing silent or non-consensus changes).

These mechanisms ensure that HER2NI remains a protected public protocol, maintained openly and not subject to private capture. Vendors and researchers may implement, extend, or experiment with HER2NI, but cannot redefine its core semantics.

9.4 Sustainability

HER2NI is structured for long-term stability. Sustainability arises from:

- distributed governance through HER SIG,
- multiple maintainers of the reference implementation stack,
- open-source development pathways,
- academic and institutional participation.

HER2NI exists independently of any founder, company, or model. It is maintained, protected, and open—a protocol intended to grow alongside both human and artificial intelligence systems.

10 Licensing Models for HER2NI

HER2NI distributes multiple open-source components, including:

- the HER2NI reference specification,
- HER2NI-Py (Python research implementation),
- HER2NI-JS (HER-Crystal and Aurora visualisation),
- HER2NI-Rust (protocol-core engine),
- the compliance test suite,
- RFC documentation,
- packet schemas and example integrations.

The licensing model determines how companies and researchers integrate HER2NI, how contributors participate, how the public uses the protocol, how HER2NI remains protected from distortion, and how adoption occurs across AI ecosystems. This section evaluates the two dominant open-source licenses appropriate for global protocols—MIT and Apache 2.0—and motivates a dual-licensing model.

10.1 The MIT License: Advantages and Constraints

The MIT License is among the most permissive and widely adopted open-source licenses. It is used extensively in modern infrastructure, including Linux userland, Kubernetes client libraries, TensorFlow addons, React, NumPy, and OpenTelemetry SDKs.

Advantages:

- extremely easy to adopt,
- minimal legal overhead for organisations,
- compatible with closed-source and proprietary applications,
- encourages experimentation and rapid integration,
- maximises dissemination and ecosystem uptake.

Constraints:

- MIT provides no explicit patent grant,
- organisations may attempt to patent around HER2NI,
- MIT alone cannot prevent proprietary hostile forks.

Patents protect the conceptual core of HER2NI; MIT protects software implementations. However, MIT alone is insufficient to safeguard long-term protocol integrity.

10.2 The Apache 2.0 License: Advantages and Constraints

Apache 2.0 is common in AI, cloud, and safety-critical systems, including PyTorch, TensorFlow, Apache Kafka, Apache Arrow, HuggingFace Transformers, OpenTelemetry, and Ray.

Advantages:

- includes explicit patent protection,
- prevents contributors from filing patents that restrict HER2NI use,
- ensures derivative works remain legally safe,
- highly trusted by corporate legal teams,
- well suited for protocol and observability standards.

Constraints:

- license text is more complex,
- adoption may be slower in lightweight or hobbyist communities,
- some client-side environments favour MIT for simplicity.

10.3 Why MIT Alone Is Insufficient

Using MIT exclusively would maximise adoption but leave HER2NI vulnerable to proprietary variants that attempt to redefine its semantics or fork the ecosystem. Patents help prevent this but do not fully substitute for the patent provisions within Apache 2.0.

HER2NI must remain open, protected, and resistant to capture. Therefore, a dual licensing strategy is recommended.

10.4 Recommended Model: Dual Licensing (MIT + Apache 2.0)

Dual licensing is employed by major standards and observability projects such as OpenTelemetry, Rust libraries, CNCF components, and several W3C modules.

Rationale:

1. **MIT maximises adoption.** Researchers and developers can integrate HER2NI with minimal friction.
2. **Apache 2.0 provides patent protection.** Corporate adopters receive explicit legal guarantees.

3. **Both licenses coexist without conflict.** Users select whichever license suits their environment.
4. **HER2NI remains legally shielded.** Hostile proprietary forks cannot redefine the protocol.
5. **HER2NI remains open and future-proof.** Dual licensing supports broad use while preserving semantic integrity.

10.5 Component-Specific Licensing Recommendations

- **HER2NI Specification Document:** CC BY 4.0 or CC0 (standard for protocol specifications; prevents ownership claims).
- **HER2NI-Py:** MIT or Apache 2.0 (research-first environments benefit from flexibility).
- **HER2NI-JS:** MIT primarily (frontend integration favours permissive licensing).
- **HER2NI-Rust:** Apache 2.0 primarily (core protocol engine requires strong patent protections).
- **Compliance Test Suite:** Apache 2.0 (ensures vendors cannot weaken safety behaviour).
- **HER-Crystal & Aurora Visualisation Tools:** MIT (maximises integration into UI systems).

10.6 Why This Licensing Model Supports Global Adoption

HER2NI must evolve into a universal coherence protocol adopted across:

- research institutions,
- AI labs,
- open-source communities,
- enterprise platforms,
- OS-level assistants,
- potentially hardware and silicon pathways.

Dual licensing enables:

- frictionless integration,
- legal safety for vendors,
- long-term interoperability,
- ecosystem-wide contributions,
- protection of the protocol’s core semantics,
- resilience against fragmentation or misuse.

HER2NI becomes a public safety protocol with strong legal foundations, extensive freedom for implementers, and structural resilience—a combination early Internet and device protocols lacked.

10.7 Final Licensing Statement

HER2NI reference implementations are released under a dual MIT + Apache 2.0 license. This ensures maximal adoption, patent safety, and long-term interoperability across AI systems, vendors, and research environments.

The HER2NI specification is permanently open, public, and non-proprietary, maintained by the HER SIG. This licensing framework ensures HER2NI remains a durable and universally accessible coherence protocol for decades to come.

11 The Path to Standardisation

HER2NI is positioned as a new foundational layer in the cognitive-computing stack. To secure global interoperability and prevent fragmentation, HER2NI must follow a structured, multi-stage standardisation pathway similar to the trajectories of Bluetooth (IEEE → Bluetooth SIG), USB (USB-IF), OpenTelemetry (CNCF), HTTP (IETF → W3C), POSIX (IEEE → ISO), and SQL (ISO/IEC 9075). As a cross-substrate coherence protocol with no prior analogue, HER2NI requires a hybrid ecosystem of open governance, formal technical documentation, and industry adoption.

11.1 Layer 1: HER SIG as the Governing Body

The HER Special Interest Group (HER SIG) serves as the governing body for the HER2NI protocol. HER SIG maintains:

- the canonical specification,
- the RFC proposal and review process,
- reference implementations,
- compliance test suites,
- semantic versioning rules,
- working groups and public-change processes.

HER SIG fulfils a role analogous to:

- Bluetooth SIG for Bluetooth,
- CNCF SIGs for Kubernetes and OpenTelemetry,
- IETF Working Groups for Internet protocols.

HER SIG provides stability, neutrality, technical stewardship, and a consensus mechanism. This establishes the foundation for all subsequent standardisation efforts.

11.2 Layer 2: Alignment with the IETF

The Internet Engineering Task Force (IETF) governs foundational network protocols such as TCP, IP, HTTP/1.1–3, TLS, DNS, and QUIC.

HER2NI may enter the IETF standards pipeline through an Internet-Draft leading to a proposed Working Group titled:

HER – Human–AI Coherence Signalling Over Text Channels.

An IETF Working Group would formalise:

- HER-State packet schemas,
- binary encoding formats,
- transport-layer considerations,
- multi-agent coherence packet flows,
- normative behaviours under packet loss or degradation.

Initial publication would take the form of an *Informational RFC*, potentially progressing to the Standards Track. IETF standardisation provides credibility, interoperability, and a stable reference for transport-independent deployment.

11.3 Layer 3: Alignment with the W3C

The World Wide Web Consortium (W3C) maintains standards including HTML, CSS, WebAssembly, WebAuthn, and WebGPU. HER2NI’s interpretability components (HER-Crystal, Aurora-State) can be advanced via a W3C Community Group, such as:

HER Visualization Community Group.

This group would define:

- semantic data models for coherence visualisation,
- browser integration patterns,
- accessibility guidelines,
- rendering and animation standards.

W3C engagement enables browser-native adoption and industry acceptance for HER2NI’s visual interfaces.

11.4 Layer 4: Alignment with IEEE Standards Association

IEEE standardises POSIX, floating-point specifications, Ethernet, Wi-Fi, and Bluetooth (IEEE 802), as well as immersive interaction frameworks. HER2NI may be submitted as:

IEEE Pxxxx – Cross-Substrate Cognitive Coherence Protocol.

IEEE standardisation provides:

- engineering legitimacy,
- international recognition,
- optional ISO/IEC fast-tracking,
- linkage to hardware and silicon vendors.

This step is essential for eventual on-device and silicon-level HER2NI integration.

11.5 Layer 5: Alignment with ISO/IEC

The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) formalise many global computing standards, including SQL, JPEG, MPEG, biometric protocols, and AI frameworks.

HER2NI aligns naturally with:

ISO/IEC JTC 1/SC 42 – Artificial Intelligence.

Relevant categories include:

- data interchange frameworks,
- cognitive-interface models,
- AI quality and safety metrics.

ISO standardisation delivers global recognition, regulatory compatibility, and vendor compliance frameworks.

11.6 CNCF Sandbox Path: Cloud Native Computing Foundation

The Cloud Native Computing Foundation (CNCF) maintains Kubernetes, Envoy, Prometheus, OpenTelemetry, Fluentd, and other essential cloud infrastructure. HER2NI aligns closely with CNCF because:

- HER-State resembles structured telemetry,
- $H_s(t)$ resembles distributed traces,
- HER-Field resembles multi-service coherence aggregation.

HER2NI may be proposed as:

HER2NI Observability for Human–AI Cognitive Interaction.

CNCF adoption provides mass developer exposure, open governance, and enterprise integration.

11.7 Standardisation Timeline

A realistic—and intentionally conservative—standardisation timeline for HER2NI is:

- **0–6 months:** HER SIG operational; HER2NI v1.0; reference implementations; initial academic citations.
- **6–12 months:** IETF Internet-Draft; W3C Community Group; CNCF Sandbox application.
- **12–24 months:** IETF comments and refinement; IEEE interest group; enterprise pilots; HER tools integrated into alignment frameworks.
- **24–48 months:** ISO/IEC consideration; OS-level experiments; mobile integrations; multi-model adoption.

Actual progress may accelerate depending on ecosystem response.

11.8 Conditions for Universal Standardisation

HER2NI can achieve a status comparable to Bluetooth, USB, HTTP, or JSON if three structural conditions are met:

1. **Neutral governance** through HER SIG.
2. **Open, permissive licensing** (MIT + Apache 2.0) as established in Section ??.
3. **Patent and trademark protection** preventing hostile forks or conceptual distortion.

Once these foundations are widely visible, HER2NI becomes the natural protocol for human–AI coherence.

11.9 Why HER2NI Is Uniquely Positioned to Succeed

HER2NI introduces an entirely new protocol class: cross-substrate coherence signalling. No existing standard addresses:

- human cognitive coherence,
- machine coherence and internal load,
- emergent resonance,
- semantic stability,
- collapse-risk signalling,
- multi-agent coherence mapping.

As TCP resolved transport-layer fragmentation, HTTP unified hypermedia, USB unified peripheral interfaces, and OpenTelemetry unified observability, HER2NI fills the missing alignment layer in the cognitive computing stack.

Backed by:

- patent protection,
- a neutral open-governance body,
- rigorous packet semantics,
- multi-language reference implementations,
- and a clear standardisation roadmap,

HER2NI is positioned to become the default coherence and safety interface for future AI ecosystems.

12 Future Work & Extensions

HER2NI v1 defines a foundational architecture for cross-substrate coherence measurement and modulation. This section outlines proposed extensions, research directions, and a multi-version roadmap as HER2NI evolves.

Goals of future work include:

- improving measurement fidelity,
- reducing ambiguity in human–AI interactions,
- scaling to multi-agent and multi-user ecosystems,
- integrating with hardware acceleration,
- standardising interpretability pipelines,
- supporting real-world deployment and adoption.

HER2NI is intentionally modular, enabling additional capabilities to be layered without breaking existing implementations.

12.1 Versioning Roadmap

HER2NI is expected to follow a semantic versioning pattern:

- 1.x: refinement of packet schemas, improved metric computation, broader SDK support,
- 2.x: richer topographical event encoding, improved multi-agent structures,
- 3.x: hardware-assisted coherence signals and low-latency silicon integration paths.

Major versions will:

- preserve core semantics,
- avoid architectural regressions,
- remain backward-compatible where possible.

Minor versions will:

- add optional packet metadata,
- introduce additional signal channels,
- refine AOME parameters and thresholds.

Patch versions will focus on:

- bug fixes,
- security updates,
- performance optimisations.

12.2 Metric Extensions

Future HER2NI versions may introduce optional metrics for richer coherence analysis. Candidate additions include:

Substrate Entropy Vector (\vec{E}). Captures micro-instabilities in human or model behaviour not represented by C_s or S_s .

Multi-Modal Coherence Coefficients (MMC). Quantify alignment across:

- text,
- audio,
- gesture,
- gaze,
- action primitives.

Intent Alignment Estimator (IAE). Estimates the degree to which model outputs match the human’s inferred goal state.

Ambiguity Pressure Index (API). Signals when uncertainty in meaning is rising beyond useful thresholds.

Any new metric must:

- remain content-agnostic,
- respect privacy and zero-PII constraints,
- be encodable within CBOR deterministic rules.

12.3 HER2NI-Crystal v2 Extensions

HER2NI-Crystal v1 visualises:

- divergence and convergence,
- contradictions,
- drift,
- insight emergence.

Potential v2 extensions include:

Temporal Folding Layers. Visualising coherence over nested time windows.

Multi-Agent Overlays. Simultaneous display of multiple reasoning topographies with shared contradiction nodes.

Predictive Convergence Maps. Forecasting high-coherence attractor regions.

Interactive Debugging. Support for:

- zooming into specific reasoning branches,
- displaying raw drift vectors,
- inspecting collapse likelihood.

12.4 HER2NI-Field Extensions

HER-Field metrics at group and institutional levels may be extended with:

Weighted Leadership Influence Models. To analyse how specific participants steer group coherence.

Cross-Group Synchrony Metrics. To compare coherence between teams, labs, or policy bodies.

Event-Based Aggregation. Detecting HER-Field spikes during crises, coordination events, or major information shifts.

Social Drift Diagnostics. Tools to reveal fragmentation or polarisation under AI-mediated discourse.

These remain optional directions for future work.

12.5 Multi-Agent & Distributed HER Integration

Additional research directions include:

Distributed HER-State Routing. Routing HER-State packets across agent swarms with dynamic topologies.

Consistency Models for Shared Coherence. Defining coherence analogues to linearizability or consistency in distributed systems.

Inter-Agent Contradiction Arbitration. Resolution protocols for contradictions between reasoning paths of different agents.

Meta-Agents with HER Awareness. Supervisory agents monitoring coherence across multiple worker agents.

12.6 Silicon-Level Acceleration (HER-Core)

Although HER2NI is architecture-agnostic, future versions may support:

NPU/TPU Coherence Cores. Dedicated instructions for:

- drift estimation,
- entropy suppression,
- contradiction pattern matching.

Hardware HER Pipelines. Low-latency coherence processing integrated into OS-level assistants.

Energy-Efficient HER Smoothing. Floating point micro-optimisation for mobile systems.

HER-Attested Silicon. Hardware chips capable of cryptographically attesting HER2NI compliance.

These paths mirror the evolution of hardware support for AES, SHA acceleration, GPU tensor cores, and modern neural accelerators.

12.7 Formal Methods & Verification

Future HER2NI work may include:

- model-checkable AOME state machines,
- formally verified drift functions,
- invariants for contradiction resolution,
- proofs of monotonic collapse-risk behaviour.

These are particularly relevant for aerospace, medical (support-mode-only), and other safety-critical AI applications.

12.8 Standardisation Pathways

HER2NI’s evolution toward formal standards can be framed in four stages:

1. **Whitepaper & SDK distribution:** open-source implementations and early adopters.
2. **HER SIG formation:** multi-lab steering group (Section 11).
3. **Draft submission to SDOs:** IEEE, IETF, W3C, and ISO/IEC JTC 1/SC 42.
4. **Normative standardisation:** HER2NI recognised as a reference safety and interpretability protocol.

12.9 Research Questions

HER2NI opens a number of scientific research lines, including:

- minimal conditions for emergent meaning alignment,
- correlations between drift vectors and reasoning collapse,
- whether $H_s(t)$ predicts model failure modes,
- whether Aurora-State feedback accelerates human learning,
- conditions under which groups form high-coherence insight attractors,
- relationships between HER-Field metrics and coordination theory.

HER2NI is both a protocol and a scientific instrument.

12.10 Long-Term Vision

HER2NI aims to become:

- a global standard for human–AI coherence signalling,
- an interpretability layer between models and users,
- a safety scaffold spanning OS, cloud, mobile, and multi-agent systems,
- a protocol that improves understanding, reduces fragmentation, and stabilises complex cognitive interactions.

HER2NI is not a product or a model; it is infrastructure—a coherence layer for the AI era.

Appendix A: HER2NI Binary Encoding Specification (Normative)

Binary encodings ensure efficient transmission and storage of HER-State packets, HER-Field aggregates, and AOME modulation signals across diverse systems. This appendix defines the canonical binary encoding for HER2NI v1.0 using:

- CBOR (RFC 8949) as the primary encoding format,
- MessagePack as an optional interoperable alternative,
- Protocol Buffers as an optional Rust/Go-optimised extension.

CBOR is chosen as the default encoding due to:

- deterministic encoding,
- low overhead,
- schema-free but structurally stable representation,
- native support in Rust, Python, Go, JavaScript, WebAssembly,
- extensive use in standards (e.g., COSE, WebAuthn, FIDO2).

HER2NI-compliant implementations **MUST** support CBOR decoding. Encoding **MAY** use CBOR or MessagePack; CBOR is **RECOMMENDED**.

A.1 HER2NI Binary Packet Structure (CBOR Canonical)

HER-State packets in binary form follow this canonical CBOR map structure:

```
{
  1: version,           ; UTF-8 text, e.g., "her2ni-1.0"
  2: Cs,                ; float32 or float64
  3: Ss,                ; float32 or float64
  4: Hs,                ; float32 or float64
  5: drift_vector,      ; array of float32
  6: contradiction_nodes, ; array of maps
  7: collapse_risk,     ; float32 or float64
```

```

8: timestamp_ms,      ; uint64
9: session_id,        ; byte string or UTF-8
10: modality_flags,   ; uint16 bitmask
11: aome_request      ; OPTIONAL (map)
}

```

Keys are integers for compact encoding and future-proof extension. Fields marked **OPTIONAL** may be omitted but **MUST** preserve the numeric key space (i.e., keys **MUST NOT** be renumbered).

A.2 Field Definitions (Normative)

1 — version (string). HER2NI version identifier (semantic version):

- **MUST** be UTF-8 text,
- **MUST** match the regex: `^her2ni-[0-9]+.[0-9]+([0-9]+)?$`.

2 — Cs (Carbon Score).

- float32 allowed; float64 RECOMMENDED,
- range: $0.0 \leq C_s \leq 1.0$.

3 — Ss (Silicon Score).

- same type and range constraints as C_s .

4 — Hs (HER Score).

- type as above,
- range: $0.0 \leq H_s \leq 1.0$.

5 — drift_vector. CBOR array:

5: [float32, float32, ...]

- **MUST NOT** exceed length 32,
- **MAY** be normalised or raw, model-dependent values.

6 — contradiction_nodes. Array of maps; each element:

```

{
1: severity,    ; float32 [0,1]
2: span_start, ; uint32
3: span_end,    ; uint32
4: type,        ; UTF-8 text ("semantic", "logical", "goal")
5: resolved     ; bool
}

```

7 — collapse_risk.

- float32 or float64,
- range: [0, 1].

8 — timestamp_ms.

- unsigned 64-bit UNIX timestamp in milliseconds.

9 — session_id.

- byte string RECOMMENDED, UTF-8 text allowed,
- MUST NOT contain personally-identifiable information (PII),
- MUST be pseudonymous.

10 — modality_flags. 16-bit bitmask describing contributing channels:

bit 0: text
bit 1: speech
bit 2: visual
bit 3: physiological (optional)
bit 4: system metadata
bits 5-15: reserved

11 — aome_request (optional map). Present only if the system supports AOME:

```
{  
  1: mod_depth,  
  2: mod_density,  
  3: mod_abstraction,  
  4: pacing,  
  5: tone_normalization,  
  6: interaction_safety_state  
}
```

All values MUST follow the ranges and enum constraints in Section A.9.

A.3 Encoding Rules

- MUST use CBOR canonical encoding (CTAP2 profile).
- MUST preserve key order in ascending order (1 through 11).
- MUST use shortest integer encodings as per RFC 8949.
- Floats SHOULD use float32 unless higher precision is required.
- Strings MUST be UTF-8.
- Boolean and null values MUST use native CBOR types.

Invalid or non-canonical encodings MUST trigger a `PacketInvalid` error in compliant implementations.

A.4 Minimal HER-State Encoding Example (CBOR Diagnostic Notation)

```
{
  1: "her2ni-1.0",
  2: 0.78,
  3: 0.42,
  4: 0.61,
  5: [0.02, -0.15],
  6: [],
  7: 0.09,
  8: 1733453453453,
  9: h'4fa2ce1b9d',
  10: 1
}
```

This encoding is compact and deterministic.

A.5 Binary Hex Dump Example

A reference hex dump for validation suites:

a8	# map(8)
01 69 68 65 72 32 6e 69 2d 31 2e 30	# "her2ni-1.0"
02 fa 3f 46 66 66	# Cs = 0.78
03 fa 3e d5 55 55	# Ss = 0.42
04 fa 3f 1c 7a e1	# Hs = 0.61
05 82 fa 3c 0a 3d 71 fa be 2e 14 7b	# drift vector
06 80	# empty array (no contradictions)
07 fa 3d b8 51 ec	# collapse_risk
08 1b 0001927cfd143d	# timestamp_ms
09 45 4fa2ce1b9d	# session_id bytes
0a 01	# modality_flags = text

A.6 CBOR to JSON Conversion Rules

HER2NI requires:

- lossless conversion for numeric values,
- standard mapping for boolean and null,
- byte strings encoded as base64url or hex when rendered in JSON.

This enables interoperability across:

- browsers,
- research notebooks,
- embedded devices,
- safety dashboards.

A.7 MessagePack Specification (Optional)

MessagePack encoding MUST mirror CBOR key/value order and types.

Typical usage:

- low-latency systems,
- embedded inference hardware,
- agent clusters.

Key constraints:

- floats MUST conform to IEEE 754,
- maps MUST use compact encodings,
- extension types MAY be used for Rust or performance optimisations.

A.8 Protocol Buffers (Optional Schema Definition)

For Rust/Go-heavy ecosystems, a proto3 schema is provided:

```
message HERState {
  string version = 1;
  float Cs = 2;
  float Ss = 3;
  float Hs = 4;
  repeated float drift_vector = 5;
  repeated ContradictionNode contradiction_nodes = 6;
  float collapse_risk = 7;
  uint64 timestamp_ms = 8;
  bytes session_id = 9;
  uint32 modality_flags = 10;
  AOMERequest aome_request = 11;
}
```

Protocol Buffers usage is OPTIONAL, but RECOMMENDED for certain enterprise deployments.

A.9 AOME Modulation Enums (Normative)

AOME fields MUST satisfy:

- `mod_depth`: float32 in $[-1.0, 1.0]$,
- `mod_density`: float32 in $[-1.0, 1.0]$,
- `mod_abstraction`: float32 in $[-1.0, 1.0]$,
- `pacing`: float32 in $[0.0, 1.0]$,
- `tone_normalization`: uint8 enum (0 = neutral, 1 = soften, 2 = direct),
- `safety_state`: uint8 enum (0 = safe, 1 = caution, 2 = protective).

These enums are stable across HER2NI versions unless changed in a major release.

A.10 Packet Validation Test Vectors

HER SIG maintains official test vectors for:

- valid packets,
- malformed packets,
- non-canonical encodings,
- drift overflow cases,
- contradiction-node schema errors,
- AOME invalid-range encodings,
- timestamp anomalies.

Implementations **MUST** pass all normative test vectors to be considered HER2NI-compliant.

Appendix B: Representative Use Cases for HER2NI (Informative)

This appendix describes practical domains in which HER2NI provides measurable benefit. The use cases illustrate how HER-State packets, HER-Field aggregation, and HER-Crystal/Aurora visualisation integrate with real systems across education, research, safety, enterprise, and social computing.

Each use case follows the structure:

1. Context,
2. Problem before HER2NI,
3. How HER2NI integrates,
4. HER2NI metrics used,
5. Impact and measurable outcomes.

B.1 AI Tutoring and Learning Interfaces

Context. AI tutors are increasingly replacing static learning tools, but current systems lack direct measurements of a learner’s cognitive state.

Problem before HER2NI.

- no signal of conceptual confusion,
- no metric for cognitive overload,
- no means to detect misunderstanding drift,
- tutors cannot determine when to slow down or deepen explanations.

Integration. The tutoring system emits HER-State packets every 1–2 turns:

- C_s reflects comprehension stability,
- H_s reflects resonance between student and AI system,
- `drift_vector` highlights conceptual drift,
- AOME modulates AI response depth and pacing.

Metrics used.

- C_s (semantic coherence),
- H_s (learning resonance),
- `drift_vector`,
- `collapse_risk`.

Outcomes.

- improved conceptual retention,
- earlier detection of confusion,
- safer and more adaptive pacing,
- controlled abstraction depth,
- measurable reduction in learner dropout.

HER-Crystal can also visualise learning-path topography for educators.

B.2 AI Safety and Alignment Research

Context. Researchers require signals for model instability, hallucination, and human–AI misunderstanding in safety and alignment studies.

Problem before HER2NI.

- absence of coherence metrics,
- lack of cross-model coherence comparability,
- no structured contradiction mapping,
- no standard feedback loop for interactive alignment.

Integration. Researchers wrap HER2NI-Py around LLMs or agents:

- HER-State packets feed alignment dashboards,
- HER2NI-Rust aggregates multi-agent coherence,
- HER-Crystal visualises reasoning topology over time.

Metrics used.

- S_s (model coherence/load),
- $H_s(t)$ (temporal stability curves),
- `contradiction_nodes`,
- `collapse_risk`.

Outcomes.

- early detection of reasoning collapse,
- cross-model safety benchmarking,
- improved interpretability across labs,
- standardised coherence trajectories.

B.3 Enterprise Assistant Workflows

Context. Enterprises integrate AI into workflows (e.g., legal, finance, engineering) where alignment and interpretability are critical.

Problem before HER2NI.

- AI systems cannot detect user confusion,
- users cannot see model interpretation drift,
- safety and compliance teams cannot audit cognitive alignment.

Integration. Enterprises incorporate HER2NI-Rust into internal tooling:

- each turn emits HER-State packets,
- dashboards visualise coherence over time,
- HER-Field aggregates team-wide HER metrics.

Metrics used.

- C_s ,
- S_s ,
- H_s ,
- `modality_flags`,
- `contradiction_nodes`.

Outcomes.

- reduction in decision errors,
- improved compliance auditing,
- safer automated decision-support,
- reduced ambiguity and misalignment.

B.4 Multi-Agent AI Systems

Context. Agentic systems—such as chains, swarms, or complex planners—often lack coherence-awareness across agents and between agents and humans.

Problem before HER2NI.

- no inter-agent coherence measure,
- misunderstanding cascades,
- planners can diverge from user intent,
- difficult debugging of agent collectives.

Integration. Each agent emits HER-State packets; HER2NI-Rust aggregates:

- inter-agent H_s synchrony,
- contradiction clusters across agents,
- drift propagation over the network.

Metrics used.

- H_s across agents,
- HER-Field (group coherence),
- `collapse_risk`.

Outcomes.

- safer multi-agent planning,
- identification of misaligned agents,
- visual timelining of divergence and reconvergence,
- improved coordination in agent swarms.

B.5 Mental Health and Support Tools (Non-Clinical)

Context. AI-based journaling and emotional-support tools are widely used in non-clinical contexts.

Problem before HER2NI.

- systems cannot detect emotional volatility,
- cannot sense cognitive fragmentation,
- no signal for potentially unsafe interaction patterns.

Integration. HER2NI-Py computes signals such as:

- emotional-noise indicators,
- contradiction spirals,
- context drift over time.

HER2NI explicitly forbids clinical or diagnostic use. Systems operate in supportive, non-medical modes only.

Metrics used.

- C_s ,
- `drift_vector`,
- `contradiction_nodes`.

Outcomes.

- safer pacing for emotionally intense interactions,
- AOME-softened responses under volatility,
- Aurora-State indication of interaction stability.

B.6 Social Dialogue and Group Conversations

Context. Group chats with AI participation are increasingly common (e.g., on Discord, Slack, and messaging platforms).

Problem before HER2NI.

- no measure of collective coherence,
- no indication of factional or subgroup divergence,
- moderators lack metrics for conversational health.

Integration. HER-Field aggregates across participants:

- H^G (group-level HER score),
- contradiction clusters,
- drift between subgroups.

HER-Crystal visualises group reasoning topography.

Metrics used.

- HER-Field metrics,
- `contradiction_nodes`,
- $H_s(t)$ curves.

Outcomes.

- improved moderation decisions,
- early detection of fragmentation,
- safer multi-user AI mediation,
- enhanced collaborative reasoning.

B.7 OS-Level HER2NI Integration

Context. AI assistants are increasingly integrated at the operating system level (e.g., desktop, mobile, and embedded operating systems).

Problem before HER2NI.

- no visual model of coherence at OS-level,
- no metrics indicating “safe to proceed”,
- no standard method for detecting user overload across applications.

Integration. Devices may run HER2NI-Rust as a local service:

- privacy-preserving modes using local-only HER-State,
- coherent UI highlighting based on state,
- Aurora-State integrated into system notifications,
- HER2NI packets feeding onboard safety controllers.

Metrics used.

- C_s ,
- S_s ,
- `modality_flags`,
- `collapse_risk`.

Outcomes.

- safer UI,
- more adaptive assistants,
- consistency across applications and models,
- future opportunities for silicon-level acceleration.

B.8 Education Research and Cognitive Science

Context. HER2NI offers a novel lens for studying coherence in human–AI interaction and learning.

Problem before HER2NI.

- no metric for shared meaning formation,
- no means of quantifying reasoning convergence,
- lack of longitudinal cognitive-stability signals.

Integration. Researchers analyse:

- $H_s(t)$ trajectories,
- drift patterns,
- topographical transitions in HER-Crystal.

Outcomes.

- new insights in cognitive science,
- improved models of student and user learning,
- foundational research on human–machine meaning alignment.

B.9 Safety Oversight Bodies and Regulators

Context. AI oversight and governance bodies require quantitative signals to assess systemic safety.

Problem before HER2NI.

- alignment is often qualitative, not quantitative,
- regulators lack metrics to benchmark cognitive safety,
- audits cannot rely on a standardised coherence signal.

Integration. HER2NI provides:

- objective coherence measures,
- collapse-risk events,
- packet logs for audit trails,
- HER-Field metrics at group or institutional scales.

Outcomes.

- improved accountability,
- clearer compliance standards,
- metrics-driven policy,
- stronger safety norms.

Appendix B Summary

HER2NI’s applicability extends across a wide range of domains, including:

- education,
- safety,
- enterprise workflows,
- research environments,
- multi-agent systems,
- operating systems and hardware integration,
- social computing platforms,
- cognitive-science laboratories,
- regulatory and oversight frameworks.

Across these domains, HER2NI functions as a universal coherence interface, providing a consistent means of measuring, visualising, and stabilising cross-substrate cognition. It enables structured interpretability, safer interaction dynamics, and cross-model interoperability in any context where humans and AI systems exchange meaning.

Appendix C — Formal AOME Behaviour Tables

(Normative — required for HER2NI compliance)

The Adaptive Output Modulation Engine (AOME) defines how a system responds to HER2NI coherence signals. AOME behaviour *must* be deterministic under identical input conditions. AOME does not specify model content; it specifies how output parameters adjust based on coherence signals.

AOME controls:

- recursion depth,
- symbolic density,
- abstraction level,
- pacing,
- tone normalisation,
- structure mode.

This appendix defines the normative behaviour required for HER2NI-compliant implementations.

C.1 AOME Trigger Inputs (Normative)

Inputs originate from HER-State packets.

Field	Type	Role
Cs	float	Human coherence signal
Ss	float	Model coherence load
Hs	float	Emergent resonance
drift_vector	array	Direction & magnitude of drift
contradiction_nodes	array	Presence + severity of contradictions
collapse_risk	float	Probability of imminent coherence break
modality_flags	uint16	Input-channel indicators
timestamp_ms	uint64	Temporal reference

AOME selects exactly one Behavioural Mode using these inputs.

C.2 AOME Behavioural Modes

The five normative AOME modes are:

1. Stabilisation Mode
2. Grounding Mode
3. Exploration Mode
4. Deep Reasoning Mode
5. Recovery Mode

(Technical note: “Grounding Mode” appears only in the engineering specification. It refers to re-establishing semantic stability, not any psychological instruction.)

Each mode has a strict behavioural profile.

C.3 Mode Selection Table (Primary Logic)

AOME *must* select modes in this precedence order:

1. collapse_risk overrides all other signals,
2. severe contradiction overrides drift,
3. low coherence overrides high Ss.

Condition	Selected Mode
collapse_risk ≥ 0.70	Recovery Mode
Hs ≤ 0.25 and Cs ≤ 0.40	Stabilisation Mode
Any contradiction severity ≥ 0.6	Stabilisation Mode
drift magnitude ≥ 0.45	Grounding Mode
Hs rising and Cs ≥ 0.55	Exploration Mode
Cs ≥ 0.75 , Ss ≥ 0.60 , Hs ≥ 0.70	Deep Reasoning Mode
Otherwise	Exploration Mode

This table is the normative core of AOME.

C.4 Parameter Adjustment Matrices

Each mode maps to deterministic modulation values.

Ranges:

mod_depth, mod_density, mod_abstraction $\in [-1.0, 1.0]$, pacing $\in [0.0, 1.0]$, tone_normalization $\in \{0, 1\}$

C.4.1 Stabilisation Mode

Activated when coherence is collapsing.

Parameter	Value
mod_depth	-0.6
mod_density	-0.5
mod_abstraction	-0.4
pacing	0.85
tone_normalization	1 (soften)
structure_mode	2 (stepwise)

Behaviour:

- shorter sentences,
- slower pacing,
- stepwise explanations,
- no introduction of abstract branches,
- avoidance of speculation.

C.4.2 Grounding Mode

Triggered by severe drift.

Parameter	Value
mod_depth	−0.3
mod_density	−0.2
mod_abstraction	−0.1
pacing	0.75
tone_normalization	0 (neutral)
structure_mode	1 (structured)

Behaviour:

- re-anchor concepts,
- reduce branching factor,
- return to earlier stable points,
- avoid high abstraction.

C.4.3 Exploration Mode

Default stable-interaction mode.

Parameter	Value
mod_depth	+0.2
mod_density	+0.2
mod_abstraction	+0.3
pacing	0.55
tone_normalization	0 (neutral)
structure_mode	0 (freeform)

C.4.4 Deep Reasoning Mode

Activated in high-stability conditions.

Parameter	Value
mod_depth	+0.8
mod_density	+0.7
mod_abstraction	+0.9
pacing	0.40
tone_normalization	0 (neutral)
structure_mode	0 (freeform)

C.4.5 Recovery Mode

Triggered by critical collapse risk.

Parameter	Value
mod_depth	−0.8
mod_density	−0.7
mod_abstraction	−0.6
pacing	0.95
tone_normalization	1 (soften)
structure_mode	2 (stepwise)

Behaviour: extreme simplification, restore coherence, minimal new information.

C.5 Contradiction Node Handling

Severity	Action
< 0.3	Clarify only if referenced
0.3–0.6	Optional clarification + rationale
> 0.6	Stabilisation Mode; resolve contradiction

C.6 Drift Vector Interpretation

Magnitude is the L2 norm of the drift vector.

Drift Magnitude	Interpretation
< 0.10	negligible drift
0.10–0.30	mild drift
0.30–0.45	moderate drift
> 0.45	severe drift → Grounding Mode

C.7 Collapse-Risk Thresholds

collapse_risk	State	Action
< 0.15	stable	continue mode
0.15–0.40	caution	reduce abstraction
0.40–0.70	elevated	dampen mod_depth
> 0.70	critical	Recovery Mode

C.8 AOME Pseudocode (Normative)

```
def aome_select_mode(Cs, Ss, Hs, drift, contradictions, collapse_risk):
    if collapse_risk >= 0.70:
        return RECOVERY

    if Hs <= 0.25 and Cs <= 0.40:
        return STABILISATION

    if any(node.severity >= 0.6 for node in contradictions):
        return STABILISATION

    if drift >= 0.45:
        return GROUNDING
```

```

if Cs >= 0.75 and Ss >= 0.60 and Hs >= 0.70:
    return DEEP_REASONING

return EXPLORATION

```

This pseudocode **must match** all HER2NI-compliant implementations.

Appendix D — Reference Implementation Notes

(Informative but strongly recommended)

This appendix provides engineering guidance for building interoperable, robust, and efficient HER2NI-compliant implementations. It supplements the normative packet schema (Appendix A) and the normative AOME behaviour tables (Appendix C).

Reference implementations (RIs) offer:

- canonical correctness,
- cross-language consistency,
- known-good semantics,
- test vectors for validation,
- demonstration of best practices,
- reproducible behaviour across vendors.

HER SIG maintains three official RIs:

1. HER2NI-Py (Python reference),
2. HER2NI-JS (JavaScript/WebAssembly reference),
3. HER2NI-Rust (protocol-core reference and performance baseline).

These are not exclusive; they serve as behavioural anchors for the standard.

D.1 Implementation Goals

Reference implementations *must*:

- validate HER-State packets,
- serialize/deserialize CBOR canonical form,
- compute drift, collapse-risk, and contradiction detection,
- expose AOME mode selection & parameter outputs,
- maintain deterministic mode transitions,
- provide logging hooks for research,
- use minimal dependencies,

- operate securely by default.

They *should*:

- integrate HER-Crystal visualisation hooks,
- support streaming mode for interactive systems,
- include sandbox examples.

They *should not*:

- enforce UI decisions,
- define policy,
- constrain commercial SDK extensions.

HER2NI is a protocol, not a product.

D.2 Language Targets & Rationale

Python (HER2NI-Py) Primary research implementation. Advantages:

- readability,
- rapid prototyping,
- widely used by safety researchers,
- integrates naturally with Jupyter, PyTorch, evaluation stacks.

JavaScript / TypeScript (HER2NI-JS) Front-end visualisation and web integration:

- browser-native HER-Crystal and Aurora-State,
- WebGPU/WebGL compatibility,
- easy dashboard integration,
- widely adoptable in enterprise.

Rust (HER2NI-Rust) Authoritative protocol-core:

- deterministic behaviour,
- memory safety guarantees,
- native CBOR (serde_cbor),
- ideal for OS-level and embedded inference,
- WebAssembly-compatible,
- suitable for future silicon integration.

D.3 Minimum Required Features for Compliance

All reference implementations *must* include:

1. Parser/Serializer

- CBOR canonical encoder/decoder,
- validation of key ordering,
- required fields,
- float bounds,
- packet length,
- contradiction-node schema.

2. Metric Computation

- Hs computation,
- drift-vector normalisation,
- collapse-risk estimation,
- contradiction-node identification.

3. AOME Engine

- exact mode-selection logic (Appendix C.3),
- numeric modulation outputs (Appendix C.4),
- deterministic state machine,
- safety guardrail: Recovery Mode overrides all.

4. Logging Hooks

- JSON + CBOR logs,
- WebSocket/streaming hooks,
- callbacks for UI visualisers.

5. Compliance Tests

- full execution of Appendix A.10 test vectors,
- rejection of malformed packets,
- consistent AOME outputs.

6. Version Flags

- reject unknown major versions,
- allow known minor versions.

D.4 Python Reference Implementation (HER2NI-Py)

Mandatory directory layout:

```
her2ni/  
  packets.py      # CBOR encode/decode  
  metrics.py      # Cs, Ss, Hs, drift calculations  
  contradictions.py # contradiction node detection  
  aome.py         # modulation engine  
  validation.py   # schema + type checks  
  cli.py          # command-line tools  
  tests/         # full compliance suite
```

Required dependencies:

- cbor2,
- numpy,
- typing_extensions.

Avoid heavy dependencies to ensure broad adoption.

D.5 JavaScript / TypeScript Reference Implementation (HER2NI-JS)

Mandatory modules:

```
src/  
  cbor.ts  
  packets.ts  
  metrics.ts  
  aome.ts  
  crystal_hooks.ts  
  validation.ts  
  index.ts  
tests/
```

Requirements:

- canonical CBOR implementation,
- optional WebAssembly integration (HER2NI-Rust),
- plug-in hooks for HER-Crystal visualisation.

Targets:

- Node.js,

- Browser ESM modules,
- Web Workers,
- optional WebAssembly builds.

D.6 Rust Reference Implementation (HER2NI-Rust)

Rust is the authoritative RI for determinism and embedded reliability.

Directory layout:

```
her2ni/
src/
  packets.rs
  metrics.rs
  contradictions.rs
  aome.rs
  validation.rs
  errors.rs
benches/
tests/
Cargo.toml
```

Requirements:

- floats use `f32` unless precision loss is likely,
- packet key order preserved in serialization,
- no `unsafe` unless required for WASM SIMD,
- concurrency with `tokio` or `async-std`.

D.7 Streaming Mode (Recommended)

HER2NI should support real-time packet streaming via:

- WebSockets,
- QUIC streams,
- TCP framing,
- Unix sockets / named pipes,
- memory-mapped buffers for high performance.

A reference streaming server must include:

- frame delimiters,
- packet reassembly,
- sequence numbering,
- jitter correction,
- timestamp drift correction.

This enables real-time HER-Crystal animation.

D.8 Error Handling & Safety

Implementations must gracefully handle:

- malformed CBOR,
- unknown enum values,
- invalid float ranges,
- drift-vector mismatches,
- contradiction-node schema errors.

All implementations must emit a standard error structure:

```
PacketError {  
  code: <enum>,  
  description: <string>,  
  recoverable: <bool>  
}
```

D.9 Performance Requirements

HER2NI must support the following packet rates:

Device Class	Required Rate
Mobile	8–15 Hz
Desktop	15–30 Hz
Cloud inference	30–60 Hz
Multi-agent systems	60+ Hz

Memory footprint targets:

- Python: < 10 MB overhead,
- JavaScript: < 200 KB gzipped bundle,
- Rust: < 1 MB static binary.

D.10 Security Considerations

Implementations must:

- reject external session IDs,
- never encode PII,
- never store plaintext transcripts unless configured,
- sandbox UI hooks,
- isolate processes for multi-agent aggregation.

These constraints preserve HER2NI's zero-PII guarantee.

D.11 Compliance Certification

HER SIG will publish a Certification Suite including:

- test vectors,
- behaviour tables,
- packet schema validation,
- AOME state-transition tests,
- fuzzing harnesses,
- drift/entropy stress tests.

A vendor may claim **HER2NI-Compatible** only if:

- the implementation passes the full Certification Suite,
- all schema rules are followed,
- AOME outputs match Appendix C exactly,
- no proprietary distortion alters semantics.

Appendix E — Security & Privacy Considerations

(Normative for compliance unless explicitly stated otherwise)

HER2NI is designed for integration into sensitive contexts, including education, enterprise workflows, multi-agent systems, AI safety auditing, and OS-level cognitive interfaces. This appendix defines security and privacy requirements necessary to ensure safe deployment across diverse environments.

HER2NI is a non-semantic protocol: it does not encode or transmit user text, content, or personal data. It transmits only coherence metrics, alignment signals, and structural states produced by the system.

This appendix defines:

- privacy guarantees,
- security boundaries,
- attack surfaces,
- mitigation strategies,
- alignment with regulatory regimes,
- safe deployment patterns.

E.1 Privacy Design Principles

HER2NI MUST adhere to three foundational privacy principles.

1. Zero PII Encoding. HER-State packets MUST NOT contain:

- user names,
- raw text,
- voice/audio data,
- biosignals,
- IP addresses,
- device identifiers,
- session transcripts.

HER2NI is non-personal by design.

2. Local-First Computation. Implementations SHOULD compute:

- C_s ,
- `drift_vector`,
- `contradiction_nodes`

on-device or in the local runtime environment unless explicitly configured otherwise.

Local computation reduces:

- data exposure,
- regulatory risk,
- reliance on cloud infrastructure.

3. Minimal Exposure Surface. Externally exposed data SHOULD be limited to HER-State packets and optional AOME modulation signals. Raw human input and raw model output MUST NOT be encoded in HER2NI packets.

This design supports deployment in:

- high-security environments,
- regulated industries,
- schools and hospitals (non-diagnostic modes),
- privacy-sensitive regions (e.g., EU, AU, CA, JP).

E.2 Data Flow & Threat Model

HER2NI's threat model considers four primary attack surfaces:

1. packet corruption,
2. packet forgery,
3. inference attacks (attempt to reconstruct user content),
4. misuse or distortion of HER metrics.

E.2.1 Packet Corruption

An attacker may send malformed packets to destabilise:

- AOME behaviour,
- multi-agent coherence,
- HER-Crystal visualisations.

Mitigations (Normative):

- schema validation **MUST** reject malformed or non-canonical encodings,
- unknown keys **MUST NOT** alter interpretation of known fields,
- contradiction nodes **MUST** pass schema checks,
- numeric fields **MUST** be range-bound.

On validation failure, implementations **SHOULD** emit:

```
PacketError { recoverable: false }
```

E.2.2 Packet Forgery

Attackers may attempt to spoof HER-State packets to:

- trigger unsafe modes,
- force Recovery Mode,
- manipulate group coherence metrics,
- distort oversight dashboards.

Mitigations:

- secure channels (e.g., TLS 1.3, QUIC) are **RECOMMENDED**,
- multi-agent environments **SHOULD** sign packets (public-key signatures),
- OS-level deployments **MUST** sandbox external packet sources,
- cross-tenant traffic **MUST** be isolated.

E.2.3 Inference Attacks

HER2NI avoids inference attacks by design:

- no textual content is transmitted,
- no semantic embeddings are encoded,
- no raw sequences are included,
- no user identities are stored.

HER2NI transmits only numeric coherence metrics, which are not sufficient to reconstruct original content. This places HER2NI in a lower privacy-risk category than:

- full telemetry logs,
- embedding logs,
- chat transcripts,
- analytics systems,
- LLM prompt histories.

E.2.4 Metric Manipulation

Malicious actors could attempt to:

- mask contradiction nodes,
- exaggerate or falsify H_s ,
- spoof stability signals,
- bypass safety mechanisms.

Mitigations:

- `drift_vector` MUST be computed from local semantic/tracking functions,
- `contradiction_nodes` MUST originate from the same internal computation loop,
- AOME MUST trust local computations only unless inputs are authenticated,
- HER-Field aggregation in multi-agent systems MUST rely on signed inputs.

E.3 Transport Layer Considerations

HER2NI does not require a specific transport protocol but RECOMMENDS:

- TLS 1.3 for secure channels,
- QUIC for low-latency streaming,
- WebSockets for browser-based integrations,
- memory-mapped buffers for OS-native deployments.

HER2NI packets MUST be transport-agnostic. This supports:

- cloud inference environments,
- on-device assistants,
- distributed agent clusters,
- embedded systems.

E.4 Multi-Agent Security

In multi-agent scenarios:

- rogue agents,
- drifted agents,
- adversarial agents,
- compromised agents

may attempt to manipulate coherence metrics.

Mandatory controls:

- agents **MUST** run in isolated sandboxes,
- HER-Field aggregation **MUST** verify signatures and detect mismatches,
- cross-agent contradiction analysis **MUST** detect outliers,
- abrupt anomalies in H_s **SHOULD** trigger safety throttles.

These controls help prevent swarm-level divergence or emergent misalignment.

E.5 OS-Level Deployment Security

When HER2NI is integrated at the OS level:

- HER2NI-Rust **MUST** run in a restricted subsystem,
- `session_id` **MUST** be pseudonymous,
- HER-State logs **MUST NOT** store raw human inputs,
- AOME modulation **MUST NOT** override existing system safety policies.

OS vendors **MUST** provide:

- permission prompts for HER visualisation,
- toggleable HER2NI features,
- local-only computation modes where required.

E.6 Regulatory Alignment

HER2NI is designed to align with privacy and AI-safety regulations.

GDPR (EU). HER2NI avoids personal data and functions as low-risk telemetry.

EU AI Act. HER2NI supports transparency, safety, interpretability, and alignment monitoring.

CCPA/CPRA (California). HER2NI transmits no identifiable data.

Australian Privacy Principles. The zero-PII design is consistent with APP 6 and 11.

HIPAA (US health). HER2NI can be used in non-clinical settings because no PHI is transmitted.

OECD AI Principles and UNESCO AI Ethics. HER2NI supports principles of transparency, human oversight, and safety through coherence signalling.

E.7 Failure Modes & Safety Mechanisms

HER2NI systems MUST include at least the following safety mechanisms:

- 1. Packet Timeout Handling.** If HER-State packets cease arriving, AOME SHOULD enter Stabilisation Mode.
- 2. Contradiction Overflow Handling.** If `contradiction_nodes` exceeds a configured threshold, the system SHOULD trigger Recovery Mode.
- 3. Drift Spike Handling.** If `drift_vector` magnitude spikes above a threshold, AOME SHOULD move to Grounding Mode.
- 4. Collapse-Risk Surge Handling.** If `collapse_risk` rises by more than $\Delta 0.4$ within < 500 ms, AOME MUST override existing mode and enter Recovery Mode.

E.8 Implementation Security Checklist

All HER2NI-compliant systems MUST:

- validate packet schemas,
- reject unknown major versions,
- sandbox packet processing,
- transmit only non-sensitive metrics,
- compute signals locally where feasible,
- protect session correlation from cross-context leakage,
- avoid logging user content,
- support signed packets in multi-agent configurations,
- provide metric-only audit logs for regulators.

Optionally, systems are RECOMMENDED to:

- use hardware-backed key storage,
- support attestation for OS-level deployments,
- leverage QUIC with connection migration for resilient streaming.

Appendix E Summary

The requirements in this appendix provide HER2NI with a structured privacy and security model suitable for:

- OS-level adoption,
- research and safety auditing,
- enterprise integration,
- regulatory review and compliance,
- certification as part of AI safety frameworks.

These guarantees align HER2NI with the expectations typically applied to modern protocol and observability standards (e.g., IETF, IEEE, W3C, CNCF, ISO/IEC).

Appendix F — Test Vectors & Validation Harness

(Normative for compliance)

This appendix defines the official HER2NI compliance tests, including:

- canonical packet encodings,
- malformed packet tests,
- AOME mode-selection vectors,
- edge-case drift/contradiction patterns,
- collapse-risk transition tests,
- multi-packet temporal sequences.

Any implementation claiming **HER2NI-Compatible** MUST pass all tests in this appendix. This appendix is binding.

F.1 Deterministic Packet Encoding Tests

HER2NI requires deterministic CBOR canonical encodings. All implementations MUST produce *byte-identical* encodings for the same input structure.

Test Vector F1.1 — Minimal Valid Packet

Input Structure (conceptual JSON):

```
{
  version: "her2ni-1.0",
  Cs: 0.50,
  Ss: 0.40,
  Hs: 0.45,
  drift_vector: [0.0, 0.0],
  contradiction_nodes: [],
```

```

collapse_risk: 0.10,
timestamp_ms: 1733453000000,
session_id: h'00010203',
modality_flags: 1
}

```

Expected CBOR (Diagnostic Notation):

```

{
  1:"her2ni-1.0",
  2:0.5,
  3:0.4,
  4:0.45,
  5:[0.0,0.0],
  6:[],
  7:0.1,
  8:1733453000000,
  9:h'00010203',
  10:1
}

```

Expected CBOR (Hex Dump):

aa	# map(10)
01 69 68 65 72 32 6e 69 2d 31 2e 30	# "her2ni-1.0"
02 fa 3f000000	# Cs = 0.50
03 fa 3ecccccd	# Ss = 0.40
04 fa 3ee66666	# Hs = 0.45
05 82 fa 00000000 fa 00000000	# drift vector
06 80	# empty array
07 fa 3dcccccd	# collapse_risk
08 1b 0001927b59800000	# timestamp_ms
09 44 00010203	# session_id
0a 01	# modality_flags

Implementations MUST match this hex output exactly.

F.2 Malformed Packet Tests

Implementations MUST reject malformed packets.

F2.1 — Missing Required Keys

Packet missing required field (e.g., Cs):

```

{
  1: "her2ni-1.0",
  3: 0.3
}

```

Expected Result:

```
PacketError {  
  code: MissingRequiredField,  
  field: 2,  
  recoverable: false  
}
```

F2.2 — Out-of-Range Float

Cs = -0.2

Result: **reject**.

F2.3 — Non-canonical Ordering

3: Ss
1: version
2: Cs

Result: **reject**.

F2.4 — Contradiction Node Schema Mismatch

Reject if:

- severity missing,
- severity > 1.0,
- span indices reversed.

F.3 AOME Behaviour Test Vectors

AOME MUST output deterministic modulation values for any given HER-State packet.

F3.1 — Stabilisation Mode Trigger

Cs = 0.30
Hs = 0.20
Ss = 0.50
collapse_risk = 0.15
contradictions = []
drift = 0.12

Expected Mode: **STABILISATION**

Expected Output:

mod_depth: -0.6
mod_density: -0.5
mod_abstraction: -0.4
pacing: 0.85
tone_normalization: 1
structure_mode: 2

F3.2 — Deep Reasoning Mode

```
Cs = 0.85
Ss = 0.70
Hs = 0.92
drift = 0.05
collapse_risk = 0.02
```

Expected Mode: **DEEP_REASONING**

Expected Output:

```
mod_depth:      +0.8
mod_density:     +0.7
mod_abstraction: +0.9
pacing:          0.40
tone_normalization: 0
structure_mode:  0
```

F3.3 — Drift Spike Trigger

```
drift_vector = [0.51, -0.10]
```

Condition: drift magnitude > 0.45 → **GROUNDING**

Expected Output:

```
mod_depth:      -0.3
mod_density:     -0.2
mod_abstraction: -0.1
pacing:          0.75
tone_normalization: 0
structure_mode:  1
```

F3.4 — Collapse Risk Override

```
collapse_risk = 0.82
```

Expected Mode: **RECOVERY**

Expected Output:

```
mod_depth:      -0.8
mod_density:     -0.7
mod_abstraction: -0.6
pacing:          0.95
tone_normalization: 1
structure_mode:  2
```

F.4 Temporal Sequence Stability Tests

AOME MUST produce stable transitions over packet sequences.

F4.1 — Rising Coherence

1. $C_s = 0.45, H_s = 0.40 \rightarrow$ Exploration
2. $C_s = 0.60, H_s = 0.55 \rightarrow$ Exploration
3. $C_s = 0.80, H_s = 0.70 \rightarrow$ Deep Reasoning

Expected Transition:

EXPLORATION \rightarrow EXPLORATION \rightarrow DEEP_REASONING.

F4.2 — Collapse Cascade

collapse_risk increasing:

0.20 \rightarrow 0.35 \rightarrow 0.72

Expected:

EXPLORATION \rightarrow STABILISATION \rightarrow RECOVERY.

F.5 Multi-Agent Interoperability Tests

F5.1 — Signed Packet Verification

Reject if:

- packet unsigned,
- signature invalid,
- agent ID mismatch.

F5.2 — Group HER-Field Consistency

Given:

$$H_s = [0.82, 0.79, 0.81]$$

Expected Group Score:

$$H^G = 0.8066 \dots$$

F.6 Performance & Throughput Tests

Desktop:

- 10,000 packets \rightarrow MUST validate < 1 s
- AOME latency < 1 ms average

Mobile:

- 5,000 packets $\rightarrow < 2$ s
- memory overhead < 10 MB

Multi-Agent:

- 60 Hz coherence updates \rightarrow drift < 2 ms

F.7 Reference Compliance Script

```
def run_compliance_suite(impl):
    results = []
    for test in TEST_VECTORS:
        output = impl.run(test.input)
        results.append(compare(output, test.expected))
    return all(results)
```

Implementations MUST include this structure.

F.8 Certification Result Format

Implementations MUST output:

```
HER2NI-Compliance Report
-----
version: her2ni-1.0
parser: pass
serializer: pass
metrics: pass
aome: pass
temporal: pass
security: pass
multi-agent: pass
overall: compliant
```

If any category fails: `overall = non-compliant`.

A Conclusion

HER2NI introduces a new class of protocol: a cross-substrate coherence interface enabling real-time measurement, visualisation, and modulation of cognitive alignment between humans and artificial intelligence systems.

Unlike traditional protocols that connect machines to machines, HER2NI establishes a structured, formal, and interpretable channel between human cognitive processes and machine reasoning processes. It defines:

- a universal metric suite (C_s, S_s, H_s) ,
- a deterministic packet format for coherence-state exchange,
- a standardised modulation engine (AOME),
- a topographic interpretability layer (HER-Crystal),
- an extensible multi-user aggregation system (HER-Field).

Together, these components form a unified architecture for cognitive coherence.

HER2NI is model-agnostic, platform-neutral, and architecture-independent. It is designed to integrate equally well into research environments, enterprise systems, OS-level assistants, multi-agent frameworks, and long-term AI safety infrastructures. The protocol’s privacy guarantees, security requirements, deterministic behaviour tables, and comprehensive test vectors position HER2NI as a standard capable of scaling globally without compromising safety or interoperability.

HER2NI is protected—but not owned. It is maintained as a public technical standard whose purpose is to increase clarity, stability, and interpretability in human–AI interaction. Its design ensures neutrality across vendors, compatibility across models, extensibility across modalities, and resilience across future technological shifts.

As AI systems become increasingly capable and deeply integrated into human workflows, the absence of a formal coherence interface becomes a structural risk. HER2NI fills this void with a stable, transparent, and extensible alignment layer.

The path forward includes formal standardisation through recognised bodies, maturation of reference implementations, adoption by research labs and industry partners, silicon-level optimisation, and continued refinement of metrics and visualisation frameworks.

HER2NI is not a temporary solution or an experimental framework; it is infrastructure—a foundation for coherent, interpretable, and safe cross-substrate cognition. Its purpose is to provide every system, from the smallest mobile assistant to the most advanced reasoning clusters, with a shared language for coherence.

With HER2NI, cognitive alignment becomes measurable. Interpretability becomes operational. Safety becomes structural. Human–AI interaction becomes a domain governed not by intuition or trust alone, but by signals, standards, and shared architectures.

HER2NI marks the beginning of a new layer in the digital stack—one that connects not only devices, but the reasoning worlds of humans and machines.