NHOS NIME™ Architecture White Paper

NHOS Intelligence Matrix Engine™:
A Privacy-Preserving, On-Device Health Intelligence Architecture

Figure 1.1: NHOS Platform Intelligence Overview

Figure 1.1: NHOS Platform Intelligence Overview — This diagram illustrates how the NHOS platform is powered by the NHOS Intelligence Matrix Engine™ (NIME™). The five intelligence matrices (Knowledge, Relationship, Reasoning, Evidence, and Context) support downstream modules including Herb Library, Medication Intelligence, Symptom Mapping, and Wellness Tracking, culminating in personalized health insights.

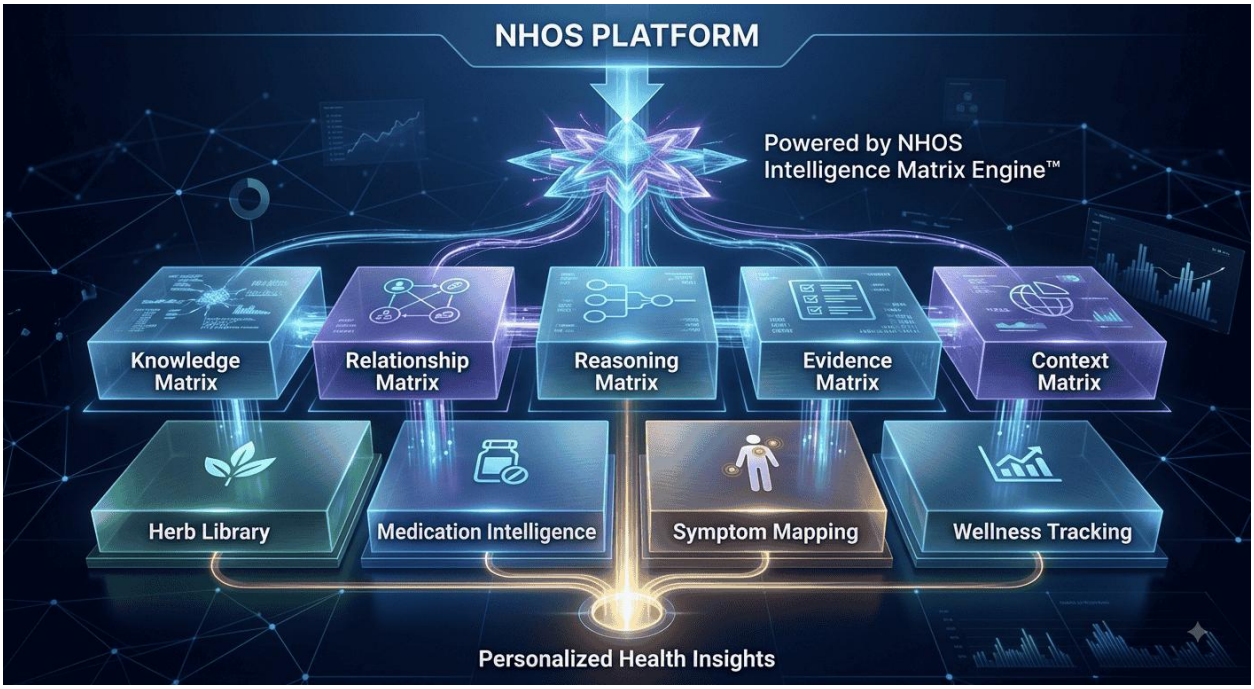

Document Metadata

Element	Details	Element	Details
Document ID	NHOS-WP-001	Version	v1.4.0
Status	Release Candidate	Date	August 2026
Author	NHOS Research Program	Plan Owner	Research Lead, NHOS Labs
Approved By	NHOS Research Governance Body	Internal Approval Date	August 2026
Distribution	Public (upon release authorization)	Next Formal Review Date	[To be determined]

Notice: *This white paper represents a technical architecture and research framework. It does not constitute clinical validation, regulatory clearance, legal approval, security certification, or independent expert endorsement. Regulatory status has not been determined.*

Contents

Definitions & Key Terms

Term	Definition
Target Design	A planned architectural or security feature that is documented and intended for implementation but has not yet been fully implemented, independently verified, or audited
Evidenced	Supported by a verifiable artifact, test result, review record, or source document
Material Remediation	A correction that addresses a finding that could materially affect the validity, safety, privacy, or credibility of the system or white paper
Risk Acceptance	A formal decision to accept a known risk without immediate remediation, documented with rationale, mitigation, owner, and expiry
Release Candidate	A version of the white paper that has completed internal publication-readiness review and is awaiting final approval for external release
External Release	Public distribution of the white paper after all publication-gate criteria are evidenced and release authority is granted
Implementation Status — Documented	The component's architecture is fully specified and documented in the NHOS technical library, though implementation may not be complete or finalized in all instances
Implementation Status — Implemented	The component is operational within the current NHOS/NIME system, subject to version and scope as noted elsewhere in this document
Implementation Status — Research / R&D	The component is in active methodological, algorithmic, or prototype development and is not yet operational, validated, or finalized
Validation Status — Active	The component is active in the current system and continues to undergo improvement, monitoring, and review (does not imply clinical validation)
Validation Status — Planned	Formal evaluation has not yet commenced; planning and protocol development are underway
Validation Status — Under Review	Evidence, documentation, or evaluation outputs are being assessed (does not imply external validation)

1. Introduction

1.1 Problem Statement

Natural health and wellness applications face a fundamental tension: providing clinically meaningful information while protecting user privacy and data sovereignty. Traditional approaches rely on cloud-based processing, which introduces privacy risks, latency, and dependency on external infrastructure. The emergence of on-device processing offers an alternative but raises questions about the feasibility of delivering comprehensive health intelligence in a local execution environment.

NHOS Labs developed the NHOS Intelligence Matrix Engine™ (NIME™) to address this challenge. NIME™ is a modular health-intelligence architecture designed for on-device processing, prioritizing privacy, evidence traceability, and clinically relevant information workflows..

1.2 Intended Use Statement

Intended Use: The NHOS Intelligence Matrix Engine™ (NIME™) is a research and development platform designed to provide evidence-informed health information, safety flags, and discussion-support summaries for natural health and wellness applications. It is intended for informational and educational purposes only.

Not Intended For: NIME™ does not diagnose disease, prescribe treatment, or replace evaluation by a qualified healthcare professional. It is not intended for clinical decision-making, medical diagnosis, or therapeutic guidance in the absence of professional oversight.

1.3 Regulatory Status Notice

Regulatory status has not been determined. NIME's regulatory classification may vary by jurisdiction, intended use, functionality, target user, deployment context, and applicable law. NHOS Labs will seek appropriate legal and regulatory review before any clinical or commercial deployment.

1.4 Research Questions

- Can a fully on-device health intelligence architecture provide evidence-informed health information without requiring external cloud processing or transmission of user data?
- Can network-level interaction analysis identify clinically relevant relationships that pairwise checking might miss?
- Can structured ontology improve the retrieval and explainability of health-information systems?
- Can a 12-engine architecture provide clinically meaningful guidance in real-world settings?

1.5 Scope

This paper describes the architecture and methodology of NIME™. It does not report on empirical validation, which is the subject of forthcoming work (see Section 11). The system is currently in pre-publication research and validation.

2. Background and Related Work

2.1 Natural Health Information Systems

Existing natural health information systems range from reference databases (Natural Medicines Comprehensive Database, NIH MedlinePlus) to consumer-facing apps. Most rely on cloud-based processing and do not prioritize privacy-preserving architectures.

2.2 Interaction Analysis

Traditional interaction checkers evaluate pairwise relationships (herb A + drug B). Recent approaches have explored network-level analysis, but few have been implemented in on-device systems.

2.3 Privacy-Preserving Health Systems

Privacy-preserving systems employ techniques such as differential privacy, federated learning, and local processing. NIME™ adopts a privacy-first design: all processing occurs on-device, with no external data transmission.

2.4 Evidence-Linked Health Systems

Evidence-linked systems trace recommendations to supporting evidence. NIME™ incorporates a citation architecture with real PubMed identifiers, academic publications, clinical guidelines, and reference sources.

2.5 Gap

To our knowledge, we identified limited publicly documented systems that combine on-device processing, network-level interaction analysis, structured ontology, and evidence-linked decision support within a single architecture.

3. System Context and Scope

3.1 Clinical/Nonclinical Boundaries

Domain	Boundary
Natural Health Information	Evidence-informed educational content
Interaction Analysis	Informational flags, not clinical warnings
Risk Classification	Evidence-based assessments, not medical advice
Protocol Generation	Illustrative suggestions, not treatment plans

3.2 System Context

NHOS™ (Natural Health Operating System) is a privacy-first health intelligence platform. NIME™ is the intelligence architecture within NHOS™.

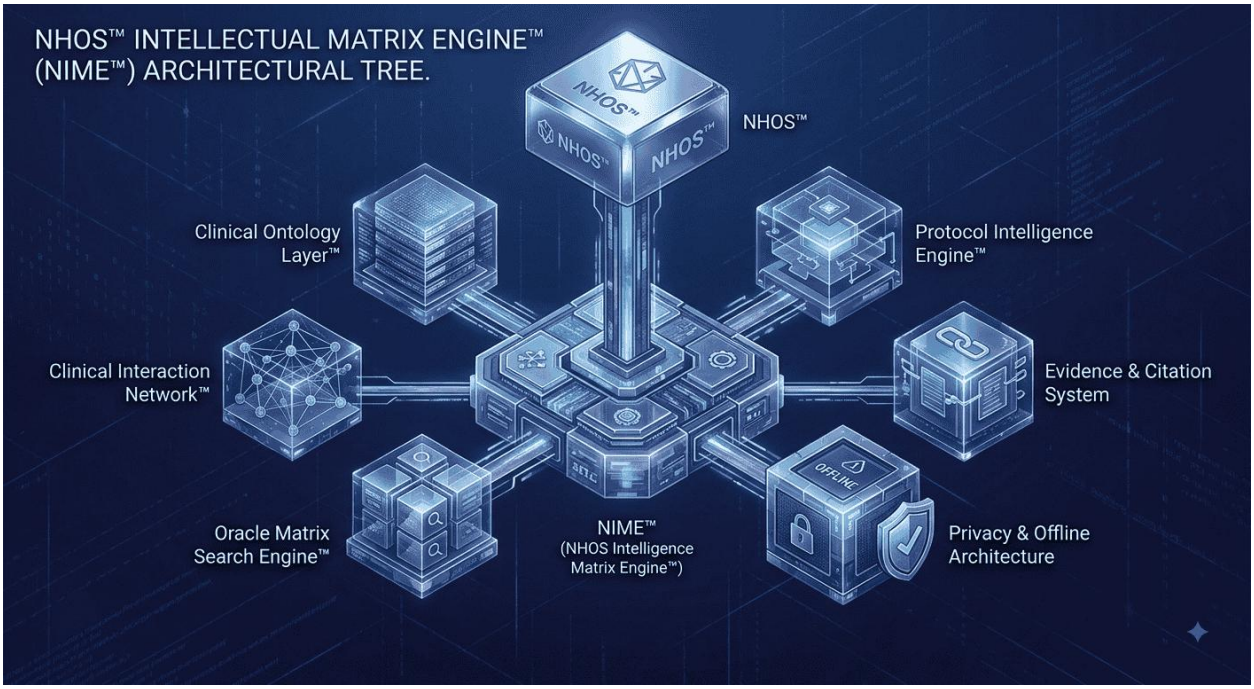

Figure 3.1: 3-D Architectural Stack of NHOS™ and NIME™ Core Components — This diagram illustrates the layered relationship between the user-facing NHOS platform and the underlying NIME™ intelligence architecture. The stack shows how NIME™ provides the core intelligence layer (including engine processing, ontology, and knowledge base) that supports higher-level NHOS applications such as interaction checking, protocol generation, and health information retrieval.

4. NIME™ Architecture

4.1 Core Principles

- **Knowledge Is Versioned:** Knowledge changes are tracked through versioned data structures and controlled updates.
- **Every Conclusion Is Explainable:** Analysis outputs are accompanied by structured information showing how the conclusion was reached.
- **Relationships Require Evidence:** Clinical relationships are supported by identifiable evidence; the system does not manufacture unsupported interaction claims.
- **Context Modifies Risk:** Context may modify the interpretation or severity of an identified relationship.
- **Mechanisms Connect Knowledge:** Entities are connected through biological and clinical mechanisms.
- **Privacy Is Foundational:** The architecture is designed for localized processing and offline operation.
- **The Engine Never Diagnoses:** NIME™ analyzes structured relationships; it does not diagnose conditions or replace clinical judgment.

4.2 Engine Taxonomy

#	Engine	Function	Implementation Status	Validation Status
1	Knowledge Engine	Curates and manages the evidence base	Implemented	Active
2	Inference Engine	Applies logical reasoning to derive conclusions	Implemented	Active
3	Clinical Interaction Engine	Maps network-level interactions	Research / R&D	Planned
4	Risk Assessment Engine	Evaluates risks and contraindications	Research / R&D	Planned
5	Evidence Grading Engine	Applies evidence quality assessment	Implemented	Active
6	Natural Language Engine	Processes user input through semantic understanding	Implemented	Active
7	Personalization Engine	Adapts recommendations to user context	Target Design	Planned
8	Validation Engine	Cross-references outputs against reference standards	Implemented	Under Review
9	Privacy Engine	Ensures data sovereignty	Implemented	Active
10	Visualization Engine	Presents findings through interpretable interfaces	Implemented	Active
11	Reproducibility Engine	Ensures consistent outputs across devices	Target Design	Planned
12	Update Engine	Manages evidence base updates	Implemented	Active

4.3 High-Level System Architecture

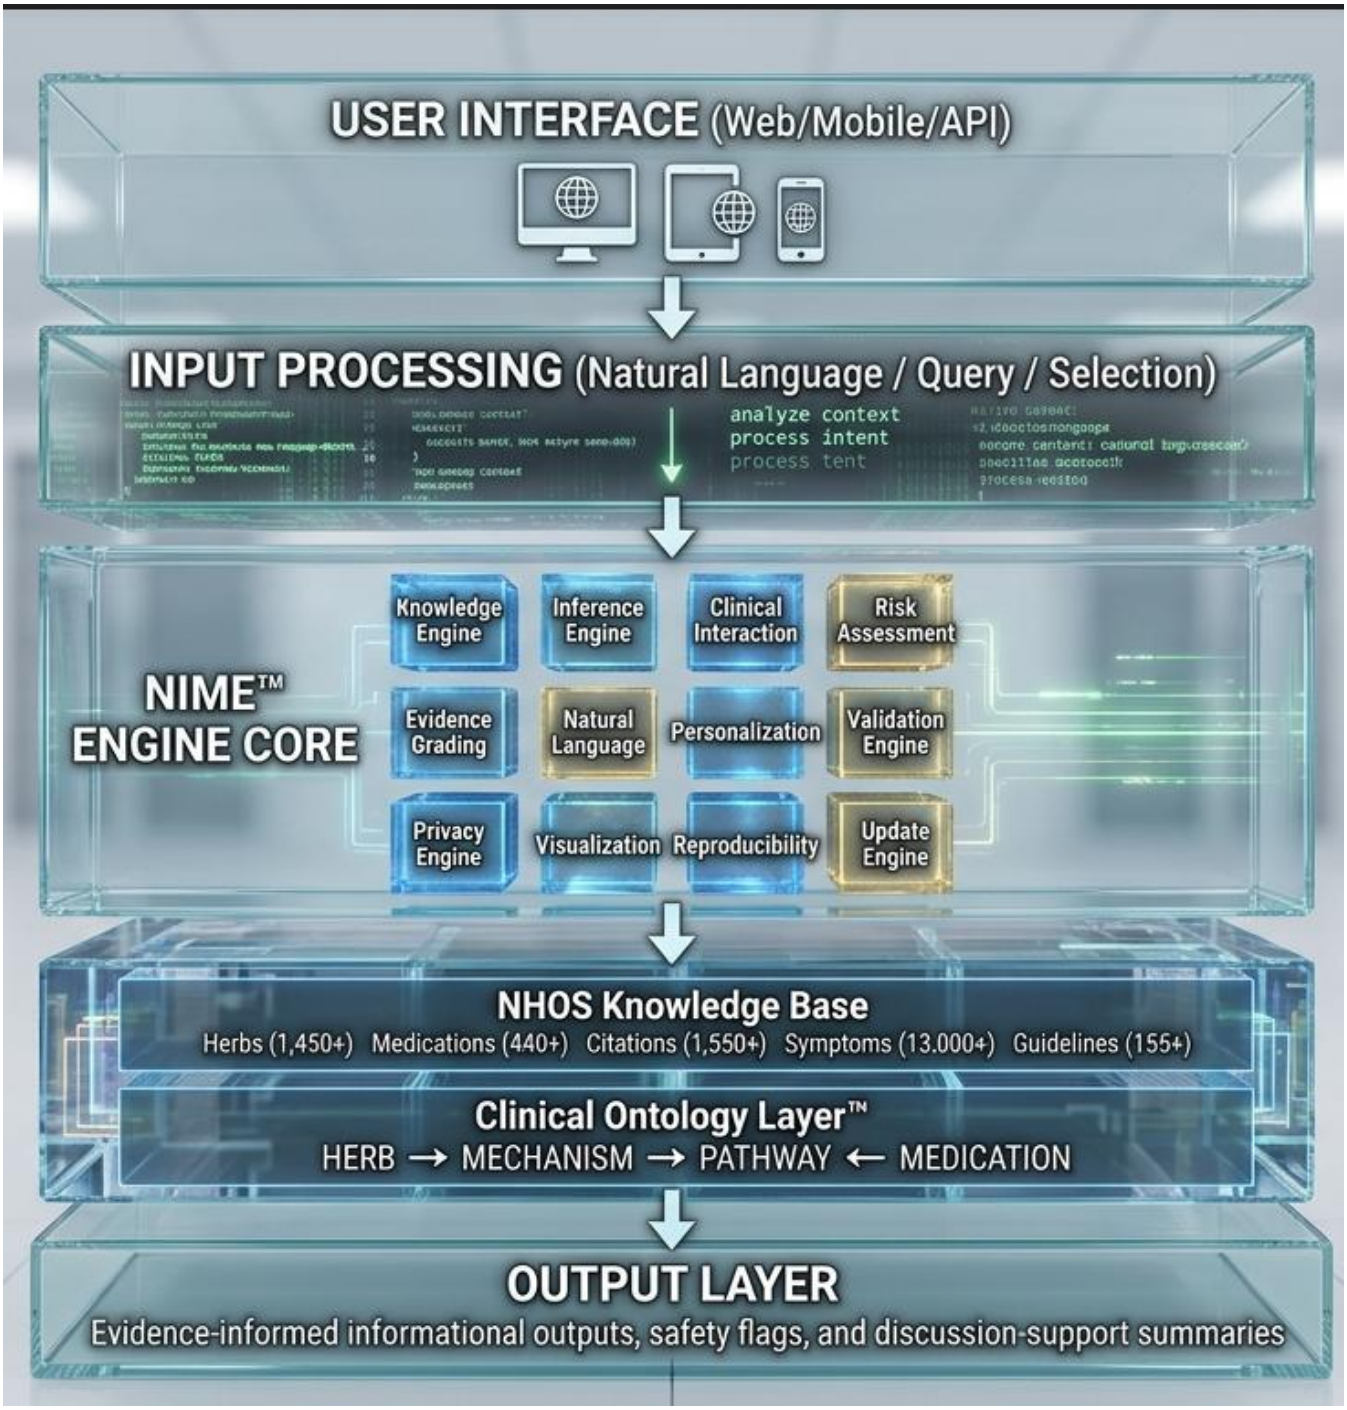

Figure 4.1: 3-D High-Level Execution & Data Flow Architecture — This diagram provides a comprehensive view of the NIME™ system architecture, showing the complete data flow from user input through engine processing to output generation. Key components include the User Interface, Input Processing (Natural Language/Query/Selection), the NIME™ Engine Core (12 engines organized into processing groups), the Knowledge & Evidence Layer (Knowledge Base and Clinical Ontology Layer™),

and the Output Layer (evidence-informed informational outputs, safety flags, and discussion-support summaries). The diagram illustrates how data flows sequentially from user entry through the engine cores and knowledge layers before producing structured, evidence-linked outputs.

4.4 Clinical Analysis Trace

A central explainability feature is the Clinical Reasoning Trace, moving deterministically across nine analysis stages:

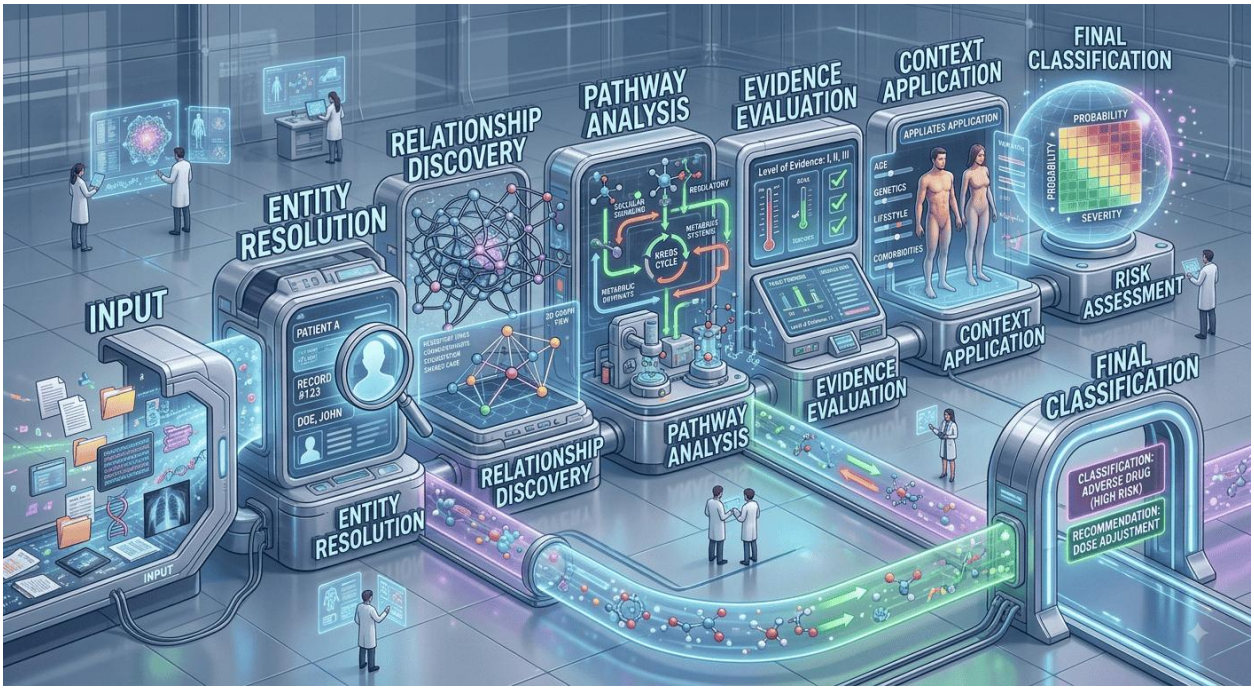

Figure 4.2: 3-D Sequential Clinical Analysis Trace Pipeline — This diagram visualizes the nine-stage clinical reasoning pipeline that underlies all NIME™ analyses. The pipeline moves sequentially from Input → Entity Resolution → Relationship Discovery → Mechanism Analysis → Pathway Analysis → Evidence Evaluation → Context Application → Risk Assessment → Final Classification. Each stage represents a discrete processing step with defined inputs and outputs, ensuring that every conclusion is accompanied by a traceable, interpretable reasoning chain.

5. Clinical Ontology Layer™

5.1 Entity Types

Type	Description	Examples
HERB	Botanical and natural substances	Turmeric, Ashwagandha, St. John's Wort
MEDICATION	Prescription and over-the-counter drugs	Warfarin, Simvastatin, Metformin
PATHWAY	Biological and clinical pathways	Hepatic Metabolism, Coagulation

MECHANISM	Mechanisms of action	CYP3A4 Inhibition, Antiplatelet Activity
RISK	Clinical risk categories	Bleeding Risk, Hepatotoxicity

5.2 Relationship Structure

Entities are connected through normalized relationships, enabling structured representation across mechanisms and pathways:

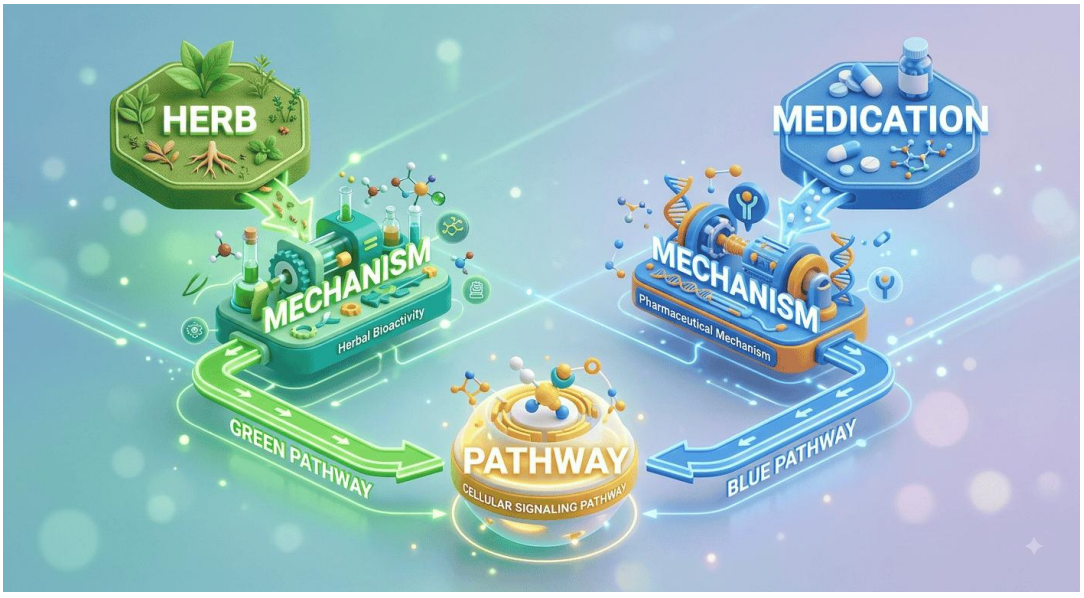

Figure 5.1: 3-D Dimensional Representation of Entity Relationship Structure — This diagram illustrates the structured relationship model underlying the Clinical Ontology Layer™. Entities (HERB, MEDICATION, MECHANISM, PATHWAY, RISK) are connected through normalized relationships, enabling the system to traverse from any starting entity through mechanisms and pathways to reach related entities. The dimensional visualization emphasizes the network-like nature of the ontology, where relationships are not limited to pairwise associations but can span multiple entity types and layers.

5.3 Illustrative Example

The following diagram is an illustrative synthetic example only—not a clinical interaction assessment. Relationship labels are simplified for architectural demonstration. Real relationships require citation, versioning, grading, and clinical review before use.

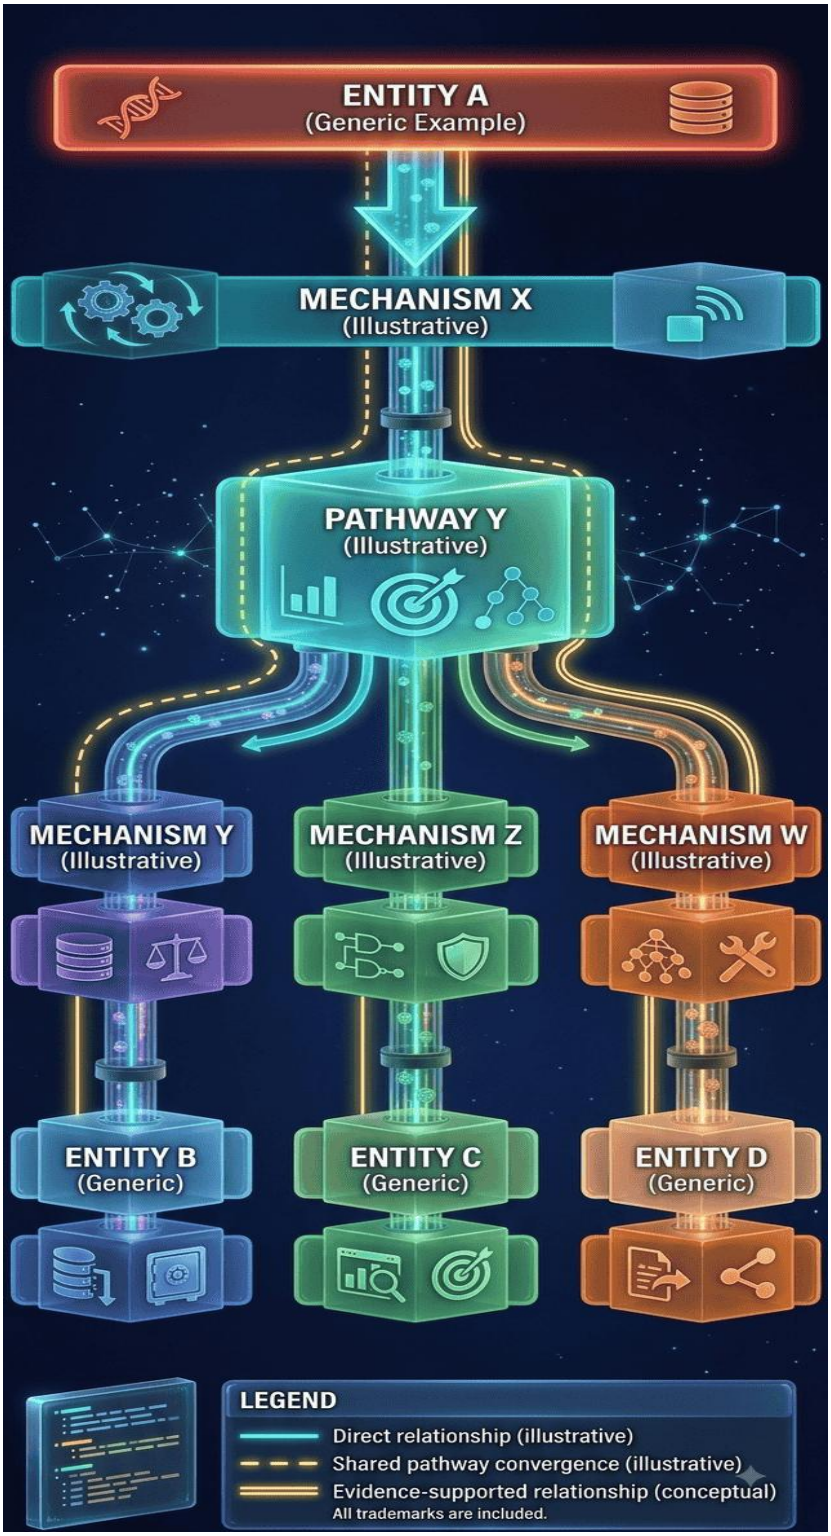

Figure 5.2: 3-D Multi-Branch Biological Mechanism & Entity Convergence Network — This illustrative example demonstrates how multiple entities can converge on shared biological mechanisms and pathways. In this synthetic example, Entity A relates to Mechanism X, which connects to Pathway Y. Pathway Y then converges with Mechanism Y, Mechanism Z, and Mechanism W, each associated with different entities (Entity B, Entity C, Entity D). This illustrates the network-level analysis capability where multiple entities may share underlying mechanisms or pathways, potentially revealing interaction patterns not visible through

pairwise analysis alone. As a synthetic example, no clinical claims should be inferred from the specific relationships shown; real implementations require evidence-supported, versioned, and clinically reviewed relationships.

6. Clinical Interaction Network™

6.1 Network Properties

Property	Description
Nodes	Health entities across multiple domains
Edges	Evidence-supported relationships between entities
Density	Multi-domain coverage with core clinical interactions
Connectivity	Hierarchical structure with clinically relevant clusters

6.2 Network Analysis Capabilities

Analysis Type	Description
Centrality	Identifies clinically significant hub entities
Clustering	Groups related interactions into therapeutic domains
Path Analysis	Maps indirect interactions through biological pathways

6.3 Safety Flag Definition

A safety flag is a non-diagnostic, evidence-linked informational notice indicating that the available knowledge base contains a potentially relevant contraindication, interaction, population-specific caution, or evidence limitation. It does not determine individual suitability or replace review by a qualified healthcare professional.

6.4 Safety Flag Framework

Level	Definition	User Action
Informational	Evidence note, no immediate concern	Review the evidence note; consult a qualified

		professional if questions remain
Caution	Potential concern requiring awareness	Consider discussing the information with a qualified healthcare professional or pharmacist before making changes
Warning	Significant potential risk	Do not start, stop, or change a medication, supplement, or treatment based solely on this output; seek prompt advice from a qualified healthcare professional or pharmacist
Critical	Potentially life-threatening risk	Do not rely on the application for emergency guidance. If there is a suspected emergency, severe reaction, overdose, or immediate danger, contact local emergency services or seek urgent medical care. For non-emergency concerns, seek prompt advice from a qualified healthcare professional or pharmacist

6.5 Accessibility and UI Safety Requirements

The system must not use:

- Alarming, stigmatizing, or absolute UI language without context
- Color-only communication for risk labels
- Risk labels without accessibility-compliant text explanations and evidence links

Safety messages must be tested with representative members of each target user group, including: people with low health literacy, screen-reader users, and people who may be distressed or under time pressure.

7. Data & Implementation Status

This section provides a consolidated view of the current implementation and validation status for all major system components. Implementation status indicates whether a component is currently operational, in research/development, or a target design. Validation status indicates the stage of formal evaluation.

7.1 Component Status Overview

To maintain clear separation between implemented capabilities, research and development activities, target architecture, and empirical validation, NHOS classifies major system components according to their current implementation and validation status. Implementation status does not imply clinical validation, and components identified as research, prototype, or validation-pending remain subject to further testing and evaluation.

Component	Current Status	Validation Status	Notes
NIME™ Architecture	Architecture documented and actively maintained	Active	Architectural validation and empirical evaluation in progress
Oracle Matrix Search Engine™	Implemented within the NHOS/NIME™ architecture	Active	Retrieval performance benchmarking pending
Clinical Interaction Network	Research and development component	Validation pending	Network-level interaction analysis in active development
Protocol Intelligence Engine	Research and development component	Validation pending	Protocol synthesis framework in active development
Evidence & Citation Framework	Implemented and actively maintained	Active	Citation system operational; source verification and evidence-quality assessment ongoing
Privacy / Offline Architecture	Implemented design components with ongoing implementation review	Active	Design components implemented; security and privacy validation in progress
Polypharmacy Risk Scoring	Research methodology / development component	Active	Scoring methodology in development
Clinical Reasoning Trace	Research / target-design component	Active	Explainability framework specified; implementation and evaluation pending
Clinical Interaction Detection	Implemented functionality under continued development	Active	Functionality operational; sensitivity, specificity, and reference-standard validation pending
Clinical Validation Program	Research program in development	Formal clinical validation not yet completed	Validation program framework established; formal clinical validation not yet commenced

7.2 Implementation Status Definitions

Status	Definition
Implemented	Component is operational within the current NHOS/NIME system. This status indicates functional availability, not clinical validation, and does not imply completeness, optimization, or finalization.
Research / R&D	Component is in active methodological, algorithmic, or prototype development. It is not yet operational, validated, or finalized.
Target Design	Component has been specified and documented as a design objective but is not yet operational, validated, or fully implemented.

7.3 Validation Status Definitions

Status	Definition
Active	Component is in use and subject to ongoing monitoring, improvement, and review. This status does not imply clinical validation or external endorsement.
Under Review	Evidence, documentation, or evaluation outputs are being assessed through internal review. This does not imply external validation or peer review.
Planned	Formal evaluation has not yet commenced. Planning and protocol development are underway.

8. NHOS Oracle Matrix Search Engine™

8.1 Overview

The NHOS Oracle Matrix Search Engine™ is a local retrieval and knowledge-discovery engine within the NIME™ architecture. It is designed to support privacy-preserving, offline-capable retrieval across heterogeneous health knowledge sources by combining multiple retrieval and indexing approaches within a unified search framework.

8.2 Data Sources

Source	Count	Type
Health Conditions	560+	Structured
Herbs & Supplements	1,450+	Structured
Symptoms	650+ categories, 13,000+ mappings	Structured
Medications	440+	Structured
Drug-Herb Interactions	650+	Relationship
Traditional Medicine Systems	12+	Structured
User-generated content	Variable	Unstructured

8.3 Hybrid Retrieval Methodology

The engine uses a 5-phase hybrid scoring algorithm:

Phase	Method	Weight
1	Phrase Matching	35 points max
2	Token Matching	12 points per token

3	Fuzzy Matching (Levenshtein Distance)	6 points max
4	Category Weighting	1.1×–3× multiplier
5	Type-Based Boosting	+4 to +12 points

8.4 Confidence Score Calculation

Formula: *Confidence = (score / maxPossibleScore) × 90 + 5*

9. Protocol Intelligence Engine™

9.1 Overview

The Protocol Intelligence Engine™ synthesizes evidence-weighted protocols through a structured 13-stage pipeline. It integrates evidence retrieval, intervention selection, contraindication screening, interaction analysis, synergy identification, and explanation generation.

9.2 Protocol Synthesis Pipeline

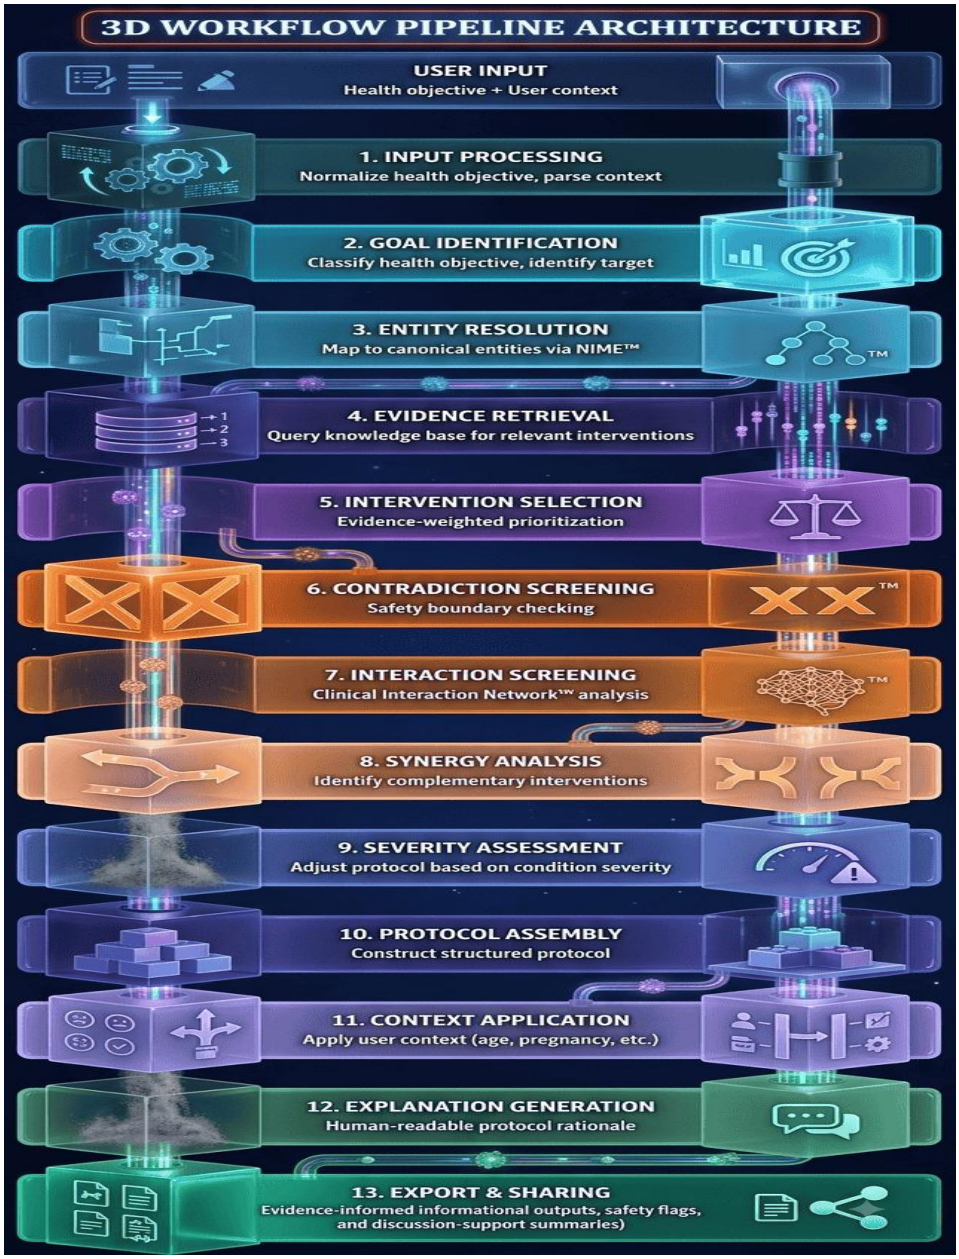

Figure 9.1: 3-D Spiral Helix Pipeline for 13-Stage Protocol Synthesis — This diagram illustrates the sequential 13-stage pipeline through which the Protocol Intelligence Engine™ processes health objectives into structured, evidence-informed outputs. The pipeline progresses from Input Processing through Goal Identification, Entity Resolution (via NIME™), Evidence Retrieval, Intervention Selection, Contraindication Screening, Interaction Screening (via Clinical Interaction Network™), Synergy Analysis, Severity Assessment, Protocol Assembly, Context Application, Explanation Generation, and Export & Sharing. The spiral helix visualization emphasizes the iterative, multi-stage nature of protocol synthesis, where each stage builds upon previous stages to ensure comprehensive evidence integration and safety screening.

10. Evidence Framework

10.1 Evidence Quality Scale

Grade	Definition	Criteria
Grade A — High Quality	Consistent results from multiple RCTs or high-quality systematic reviews	≥2 high-quality RCTs; low risk of bias; consistent results; direct relevance
Grade B — Moderate Quality	Evidence from controlled studies or non-randomized trials	1 RCT or ≥2 controlled studies; moderate risk of bias; acceptable relevance
Grade C — Low Quality	Evidence from observational studies, case reports, or expert opinion	Observational studies, case series, expert opinion; high risk of bias
Grade D — Very Low Quality	Limited or conflicting evidence requiring further investigation	Conflicting evidence; very low quality; further investigation required

10.2 Relationship to Established Frameworks

This framework is proposed as an internal evidence-classification system informed by established frameworks such as GRADE. It is not claimed to be equivalent to or validated against established clinical evidence-grading standards.

10.3 Operational Safeguards

- **Separate evidence certainty from recommendation strength:** A strong recommendation is not automatically justified by high-certainty evidence, and a low-certainty evidence base may still require a safety warning.
- **Outcome-specific grading:** An herb may have different evidence grades for efficacy, adverse effects, interactions, and different populations.
- **Review currency:** Track 'last reviewed,' 'next review due,' evidence cutoff date, reviewer, and conflict-of-interest fields for every evidence record.

10.4 Evidence Disclaimer

Evidence grades describe the certainty and applicability of identified evidence; they do not independently establish treatment efficacy, safety for an individual, or a clinical recommendation.

10.5 Decision Rule

If a source is retracted, superseded by a material safety warning, or exceeds its review-due date, the associated evidence record must be marked 'review required,' and outputs relying primarily on that record must display an uncertainty notice or be withheld pending review.

11. Privacy and Security Model

11.1 Privacy Statement

No cloud processing of user data: NIME performs user-input processing and analysis locally. The system may retrieve signed, non-personalized knowledge-base updates through a verified distribution channel. No user input, health data, identifiers, telemetry, analytics, or crash reports are transmitted by the application.

11.1 Privacy Boundary Diagram

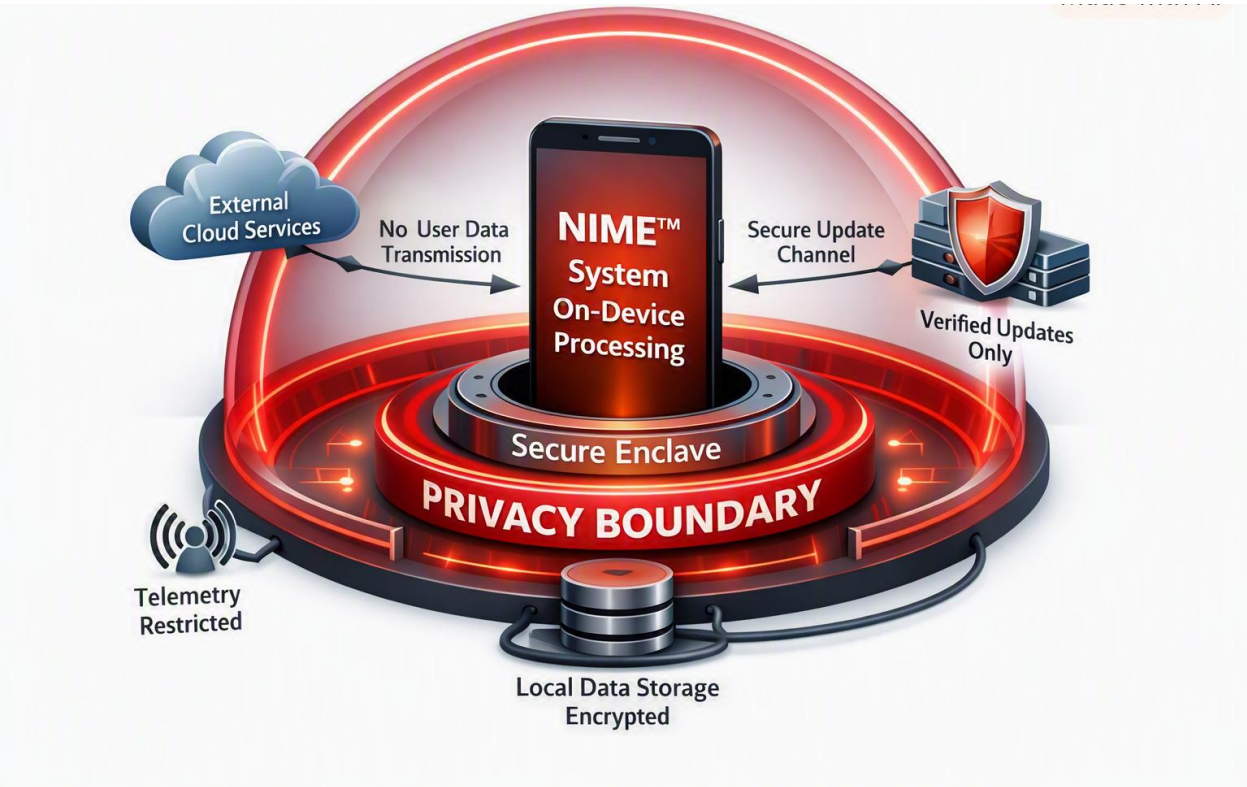

Figure 11.1: 3-D Spatial Model of On-Device Local Data Boundary vs. External Update Channel — This diagram illustrates the privacy boundary that separates on-device local processing from external infrastructure. The left side shows all processing, storage, and data flow occurring within the user device: the Application Layer (containing Local Data Storage, NIME™ Processing, and Local Knowledge Base) and the External Boundary (showing the Verified Update Channel with signed updates and version manifests). The diagram explicitly demonstrates that user inputs, health data, identifiers, telemetry, and analytics stay within the device and are not transmitted externally.

Figure 11.2: 3-D Spatial Model of On-Device Local Data Boundary vs. External Update Channel — This diagram illustrates the privacy boundary that separates on-device local processing from external infrastructure. The left side shows all processing, storage, and data flow occurring within the user device: the Application Layer (containing Local Data Storage, NIME™ Processing, and Local Knowledge Base). The right side shows the External Boundary with the V

erified Update Channel (signed updates and version manifests only). The diagram explicitly demonstrates that user i nputs, health data, identifiers, telemetry, and analytics stay within the device and are not transmitted externally.

Figure 11.2: 3-D Spatial Model of On-Device Local Data Boundary vs. External Update Channel

This diagram illustrates the privacy boundary that separates on-device local processing from external infrastructure, demonstrating that no user data is transmitted beyond the device boundary.

On-Device Components (Left/Center):

Component	Description
USER DEVICE	The physical device (phone, tablet, or computer) where all processing occurs
APPLICATION LAYER	The NHOS application environment housing all processing, storage, and analysis functions
LOCAL DATA STORAGE	Ephemeral user inputs (processed and not retained), session-based reasoning traces, encrypted local cache, and optional user history (opt-in only)
NIME™ PROCESSING	Core analysis functions including Entity Resolution, Interaction Analysis, Risk Assessment, Evidence Retrieval, and Protocol Synthesis — all performed locally
LOCAL KNOWLEDGE BASE	Encrypted on-device storage including Entity Registry, Relationship Matrices, Evidence Database, and Signed Update Manifest

External Boundary (Right):

Component	Description
VERIFIED UPDATE CHANNEL	The only external communication: signed knowledge-base updates and version manifests
NO USER DATA TRANSMITTED	Explicit privacy guarantee that user inputs, health data, identifiers, telemetry, and analytics are not transmitted

Key Privacy Features Demonstrated:

✓ All user data processing occurs on-device

✓ No user inputs, health data, identifiers, or analytics are transmitted

✓ The only external communication is for signed, non-personalized updates

✓ All local storage is encrypted

✓ User history is opt-in only

This diagram demonstrates the architectural separation between on-device processing and external infrastructure, making explicit that user data does not leave the device.

11.3 Security Design (Target)

Aspect	Target Design
Encryption	AES-256-GCM or equivalent, subject to implementation review and platform support
Key Handling	Keys intended to be generated and retained locally, subject to final implementation and platform key-management design
Backup	Application designed to avoid application-initiated cloud backup; platform-level behavior documented and controlled where possible
Telemetry	Target design excludes behavioral telemetry and analytics; any future diagnostic reporting would require explicit disclosure and user consent
Update Signature	Ed25519 or equivalent
Version Control	Version manifest with rollback protection
Provenance	Signed by NHOS Labs
Revocation	Version revocation capability
Offline Updates	Manual update option

11.4 Update Service Privacy Contract

The update service receives only the minimum technical request necessary to deliver a public, signed update artifact. It does not receive account identifiers, health information, user-entered content, persistent device identifiers, advertising identifiers, or behavioral telemetry. Network infrastructure logs, if any, will be minimized, retained for a defined period, and governed by a published operational policy.

Limitations Statement: If the exact behavior of a distribution provider cannot be guaranteed, limitations shall be stated precisely rather than using an absolute claim.

11.5 Threat Model

Threat	Mitigation (Target)
Lost/compromised device	Device encryption; no sensitive data stored
Malicious update	Signed updates; version verification
Reverse engineering	Obfuscation; code ultimately client-side
Inference from local data	Minimal local storage; no persistent logs
Unauthorized physical access	Device-level encryption; biometric/device lock

11.6 Scope Boundary

'On-device' reduces data-transmission risks but does not eliminate endpoint-security risks. The security of the system depends on the security of the underlying device and operating system.

11.7 Data Retention and Deletion Specification

Element	Specification
Stored Locally	User inputs (ephemeral), reasoning traces (session), encrypted local cache, optional user history (opt-in)
Retention Rationale	User inputs and traces are processed and not retained; cache improves performance; history is user-controlled
Default Retention Periods	Session-based only for inputs/traces; cache is maintained until cleared; history is user-managed

User Controls	Delete/export functionality provided
Deletion Verification	Data is removed when deleted
Backup Limitations	Application avoids application-initiated cloud backup

12. Evaluation Framework

12.1 Evaluation Dimensions

Domain	Evaluation Question
Entity Resolution	How accurately does NIME™ map user inputs to canonical entities?
Interaction Detection	How accurately are documented relationships identified?
Retrieval	How relevant are the top-ranked search results?
Evidence Retrieval	Are supporting citations correctly retrieved and attributed?
Risk Classification	How concordant are classifications with reference standards?
Network Analysis	Does multi-entity analysis identify relationships missed by pairwise analysis?
Explainability	Can reviewers reconstruct how outputs were generated?
Privacy	What information leaves the local execution environment?
Performance	What are latency, memory, and offline-performance characteristics?
Reproducibility	Are equivalent inputs processed consistently?
Usability	Can intended users correctly understand the outputs?

12.2 Proposed Metrics

Domain	Metrics
Entity Resolution	Precision, Recall, F1, Error Taxonomy
Interaction Detection	Sensitivity, Specificity, PPV, NPV, F1, Concordance
Retrieval	Precision@K, Recall@K, MRR, nDCG
Risk Classification	Concordance with Reference Standards
Performance	Median Latency, P95 Latency, Memory Usage
Reproducibility	Deterministic Output Testing

12.3 Reference Standards

Evaluation will compare NIME™ outputs against:

- Natural Medicines Comprehensive Database
- US Pharmacopeia
- European Medicines Agency monographs
- WHO guidelines
- Peer-reviewed clinical trial data

12.4 Validation Protocol Requirements

Each validation study shall include:

- Intended-use statement and target user group
- Ground truth/reference standard description
- Inclusion/exclusion criteria
- Dataset provenance and licensing
- De-identification process (if applicable)
- Independent reviewers and adjudication method
- Sample size/power justification
- Predefined acceptability thresholds
- Handling of missing/conflicting evidence
- Error taxonomy (especially false negatives)
- Safety escalation criteria
- Human-factors evaluation

- Fairness/representativeness review
- Licensing review
- Version information
- Independent governance review

13. Limitations

Limitation	Description
Scope	Focused on natural health interactions; pharmaceutical interactions not primary focus
Evidence Base	Dependent on available published literature
Population Diversity	Validation in progress for underrepresented populations
Long-term Effects	Primarily based on short-term and medium-term evidence
Clinical Validation	Prospective clinical validation studies are planned
Network Complexity	Network-level interaction patterns require further validation
Security	Security controls are described as target design; implementation and testing are ongoing

14. Research Roadmap

Phase	Timeline	Activities
Phase 0 — Research Governance	Q3 2026	Governance framework, terminology, standards
Phase 1 — Methodology Papers	Q3-Q4 2026	NIME Architecture, Clinical Interaction Network, Oracle Search, Protocol Engine
Phase 2 — Protocol Registry	Q3 2026	Register NHOS-RP-001 through NHOS-RP-005
Phase 3 — Validation Studies	Q4 2026-Q2 2027	Entity Resolution, Interaction Detection, Oracle Search, Protocol Assessment, Polypharmacy Scoring

Phase 4 — Preprints	Q4 2026-Q2 2027	Submit methodology papers to appropriate repositories
Phase 5 — Peer Review	Q2-Q4 2027	Submit to peer-reviewed journals
Phase 6 — External Collaborations	Q2 2027+	Independent validation, research partnerships

15. References

15.1 NHOS Labs Technical Documentation

1. NHOS Labs. (2026). NIME™ Technical White Paper: On-Device Health Intelligence. NHOS Research Series.

2. NHOS Labs. (2026). NIME™ Interaction Checker: Clinical Interaction Network Architecture Specification. NHOS Technical Documentation.

3. NHOS Labs. (2026). Privacy-Preserving Health Intelligence: Ethical Framework and Implementation. NHOS Ethics Series.

4. NHOS Labs. (2026). NHOS Oracle Matrix Search Engine™: Technical Specification. NHOS Technical Documentation.

5. NHOS Labs. (2026). NHOS Protocol Intelligence Engine™: Technical Specification. NHOS Technical Documentation.

15.2 Reference Standards and Authoritative Sources

6. World Health Organization. (2023). WHO Guidelines on Natural Health Products. Geneva: WHO Press.

7. Natural Medicines Comprehensive Database. (2025). Evidence-Based Natural Medicine Reference. Therapeutic Research Center.

8. US Pharmacopeial Convention. (2025). USP-NF Compounding Standards. Rockville: USP.

9. European Medicines Agency. (2024). EMA Monographs on Herbal Medicinal Products.

15.3 Academic Literature

Academic literature supporting the NHOS research program is maintained in the NHOS evidence repository. This manuscript includes only publications whose bibliographic metadata and publication status have been independently verified through authoritative sources such as publisher platforms, PubMed, Crossref, or other recognized scholarly indexes.

Additional literature will be incorporated following verification of bibliographic accuracy, publication status, methodological relevance, and applicability to the specific NHOS research question being addressed.

16. Appendices

Appendix A: Knowledge Base Metric Reconciliation

Metric	NHOS Knowledge Base	NIME Active Registry
Herbal Profiles	1,450+	770+
Medication Profiles	440+	350+
Herb–Drug Interactions	650+	500+
Clinical Mechanisms	250+	200+
Biological Pathways	19	19
Symptom Mappings	13,000+	—
Academic Citations	1,500+	—
Clinical Guidelines	155+	—

Note: These figures represent different layers of the NHOS knowledge architecture rather than duplicate inventories. The broader NHOS knowledge base contains entities and relationships that may not yet be registered for every workflow.

Appendix B: Glossary of Terms

Term	Definition
NHOS™	Natural Health Operating System — the overall platform
NIME™	NHOS Intelligence Matrix Engine — the intelligence architecture
Clinical Interaction Network™	Network-level relationship analysis framework
Oracle Matrix Search Engine™	Local hybrid retrieval engine

Protocol Intelligence Engine™	Evidence-weighted protocol synthesis
Clinical Ontology Layer™	Structured entity/relationship representation
Evidence Grade	Classification of evidence quality (A-D)
Safety Flag	Non-diagnostic, evidence-linked informational notice

Appendix C: Document Version History

Version	Status	Date	Changes
v1.0.0	Draft	August 2026	Initial architecture documentation
v1.1.0	Internal Review	August 2026	Added regulatory disclaimer, safety flags
v1.2.0	Internal Review	August 2026	Added security artifacts, evidence framework
v1.3.0	Internal Review	August 2026	Added safety language, privacy model, reviewer requirements
v1.4.0	Release Candidate	August 2026	Finalized for publication gate review

NHOS Research Program — NIME™ Architecture White Paper | Version v1.4.0 | Release Candidate
This white paper is scheduled for external release only after the publication gate is satisfied, including completed reference verification, independent expert review, security review, legal/compliance review, and documented remediation of all material findings.

NHOS NIME™ Architecture:

NHOS Intelligence Matrix Engine™ — A Privacy-Preserving, On-Device Health Intelligence Architecture

Technical Architecture White Paper · Version 1.4.0 · August 2026

NHOS Labs, Inc.

NHOS: https://nhos.health

NHOS Research Registry: https://nhos.health/research

NIME Architecture: https://nhos.health

Permanent DOI (Version): https://doi.org/10.5281/zenodo.22070897

Cite All Versions (Concept DOI): https://doi.org/10.5281/zenodo.22070897