Computable Nature Dependency in a Watershed Knowledge Graph

Jay Gutierrez[1] and Javier C. Alvarez[2]

[1] Independent Researcher, Antioquia, Colombia.

[2] Techlife Research Group, School of Applied Sciences and Engineering, EAFIT University,
Medellin 050022, Colombia.

Contact email: Jay Gutierrez, jg@graphoflife.com

Highlights

- Auditing nature's value: We move beyond simple "nature scores" by using a transparent, evidence-backed graph that clearly shows the entire pathway, from how ecological processes work to their financial relevance.
- Rigorously built and tested: Our knowledge graph is a fully reproducible snapshot, containing 225 nodes and 1,807 connections. It passes 33 integrity checks, guaranteeing the data's reliability and complete transparency of its origins.
- Tracing the service chain: A key query demonstrates the full service chain in action, linking a specialized ecological interaction of *Ensifera ensifera* and *Passiflora mixta* directly to the Aburra Valley's named water-service infrastructure.
- Systemic ecological support: Our analysis shows the ecological system is supported by a broad "backbone" of 50 interconnected species (the shell-23 k-core plateau). This approach helps us screen the entire system's structure, rather than relying on just one keystone species.
- A foundation for trust: To ensure credible results, we use a transparent set of rules (bridge axioms, evidence tiers, and governed queries). This framework establishes a reliable protocol that avoids making unproven claims about future predictions or financial pricing.

In Brief

Nature-risk and natural-capital workflows increasingly ask institutions to explain how ecological change can affect infrastructure, service continuity, and financial exposure. Existing tools often help with sector screening, spatial ecosystem-service estimation, or biodiversity data access, but they rarely preserve a named route from ecological mechanism to named infrastructure and finance-facing interpretation. This manuscript presents a reproducible Medellin watershed knowledge graph as a proof of concept for computable nature dependency. The contribution is not a calibrated hydrological model or an automated investment engine. It is an auditable representation pattern in which ecological interactions, ecosystem processes, infrastructure assets, economic exposure boundaries, evidence provenance, bridge assumptions, and reviewed queries remain inspectable as connected graph objects.

Abstract

Resolving nature-related financial risk and watershed resilience requires a representational shift from opaque, aggregate indices to auditable, mechanistic pathways. Current natural-capital workflows often succumb to an "epistemic collapse", where granular ecological interactions are erased in favor of sector-level proxies or spatial heatmaps that lack a downstream asset trace. Here, we present a reproducible architectural proof of concept for computable nature dependency within the Medellin-Aburra watershed of Colombia. Using a YAML-governed Neo4j property graph, we materialize a multidimensional system comprising 225 nodes and 1,807 relationships, substantiated by a four-tier evidence registry and governed by 33 domain-specific integrity checks. Our headline query reveals a precise, eight-edge service chain that preserves mechanistic continuity: beginning with the specialized pollination mutualism of *Ensifera ensifera* and *Passiflora mixta*, the path traverses canopy recruitment, fog interception, and catchment regulation to terminate in named infrastructure assets, e.g., the La Fe reservoir, La Ayura treatment plant, and the EPM distribution network. Structural analysis of the 1,023-edge biological substrate identifies a complex ecological backbone (k-core shell 23), while compound-stress simulations indicate a 67.5% loss in pathway integrity under urbanization and climate pressures. By enforcing the ontological separation of evidence, bridge axioms, and risk metrics, this architecture formally establishes a rigorous, standards-governed protocol for the continuous, verifiable auditability of ecological dependencies. We demonstrate that nature dependency is most accurately managed not as a static score, but as a transparent, query-reviewed evidence agenda that maintains the auditability of ecological complexity across the entire service-to-finance continuum.

Keywords: biodiversity knowledge graph; nature-related financial risk; ecosystem services; watershed infrastructure; k-core decomposition; evidence provenance; natural capital; graph-based decision support

Graphical Abstract

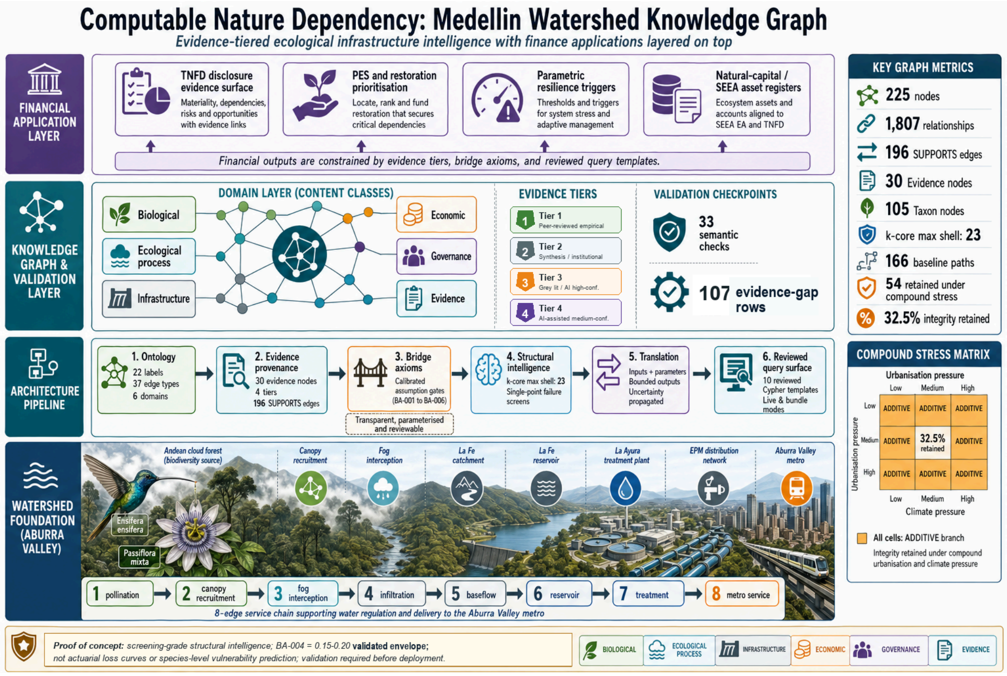

Graphical Abstract. The graphical abstract summarizes the central architecture: ecological, infrastructure, governance, evidence, and finance-facing objects are represented in a single evidence-tiered graph while validation checks, query governance, and proof-of-concept boundaries remain visible. Values shown are canonical validation outputs or explicitly screening-grade scenario outputs.

Introduction

The quantification of nature-related risk is fundamentally constrained by a denominative challenge. While financial institutions, utilities, and conservation funds may qualitatively assert dependencies on ecosystem services such as water regulation or pollination, such assertions often lack the resolution required for rigorous technical audit. A robust evaluative framework must distinguish whether a dependency is localized to a broad economic sector, a specific geographic site, a discrete catchment, or a named infrastructure asset supported by specific biological interactions. Furthermore, the epistemic basis of these claims, ranging from e.g., direct empirical observation and institutional records to neurosymbolic ecological network reconstruction outputs and theoretical bridge axioms, must be explicitly characterized to facilitate transparent risk disclosure.

This representational gap is critical as nature finance increasingly shifts toward standardized disclosure regimes, such as the Taskforce on Nature-related Financial Disclosures (TNFD). Parametric instruments and natural-capital registers require a precise mapping of ecological assets to service boundaries; without such clarity, service scores remain ambiguous. An effective evidence architecture must therefore differentiate between observed evidence and scenario-based proxies, ensuring that structural pathways are not conflated with calibrated operational responses. Consequently, the challenge of measuring nature dependency is as much an epistemic requirement as it is a computational task.

While existing methodologies, including sector-level exposure frameworks, spatial ecosystem-service modeling, and biodiversity databases, address facets of this problem, they often fail to preserve mechanistic continuity across domains. The primary deficit is not a lack of data, but the absence of a governed representational substrate (e.g., a cross-domain connective tissue). Such a substrate must enable reviewers to traverse the entire causal chain, from primary biological interactions to terminal infrastructure endpoints, while maintaining visibility of evidence tiers, translation assumptions, and inherent query boundaries.

This work introduces the concept of "computable nature dependency" as a formal solution to this representational need. In this framework, a dependency is defined as a claim materialized through graph objects and relationships with explicitly declared semantics and evidentiary support. The Medellin-Aburra implementation utilizes a Neo4j property graph to leverage advanced Cypher traversal and visualization capabilities. However, governance is maintained through a YAML-based registry that defines ontology, bridge axioms, and evidence tiers. This architecture ensures that while the runtime environment is optimized for performance, the system remains anchored to a standards-governed audit contract through exports aligned with PROV-O, DCAT, and SHACL protocols.

This work resolves a core representational challenge in natural-capital assessment, establishing a formal architecture for computable nature dependency. The contribution is three-fold, addressing critical gaps in scientific rigor and institutional auditability. Architecturally, we introduce a unified, reproducible knowledge graph that achieves mechanistic continuity by seamlessly integrating heterogeneous domains, from highly specific biological interactions to downstream economic exposure boundaries, within a single topological space. Empirically, we materialize and auto-validate the canonical service chain for the Medellin-Aburra watershed, providing the first auditable, named path from specialized ecological mutualism to named municipal infrastructure assets. Methodologically, the system enforces a novel protocol for epistemic assurance: graph paths are constrained as verifiable claims, k-core decomposition is strictly utilized as a structural screening primitive, and cross-domain translation is governed by inspectable bridge axioms. Crucially, this governed structure prevents the conflation of topological descriptors with causal predictions or actuarial pricing logic, offering a rigorous, standards-compliant framework for evidence-based diligence.

Related Work

Nature-related financial disclosure and natural-capital frameworks provide the institutional vocabulary for the problem. TNFD, ENCORE, SEEA Ecosystem Accounting, and the Dasgupta Review make dependencies and impacts visible to organizations that do not normally work in ecological mechanism language (Dasgupta, 2021; Natural Capital Finance Alliance, 2024; Taskforce on Nature-related Financial Disclosures, 2023; United Nations et al., 2021). Arguably, their strength is governance legibility. Their limitation for this manuscript's purpose is resolution. They do not, by themselves, encode a named route from a local biological interaction to a named reservoir or treatment plant.

Ecosystem-service models provide a second foundation. InVEST-style and ARIES-style workflows can estimate service supply, demand, or flow across space (Chaplin-Kramer et al., 2019; Villa et al., 2014). These methods are essential for many valuation and planning tasks, but their usual analytical object is a mapped service quantity or modelled service process. A graph representation, however, addresses a different need: retaining the identity of each biological, ecological, infrastructural, and economic object in a mechanism chain so that evidence and assumptions can be inspected edge by edge.

Biodiversity data infrastructures provide the biological substrate required for any species-level system. GBIF and GloBI expose occurrence and interaction data at scales that make ecological graph construction possible (GBIF Secretariat, 2026; Poelen et al., 2014). Network ecology supplies the structural concepts used to interpret interaction architecture, including mutualistic-network dependence, food-web robustness, and graph centrality (Allesina and Pascual, 2009; Bascompte and Jordano, 2007; Dunne et al., 2002; Vanbergen et al., 2017). The present graph uses those ideas conservatively. A high k-core shell is a structural descriptor inside the represented substrate, not a demographic vulnerability estimate.

Moreover, knowledge graphs, ontologies, and Semantic Web standards provide the representational discipline. ENVO, BCO, OpenBiodiv, and related biodiversity knowledge-graph efforts show how ecological and biodiversity knowledge can be made semantically explicit (Buttigieg et al., 2013; Dimitrova et al., 2021; Penev et al., 2019; Walls et al., 2014). NatureKG and adjacent nature-finance graph work indicate increasing interest in ontology-backed financial risk systems (Kushwaha et al., 2025). However, the Medellin graph introduced here contributes a different case: an asset-resolution watershed dependency graph that combines species interaction structure, infrastructure traversal, evidence-tiered relationship support, bridge-axiom calibration status, post-load validation checks, and a reviewed Cypher query surface.

A systematic synthesis of the extant literature and operational landscape reveals a profound structural void, given that no existing architecture successfully integrates these discrete functional requirements into a cohesive, city-scale watershed intelligence system. This observation defines the precise epistemic positioning of the current study, representing a targeted architectural intervention rather than a claim of absolute singularity. While peripheral frameworks remain indispensable as computational complements, the primary contribution of this work lies in its foundational architectural proposition. We introduce a proof-of-concept implementation engineered to maintain the uninterrupted continuity of the service chain, preserving the rigorous identity of each link from specific ecological mechanisms to named infrastructure and financial endpoints. By refusing to collapse such complexity into opaque, aggregate metrics, this system establishes a new standard for the auditable management of ecological uncertainty.

Conceptual Model

The methodology for computable nature dependency is predicated upon the rigorous epistemological separation of operational constructs that are frequently conflated within descriptive nature-risk discourse. Specifically, the framework enforces a mandatory differentiation between evidential support and translation protocols, graph-structural metrics and inherent ecological vulnerability, computed scenario pathways and predictive forecasting, and bounded financial exposure and realized loss. The Medellin graph architecture is explicitly designed to maintain the visibility and discrete inspectability of these categories for any external reviewer.

The analytical sequence is strictly sequential and governed by this conceptual separation. Primary source registries formalize the ontology, evidence tiers, bridge axioms, and query manifest boundaries. Following ETL materialization in Neo4j, the system executes post-load validation checks to confirm graph invariants, critical-path existence, provenance coverage, and orphan node absence. Structural metrics (e.g., k-core) and reviewed Cypher queries then generate computational outputs (e.g., path counts, scenario retention). Crucially, interpretation of these outputs is contingent upon the documented evidence tier, the calibration status of the relevant bridge axiom, and the explicit query caveats.

This constrained model serves as the primary mechanism for preventing epistemological overclaim. Reviewer challenges are systematically routed to specific, inspectable graph components: a contested biological or infrastructural link is addressed by the relationship's edge-specific SUPPORTS rows and the associated evidence registry; a challenged cross-domain translation is governed by the functional form and calibration status of the BridgeAxiom node; and a finance-facing output is traceable to a RiskMetric node or declared bridge-axiom parameter. Consequently, the graph functions not as a persuasive predictive device, but as a verifiable, object-oriented audit structure.

Concept	**Operational representation**	**Interpretation boundary**
Observed or materialized graph object	Typed Neo4j node or relationship governed by YAML ontology files	The object exists in the graph denominator, not necessarily in a complete real-world census
Evidence support	Evidence nodes and SUPPORTS relationships with confidence tiers and relationship-level metadata	Evidence quality varies by tier and source; support is not uniform certainty
AI-assisted substrate	Project-maintained biological interaction and structural-metrics inputs	Useful for proof-of-concept topology, requiring expert validation for operational claims
Bridge axiom	Declared translation rule with input domain, output domain, functional form, evidence IDs, and calibration status	Translation assumption, not observed fact
Structural metric	k-core shell, path membership, pathway counts, and scenario-retention metrics	Screening descriptor, not causal vulnerability or hydrological elasticity
Finance-facing interpretation	RiskMetric nodes, query outputs, and bounded exposure envelopes	Disclosure and diligence support, not actuarial

Concept	Operational representation	Interpretation boundary
		pricing or investment advice
Query governance	Read-only Cypher templates recorded in a manifest with outputs and misuse caveats	Reviewed retrieval surface, not autonomous arbitrary Text2Cypher
Semantic interoperability	Generated RDF/OWL, SHACL-style, PROV-O-aligned, and DCAT/VoID-compatible artifacts	Standards-governed audit and federation contract, not a public linked-data endpoint

Table 1. Conceptual Model. Definition of the graph's core epistemological constraints, specifying the operational representation for key concepts and governing their interpretation boundaries for auditable nature dependency claims.

Results

Evidence substrate and system denominator

The canonical validation snapshot establishes the precise empirical denominator for all subsequent analysis. This materialized structure consists of 225 nodes and 1,807 relationships, rigorously constructed with zero orphan nodes to ensure foundational integrity. Architectural fidelity is secured by the passage of all 33 post-load integrity checks, confirming the graph's invariants. The comprehensive topology includes 105 Taxon nodes, 30 Evidence nodes, and six BridgeAxiom nodes, leveraging 37 distinct relationship types, including 196 SUPPORTS links and 40 TRANSLATES links that explicitly map evidence provenance and cross-domain translation protocols. This meticulously detailed structure imposes a critical epistemological constraint: the system is presented as a reproducible, inspectable test fixture, not a comprehensive, calibrated representation of the entire Aburra Valley watershed. This carefully circumscribed scope, however, is demonstrably sufficient to bridge six critical domains, ensuring that every central assertion is fully auditable and traceable back to the source code, input data, or a manifest of reviewed Cypher queries.

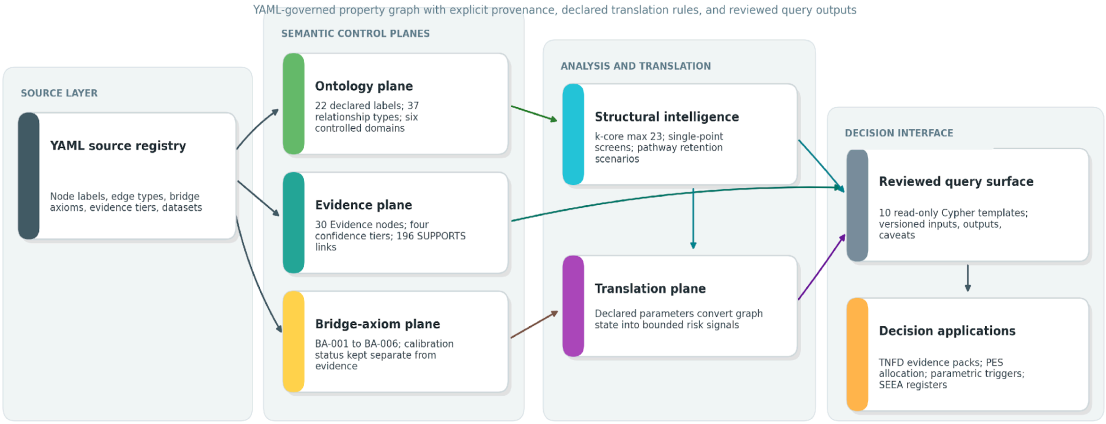

Figure 1. Architectural planes and dependencies. The figure shows how the YAML source registry constrains ontology, evidence, and bridge-axiom control planes. Those planes govern structural intelligence, translation, and reviewed queries, while semantic exports provide the standards-governed audit contract for interoperability and future federation.

A named biological interaction traverses to named water-service infrastructure

The execution of our primary traversal query identifies a canonical service chain that preserves mechanistic continuity across eight distinct directed edges, effectively mapping the dependency route from a discrete ecological interaction to terminal metropolitan infrastructure. This sequence begins with a specialized biological mutualism where *Ensifera ensifera* pollinates *Passiflora mixta* (Lindberg and Olesen, 2001). The chain then traverses the ecological substrate as *Passiflora mixta* contributes to canopy recruitment, which in turn enables fog interception within the high-elevation montane forest. This captured moisture infiltrates through the La Fe catchment, which regulates baseflow to the La Fe reservoir. Moving into the engineered domain, the reservoir supplies the La Ayura treatment plant, which delivers to the EPM distribution network, ultimately serving the Aburra Valley metropolitan area.

Figure 2. Canonical 8-edge service chain
Ensifera ensifera (k=23) to Aburra Valley metro; 13 SUPPORTS across 8 critical-path edges

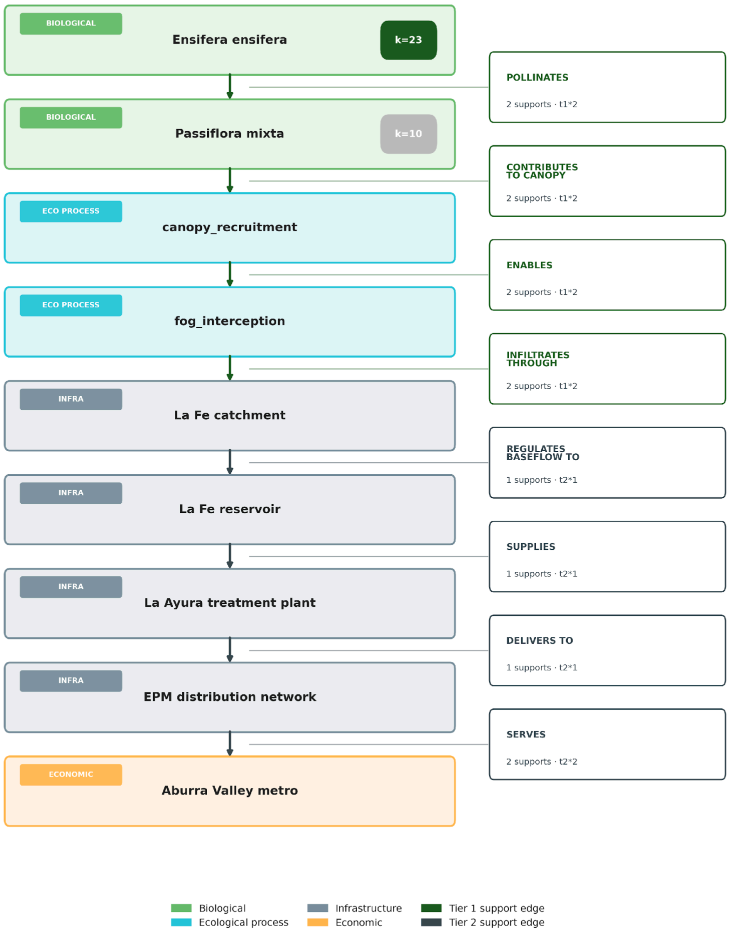

Figure 2. Canonical 8-edge service chain. The chain links a named biological interaction to named catchment, reservoir, treatment, distribution, and metropolitan service nodes. Edge callouts show relationship type and tiered SUPPORTS counts.

Mechanistic substantiation is guaranteed by requiring every critical edge in the service chain to be anchored by at least one relationship-specific SUPPORTS record. The complete eight-edge path is supported by a total of 13 edge-specific SUPPORTS rows, with quantified support across relationship types: two for POLLINATES, two for CONTRIBUTES_TO_CANOPY, two for ENABLES, two for INFILTRATES_THROUGH, one for REGULATES_BASEFLOW_TO, one for SUPPLIES, one for DELIVERS_TO, and two for SERVES. This architectural design fundamentally resolves the ambiguity of node-level attribution by explicitly pinning evidentiary proof to the causal mechanism represented by the relationship. Furthermore, each SUPPORTS record rigorously documents source identifiers, relationship type, evidence scope, confidence tier, and qualitative support notes, thereby ensuring that the entire pathway constitutes a verifiable sequence of auditable claims.

Crucially, this formal traversal does not assert deterministic control by a single biological entity over regional water security. Rather, it is the formal materialization of a verifiable dependency route strictly within the system’s defined denominator. The central scientific utility of this framework resides in its capacity to seamlessly integrate specialized biological interactions, ecological regeneration processes, hydrological mechanisms, and institutional infrastructure assets into a singular, inspectable, and evidence-tiered structural representation.

Structural backbone and k-core interpretation

While the primary traversal confirms a discrete dependency route, a rigorous assessment of the system requires characterizing the topology of the broader ecological network. In this structural metrics snapshot, the biological substrate comprises 105 Taxon nodes and 1,023 validated interaction edges. The initial parser identified 1,037 candidate interactions through the neurosymbolic network reconstruction engine; following de-duplication and type-consistency filtering, 1,023 edges were committed to the k-core computation. The resulting decomposition reveals a maximum shell of 23, with a substantial plateau of fifty taxa, including Ensifera ensifera, occupying this highest core. This topological configuration is interpretatively significant because the graph does not isolate a single keystone species but rather identifies a densely interconnected interaction backbone where numerous high-shell taxa share equivalent structural positions.

This structural density necessitates a shift in visualization governance. Although the eight-edge service chain is legible as a linear path, the biological substrate forms a non-tree-like network that demands interactive inspection rather than static representation. Consequently, the architecture treats ecological visualization as a functional primitive of governance. This approach ensures that reviewers can directly interrogate layer membership, high-shell clusters, evidence attachment, and the surfaces of cross-domain translation without reducing the inherent complexity of the biological system to a mere decorative diagram.

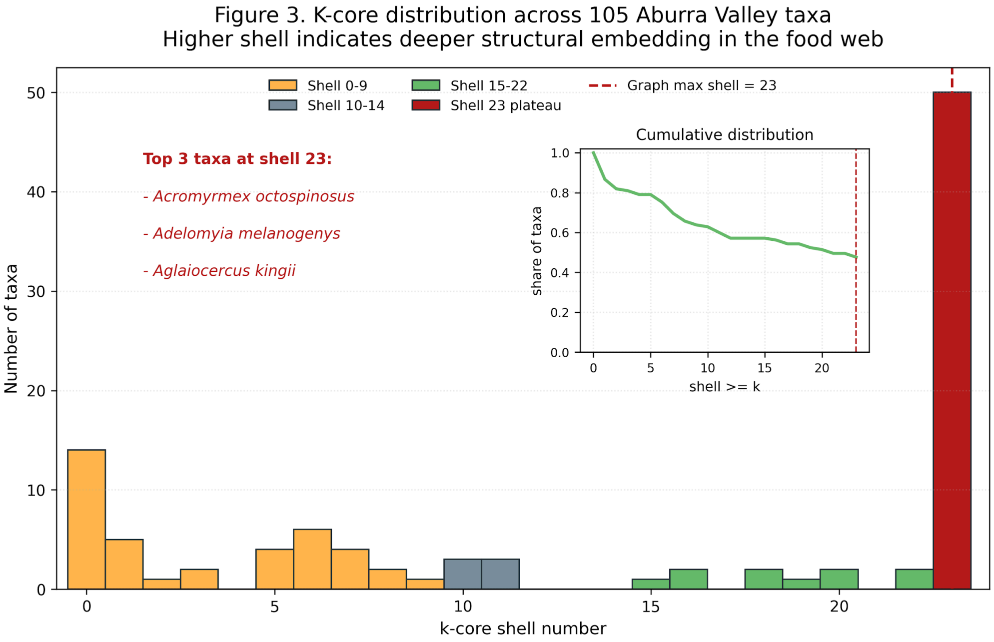

Figure 3. K-core distribution across 105 Aburra Valley taxa. The dashed reference indicates the shell-23 maximum. The broad shell-23 plateau supports structural-backbone interpretation and cautions against single-keystone language.

The k-core decomposition provides a robust measure of embeddedness within the ecological network topology. Its application, however, is subject to rigorous epistemological constraints, consistent with the framework's intent to prevent overclaim. Specifically, the shell value serves as a structural screening primitive, facilitating targeted path review and sensitivity analysis design. It is not an estimation of extinction risk, local demographic vulnerability, interaction strength, or replacement probability. Consequently, a species' structural centrality within the represented substrate must not be interpreted as validated ecological certainty, necessitating subsequent expert review prior to any field or management conclusion. Complementing this structural analysis, the system materializes the fragility_ratio RiskMetric across three distinct stress scenarios: current (1.30), climate-stress (1.28), and urbanization-stress (1.70). Derived through our ecological network inference engine, these ratios quantify the relative structural change in mean k-core shell under each scenario; as with the core k-shell metric, they function exclusively as structural-screening descriptors, not as indices of ecological vulnerability or extinction probability. Furthermore, a complementary single-point-of-failure screening query identifies seven taxa whose represented connectivity to hydrological processes is mediated by a single relationship, predominantly endangered amphibians residing in the k = 23 shell with sole

INDICATES links to water-supply provisioning. These species are categorized as low-redundancy candidates for prioritized field validation, rather than confirmed system bottlenecks.

Bridge-axiom calibration gradient and evidence density

The graph explicitly defines six Bridge Axioms which function as formal, inspectable translation protocols between domains. BA-001 translates low-redundancy pollination and trophic integrity into a theoretical regeneration-risk descriptor and is marked theoretical. BA-002 links regeneration and canopy integrity to cloud-water capture, designated as empirically grounded; while the directional mechanism is supported by tropical-montane-cloud-forest field literature (Bruijnzeel et al., 2011; Villegas et al., 2008), the precise La-Fe-specific elasticities remain a critical calibration gap. BA-003 provides a partially calibrated linkage between soil/catchment integrity and baseflow delivery. BA-004 maps reservoir stress into a partially calibrated treatment-cost sensitivity envelope. BA-005 translates water-service continuity into a partially calibrated economic exposure footprint. BA-006 establishes the institutional linkage connecting finance and governance objects to protected service-chain coverage.

This calibration gradient represents a fundamental result: the architecture explicitly preserves the epistemic heterogeneity of the dependency chain. Diverse inputs, including Tier 1 peer-reviewed ecological literature, Tier 2 institutional records, and partially calibrated translation protocols, contribute to the overall path structure, yet they are not conflated into a unitary confidence metric. Consequently, the differential integrity of evidence density and bridge-axiom calibration status is formalized as an inspectable component of the dependency claim surface. Furthermore, the financial exposure envelope defined by Bridge Axiom BA-004 is formally calculated via the product of annual_treated_volume, treatment_cost_per_m3, and a defined stressed_cost_increase_pct. Specifically, this is instantiated as 518,000 m3/day x 365 x USD 0.12/m3 x {0.15, 0.20}, resulting in a projected range of USD 3.40 million to USD 4.54 million per year. Crucially, the stress multiplier is parameterized by a historical 2015-2016 El Niño operational stress envelope and functions exclusively as a screening parameter, explicitly negating the status of a calibrated response function.

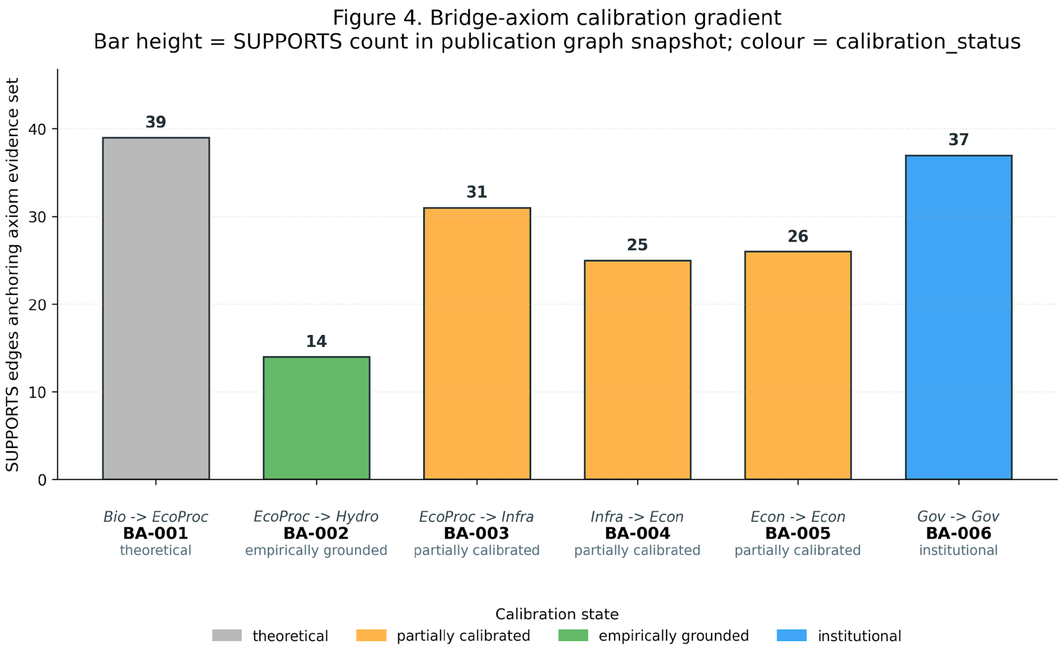

Figure 4. Bridge-axiom calibration gradient. Color indicates calibration status and bar height reports the number of (:Evidence)-[:SUPPORTS]->() relationships in the publication graph snapshot anchored on the evidence nodes named in each axiom's evidence_ids set. The figure separates theoretical, empirically grounded, partially calibrated, and institutional translation layers.

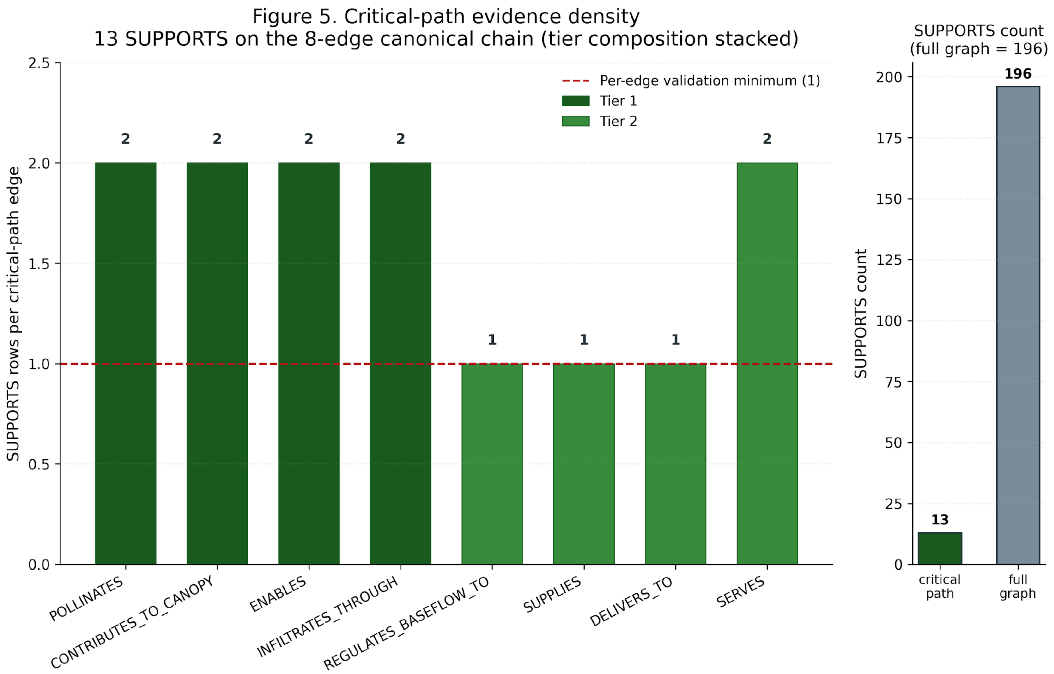

Figure 5. Critical-path evidence density. The figure reports tier composition for the critical-path subset and highlights the validation minimum of one SUPPORTS row per edge. Its denominator is the eight-edge service chain, not the full graph.

The evidence-gap query returns 107 rows, with the first row identifying *Weinmannia tomentosa* as a Taxon node with no direct SUPPORTS relationship. That output should be read as a research-prioritization surface, not as evidence against the entity. Absence or weakness of evidence in the graph indicates where the proof-of-concept denominator needs review, not where ecological truth is absent. Ultimately, the scientific value of this evidence-tiered, axiom-governed architecture resides in its capacity to formally render the complete dependency claim surface, incorporating all identified knowledge gaps and cross-domain translation assumptions, fully inspectable for rigorous audit. By making these uncertainties explicit rather than opaque, the system solidifies the claim constraint, ensuring that every assertion remains anchored to its specific evidentiary and logic-based denominator.

Protocol-governed structural screening: compound-stress scenario and retained-integrity matrix

This section details the application of the structural backbone analysis from the k-core decomposition to a formal screening scenario, establishing a protocol for calculating retained integrity, that is the percentage of 166 baseline species-to-reservoir paths that survive transparent k-core threshold filters used as non-causal proxies for compound stress. Following the architecture's protocol for structural screening, the default compound-stress query uses transparent k-core threshold filters as screening proxies. The baseline species-to-reservoir denominator contains 166 paths across 29 species and three reservoirs. The urbanization branch retains 60 paths and loses 63.9% of paths. The climate branch retains 160 paths and loses 3.6%. The compound branch retains 54 paths and loses 67.5%, leaving 32.5% retained integrity. Under the current threshold rules, the classifier branch is additive because the urbanization-removed ($k < 5$) and climate-removed ($10 <= k <= 18$) species sets are disjoint by construction; the additive label therefore reports the geometry of the threshold partition, not an empirical interaction structure. Different threshold definitions could yield synergistic or buffered branches. The matrix is deliberately framed as a screening envelope. It does not model land-use dynamics, climate exposure, species demography, hydrological response, or reservoir operation. Its value is transparency: reviewers can see the thresholds, path denominator, residual paths, retained-integrity percentage, and classifier rule. That makes the scenario contestable and improvable.

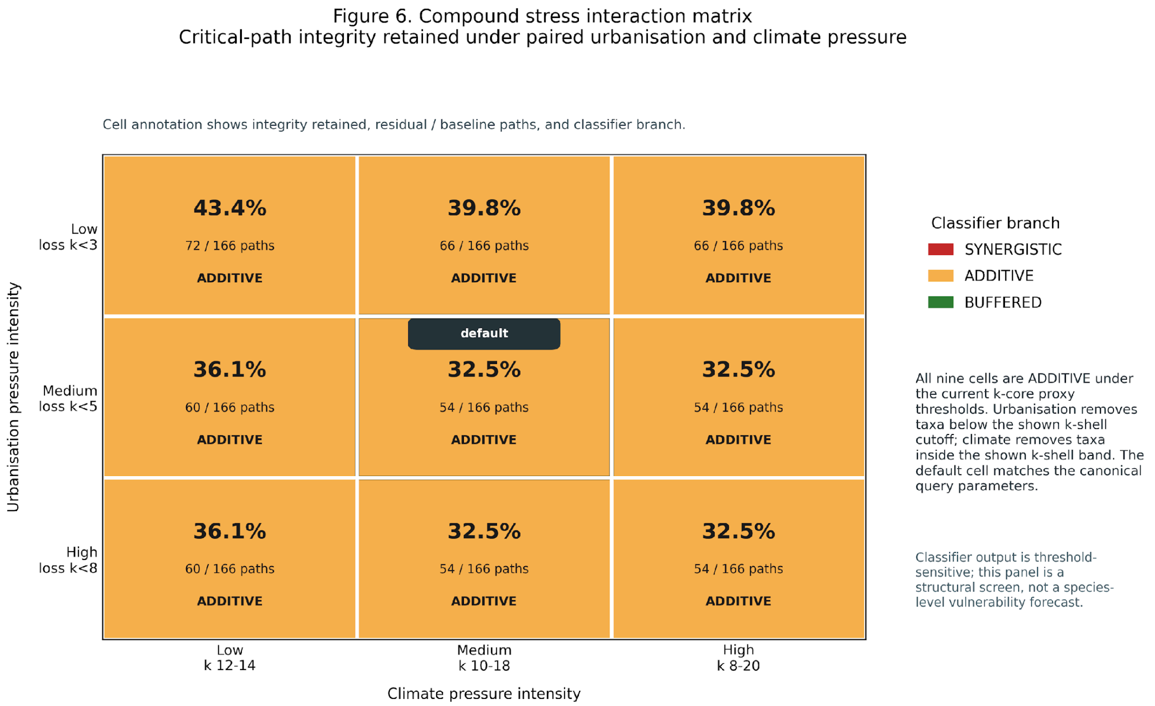

Figure 6. Compound stress interaction matrix. Cell color reports classifier branch and text reports retained integrity and residual/baseline path counts. The medium-medium default cell retains 54 of 166 baseline species-to-reservoir paths.

Validation, query governance, and semantic interoperability as claim constraints

Validation is an explicit component of the scientific claim, enforced by 33 post-load Cypher checks that operate as a SHACL-analogous runtime contract. These checks confirm critical invariants such as graph integrity, critical-path existence, provenance coverage, and semantic-export readiness after materialization. The query manifest is a key governance element, recording ten read-only Cypher templates with their intended purpose, expected outputs, manuscript claim boundaries, and misuse caveats. Further, semantic exports constitute a second audit layer. Specifically, the system, while running on a Neo4j property graph, generates artifacts aligned with RDF/OWL, SHACL, PROV-O, and DCAT/VoID standards. This formalization includes mapping Evidence records to prov:Entity and bridge axioms to prov:Plan, with support/translation links utilizing provenance concepts like prov:wasDerivedFrom. These exports, accompanied by a pySHACL validation report, ensure the ontology's legibility for Semantic Web reviewers without asserting public linked-data deployment.

Interpretation Protocol

This graph serves as a standardized protocol for reviewing evidence. To use it, reviewers should first verify the data snapshot, which in this case is set at the May 2026 version, which defines the scope of all claims. Second, reviewers must check the "epistemic status" of each connection: while some links are based on peer-reviewed science or utility records, others rely on AI-powered inference and mathematical models, or theoretical assumptions. The system's primary value is making these differences in evidence quality visible. Finally, the graph distinguishes between observed facts and computed estimates, ensuring that path existence or statistical scores are treated as screening tools rather than absolute causal proof. This transparent structure allows experts—from engineers to ecologists—to pinpoint and challenge specific data points or assumptions, turning complex nature-risk debates into a concrete, auditable evidence agenda.

Research, Governance, and Infrastructure Implications

The system's architecture demonstrates distinct implications across multiple domains, creating a seamless route from fundamental ecological data to applied governance. For biodiversity informatics, it establishes a clear path from species interaction data to downstream service-chain interpretation, preserving critical details like taxon identity, edge type, and evidence. Building on this, the system benefits ecological network science by showing how k-core and path analysis can be used conservatively as structural measures within a cross-domain graph, rather than being converted into unsupported vulnerability claims. This structural clarity extends to ecohydrology, making the translation from canopy and catchment processes to infrastructure explicit enough to identify necessary calibration gaps. Consequently, for watershed governance, the graph supports the creation of evidence packs around named infrastructure assets. Moreover, a utility or public agency can use this to quickly pinpoint which ecological objects support an asset like La Fe, which evidence tiers support each step, which bridge axioms remain theoretical or partially calibrated, and which monitoring data would most improve confidence. This constitutes a governance use because it defines a validation agenda, rather than automating a decision.

Finally, for nature finance, the practical contribution of this architecture is disciplined translation, given that the graph is capable of representing candidate trigger species, catchments, reservoirs, interventions, RiskMetric nodes, and evidence gaps before any financial instrument is structured, thereby supporting diligence language, scenario screening, restoration prioritization, and the Locate and Evaluate phases of TNFD LEAP. However, it is essential to note that this architecture is *not* a substitute for the Assess and Prepare phases, which require institutional review, calibrated response functions, and event-level data; therefore, it should not be used as pricing logic, underwriting evidence, regulatory-disclosure evidence, or an investment recommendation without adequate empirical calibration and institutional review.

Design Principles

The system's integrity and auditable constraints are enforced by seven core design principles. These principles establish the mandatory epistemological boundaries and structural requirements for representing nature dependency without collapsing complexity or overclaiming predictive certainty. The first design principle is *Preservation of Ontological Distinctness*. Biological Taxa, Infrastructure Assets, and Bridge Axioms must retain their distinct graph identity until a formal, declared relationship makes a translation explicit. Collapsing these heterogeneous domains into a singular, ambiguous nature score compromises the traceability and auditability of the underlying claims. The second principle is *Epistemological Separation of Evidence and Translation*. A relationship's support (Evidence) must be formally distinguished from the cross-domain rule that interprets its implication (Bridge Axiom). This mandatory separation maintains the reviewability of the claim, allowing the system to explicitly tag axioms as theoretical, empirically grounded, or institutional. The third principle is *Conservative Epistemology in Structural Metrics*. Structural intelligence tools, such as k-core, pathway counts, and stress retention metrics, function solely as screening primitives and topological descriptors. They must not be conflated with field-validated demographic vulnerability, causal hydrological prediction, or actuarial loss estimation. The fourth principle is *Query Governance as Claim Constraint*. All externalized and publication-grade results must originate from a reviewed, read-only Query Manifest. This protocol defines the claim surface and prevents unconstrained execution (e.g., via arbitrary Text2Cypher), thereby enforcing known parameters, caveats, and misuse boundaries. The fifth principle is *Pragmatic Deployment with a Standards-Governed Audit Contract*. The use of Neo4j prioritizes operational performance and traversal capabilities. Governance, however, is anchored by standards-aligned semantic exports (RDF/OWL, SHACL, PROV-O), which provide a formal, machine-readable audit contract for interoperability and future federation without asserting a public linked-data service. The sixth principle is *Inspectable Schema Extensibility*. New relationship types, node labels, and bridge axioms must be registered in the YAML-based governance system, ensuring that schema extension remains explicit and inspectable. This process maintains a formal TBox concept for external audit and ensures consistency across the runtime graph and semantic exports. The seventh principle is *Network Visualization as a Governance Primitive*. For dense or non-tree-like substrates (e.g., the biological interaction network or k-core plateau), interactive visualization tools must be provided. This ensures that reviewers can inspect high-shell membership, layer controls, and evidence attachment, thereby preventing the complexity of the full graph from being obscured by over-simplified static figures.

Translation Scenario

The architecture transforms general nature-risk statements into a disciplined institutional due-diligence protocol. For a watershed utility preparing an evidence pack for an asset like the La Fe reservoir or La Ayura treatment plant, a reviewer can execute the headline and evidence-gap templates to retrieve the full dependency path from a named ecological interaction to the service boundary. Crucially, the system's output is not a definitive action plan but a robust audit surface that organizes critical questions: Which biological and hydrological steps are substantiated by Tier 1 literature? Which infrastructure links rely on institutional records? Which bridge-axiom translations remain partially calibrated assumptions? By making these evidentiary and logical gaps explicit, the graph shifts the focus from qualitative risk assessment to a concrete, traceable evidence agenda, fundamentally supporting rigorous governance and validation planning.

This same architecture can provide a robust mechanism for conservation-finance teams seeking to optimize restoration prioritization. Moving beyond conventional metrics like species richness or hectares, the graph enables a multi-criteria inquiry: does a proposed intervention intersect high-shell taxa (k-core analysis), critical canopy recruitment processes, and an existing reservoir service path, while also addressing a documented evidence gap? This structured approach demands a fusion of ecological expertise and governance review, ensuring strategic deployment. Ultimately, the architecture's most significant contribution is to enforce the reviewability of the underlying ecological mechanism and its associated uncertainties before financial instruments are structured or investment claims are attached, establishing a rigorous pre-contractual due-diligence framework.

Discussion

The Medellin proof of concept demonstrates that nature dependency is computable when represented as a named-object graph rather than an opaque, aggregated score. By preserving the identity of each biological interaction, infrastructure asset, and evidentiary link, the architecture moves beyond narrative assertion into a formal, inspectable claim surface. The value of this approach lies in its capacity to organize uncertainty, as it does not claim to predict precise hydrological yields or actuarial losses, but instead materializes the specific evidence gaps and translation assumptions that a reviewer must evaluate to reach a defensible conclusion. This distinction between structural connectivity and causal certainty is a primary methodological strength. For instance, identifying *Ensifera ensifera* as a high-shell occupant in a service-chain path provides a rigorous screening primitive for risk disclosure, without overclaiming demographic vulnerability. In an environment of evolving nature-related standards, such as TNFD, the honesty of this representation sustains auditability, because it exposes exactly where the logic of a dependency claim is anchored in literature and where it relies on theoretical bridge axioms requiring further calibration.

Practically, this framework has the potential to transform nature-risk debates into concrete diligence agendas. For institutions, the graph functions as a connective tissue that allows ecologists, engineers, and financial analysts to inspect the same causal route from different professional vantage points. By providing a standards-aligned audit trail, it shifts the focus from estimating nature in the abstract to managing specific, named ecological-to-economic dependencies. This makes the architecture a robust foundation for building institutional trust in natural-capital registers and disclosures. The immediate horizon involves populating this inspectable structure with longitudinal monitoring and field-validated records. In this sense, the graph is not a static model to be believed, but a dynamic fixture designed to be improved, with every new data point strengthened at the node or edge level increasing the overall confidence of the terminal financial interpretation. Ultimately, the contribution is a disciplined architecture that sustains the complexity of nature while making it legible for rigorous institutional governance.

Future Analysis Extensions

The most critical extension is the expert ecological validation of the AI-powered ecological network reconstruction, which provides the biological substrate. Regional ecologists should review high-shell taxa and critical-chain interactions against local records to distinguish graph-structural importance from verified interaction and persistence. This empirical anchoring should extend to longitudinal hydrology and operations data, integrating event-level rainfall, turbidity, and treatment-cost time series to calibrate bridge axioms BA-002 through BA-005 against observed events rather than static parameters. Furthermore, future versions must replace threshold-only scenarios with explicit land-cover, climate, and species-vulnerability response functions attached to graph objects, ensuring stronger calibration while maintaining architectural inspectability.

To test the framework's falsifiability, multi-watershed replication will evaluate whether the audit pattern remains coherent as taxa, infrastructure, and evidence density change. This expansion necessitates a transition toward public semantic namespace governance, replacing internal URNs with dereferenceable IRIs and linked-data endpoints once provenance, update policies, and data-rights obligations are settled. These steps could collectively move the architecture from a single-watershed proof of concept toward a robust, standards-governed system for computable nature dependency.

Limitations and Threats to Validity

The study’s primary limitations stem from its bounded scope and the nature of the represented data. Centered on the La Fe service chain, the graph functions as a proof of concept rather than a comprehensive census of the Aburra Valley's ecological or hydrological processes. The biological substrate relies on AI-powered ecological network inference, where structural metrics like k-core shells serve as research prioritization tools rather than field-validated indicators of species vulnerability or demographic persistence. Geographically, the evidence distinguishes between site-specific local data and regional or neotropical analogues, ensuring that mechanism analogues are not mistaken for local validation.

Causal and construct validity are constrained by the absence of counterfactual identification and calibrated hydrological response models. While topological counts and query outputs are directly recomputed, claims regarding watershed function require additional empirical evidence. The directed paths and bridge axioms declare dependency structures and translation assumptions rather than proving deterministic control. Notably, while segments like reservoir supply and treatment delivery are supported by institutional reporting, critical calibration gaps remain in the peer-reviewed catchment-yield and treatment-cost time series required to move beyond screening-grade descriptors.

Finally, the system's financial and external validity is deliberately bounded. Risk metrics and exposure footprints, such as the USD 12.6 billion GDP exposure, utilize weighted sector outputs and screening parameters rather than realized losses or actuarial pricing. Replicating the architecture in other watersheds would preserve the audit pattern but require entirely new local datasets. To prevent misuse, the manuscript emphasizes that a result is only as mature as its underlying evidence tier and calibration status; without institutional sign-off and empirical recalibration, these outputs should not be used as investment recommendations or regulatory disclosure evidence.

STAR Methods / Methods

Resource Availability

Materials Availability. No physical materials were generated. Digital materials include graph source data, processed artifacts, ontology YAML files, evidence registries, Cypher templates, semantic exports, figure scripts, generated figures, validation results, and manuscript package files.

Data and Code Availability. Code, source registries, query templates, semantic exports, figures, and reproducibility notes are maintained in the project repository and mirrored in this manuscript package. Code upon request. Public archive metadata and final repository DOI should be attached when the manuscript is frozen for external distribution.

Study design and system boundary

The study is a reproducible computational proof of concept for a single watershed knowledge graph. The graph represents selected objects relevant to the La Fe water-service chain across biological, ecological-process, infrastructure, economic, governance, evidence, bridge-axiom, and dataset domains. The canonical headline path is defined in code and validated after rebuild. La Fe catchment ranges approximately 2,000-3,000 m elevation in the Aburra Valley region; cloud-forest mechanism literature is applied as the regional framework for fog interception, not as a verified site characterization of La Fe specifically.

Data sources and evidence generation

Inputs include an AI-assisted ecological interaction substrate (the parser yields 1,037 candidate edges, of which 1,023 enter the structural-metrics computation after de-duplication and typing), structural metrics, stress-test files, a literature evidence registry, institutional records from EPM, Cuenca Verde, and DANE, nature-finance and natural-capital framework references, and ontology registries for node labels, relationship types, semantic prefixes, bridge axioms, evidence tiers, and query governance. The supplementary geographic-evidence pack classifies evidence sources by geographic provenance so that local, regional, institutional, and mechanism-analogue support remain visible.

Canonical graph model

The runtime graph is a Neo4j property graph selected for pragmatic deployment, performance, Cypher traversal, visualization integration, and compatibility with graph data-science tooling. Nodes carry stable node_id identifiers and labels such as Taxon, HydrologicalProcess, Catchment, Reservoir, TreatmentPlant, MetropolitanArea, Evidence, BridgeAxiom, Dataset, and RiskMetric. Relationships encode biological interactions, ecological processes, hydro-infrastructure flow, governance, finance, evidence support, and bridge translation. Loading uses parameterized UNWIND batch operations and programmatic edge creation where relationships are deterministic from source registries. Some committed Cypher templates carry :param directives interpreted by cypher-shell; programmatic execution by the Python Neo4j driver requires equivalent parameter substitution. Runtime pragmatism does not replace semantic governance: node labels, relationship types, bridge axioms, evidence tiers, dataset catalog entries, and semantic prefixes remain registry-governed and are exported for standards-aligned audit.

Evidence tiers and SUPPORTS relationships

Evidence nodes store source metadata, confidence tier, source URI, retrieval date, curation method, derived-from URI, confidence basis, and provenance role. SUPPORTS relationships attach evidence to graph objects and, for critical-path claims, carry relationship-specific metadata: supported source identifier, supported target identifier, supported relationship type, support scope, and support note. The supplementary geographic-evidence pack classifies every Evidence node by provenance category and intended use. This keeps local institutional sources, Colombian or Andean mechanism sources, global syntheses, methodological sources, and non-Andean analogues visible during review.

Bridge axioms

The system's translation protocols, formalized as Bridge Axioms, are defined within a central governance registry. This registry records the input and output domains, calibration status, key uncertainties, and supporting evidence identifiers for each axiom. The system materializes these definitions as distinct nodes within the graph connected by explicit translation links, ensuring that every cross-domain mapping is traceable. These bridge axioms are strictly interpreted as declared translation rules and logical assumptions rather than observed physical facts.

Schema governance and extension

The YAML registries are the schema-control surface for technical users. Node labels and properties are governed by ontology/node-labels.yaml; relationship types, source/target constraints, and relationship properties are governed by ontology/edge-types.yaml; bridge axioms and critical-path SUPPORTS metadata are governed by ontology/bridge-axioms.yaml; evidence tiers and dataset records are governed by config/evidence-tiers.yaml and data/catalog.yaml. A new bridge axiom is introduced by editing the registry and rerunning rebuild, validation, and semantic export; its calibration status remains part of the axiom definition. New node labels or relationship types are similarly registered for semantic audit, but runtime materialization also requires aligned ETL, processed-artifact, and schema changes. This exposes control to the knowledge architect while keeping changes subject to the same validation and provenance contract as the original graph.

Query governance

The query interface is restricted to a set of reviewed, read-only templates governed by a central manifest. This manifest defines the specific purpose, intended parameters, and manuscript claims supported by each query, while explicitly documenting misuse caveats and necessary data requirements. Automated tests ensure that all accessible queries are properly documented and strictly avoid unauthorized data modification operations.

K-core decomposition

K-core decomposition follows the standard graph-theoretic definition of a maximal subgraph in which all nodes have degree at least k (Batagelj and Zaversnik, 2003). Shell values are computed in the biological structural-metrics substrate and written to Taxon nodes. In this manuscript they are interpreted only as structural embeddedness indicators.

Semantic interoperability exports

This process ensures the integrity of the system by generating official audit files from the core governing data (the ontology, evidence, and translation rules). Although the main database is used for performance, these standardized exports align with international protocols for data provenance and interoperability. This allows external reviewers to formally audit and validate the system's structure and data against established governance standards, ensuring that the claims are transparent and verifiable.

Quantification and statistical analysis

No inferential statistics are reported. Quantification consists of graph counts, relationship counts, validation counts, test counts, k-core shell values, path counts, retained-integrity percentages, evidence counts, and RiskMetric values. Uncertainty is represented through evidence tiers, bridge-axiom calibration status, declared parameters, and limitations rather than confidence intervals.

Acknowledgments

The author acknowledges the public and institutional resources cited in the manuscript, including GBIF, GloBI, TNFD, ENCORE, SEEA Ecosystem Accounting, EPM, Cuenca Verde, and DANE.

Author Contributions

Jay Gutierrez: conceptualization, methodology, software, validation, formal analysis, investigation, data curation, visualization, writing, original draft, writing, review and editing.

Declaration of Interests

The author may pursue commercial development of watershed knowledge-graph, nature-risk screening, and natural-capital diligence tools based on the architecture described here. No incorporated entity, customer contract, external investor financing, or paid counterparty engagement is declared for this manuscript package.

Funding

No external funding was declared for this manuscript package.

References

Abrahamczyk S, Souto-Vilaros D, Renner SS. 2014. Escape from extreme specialization: passionflowers, bats and the sword-billed hummingbird. *Proceedings of the Royal Society B: Biological Sciences* 281:20140888. https://doi.org/10.1098/rspb.2014.0888.

Allesina S, Pascual M. 2009. Googling food webs: can an eigenvector measure species' importance for coextinctions? *PLOS Computational Biology* 5:e1000494. https://doi.org/10.1371/journal.pcbi.1000494.

Bascompte J, Jordano P. 2007. Plant-animal mutualistic networks: the architecture of biodiversity. *Annual Review of Ecology, Evolution, and Systematics* 38:567-593. https://doi.org/10.1146/annurev.ecolsys.38.091206.095818.

Atkins JL, Long RA, Pansu J, Daskin JH, Potter AB, Stalmans ME, Tarnita CE, Pringle RM. 2019. Cascading impacts of large-carnivore extirpation in an African ecosystem. *Science* 364:173-177. https://doi.org/10.1126/science.aau3561.

Batagelj V, Zaversnik M. 2003. An O(m) algorithm for cores decomposition of networks. arXiv:cs/0310049. https://doi.org/10.48550/arXiv.cs/0310049.

Bruijnzeel LA, Mulligan M, Scatena FN. 2011. Hydrometeorology of tropical montane cloud forests: emerging patterns. *Hydrological Processes* 25:465-498. https://doi.org/10.1002/hyp.7974.

Buttigieg PL, Morrison N, Smith B, Mungall CJ, Lewis SE, ENVO Consortium. 2013. The environment ontology: contextualising biological and biomedical entities. *Journal of Biomedical Semantics* 4:43. https://doi.org/10.1186/2041-1480-4-43.

Cavelier J, Solis D, Jaramillo MA. 1996. Fog interception in montane forests across the Central Cordillera of Panama. *Journal of Tropical Ecology* 12:357-369. https://doi.org/10.1017/S026646740000955X.

Chaplin-Kramer R, Sharp RP, Weil C, Bennett EM, Pascual U, Arkema KK, Brauman KA, Bryant BP, Guerry AD, Haddad NM, Hamann M, Hamel P, Johnson JA, Mandle L, Pereira HM, Polasky S, Ruckelshaus M, Shaw MR, Silver JM, Vogl AL, Daily GC. 2019. Global modeling of nature's contributions to people. *Science* 366:255-258. https://doi.org/10.1126/science.aaw3372.

Cuenca Verde. 2023. *Informe de Gestion 2023*. Medellin, Colombia: Cuenca Verde Fund. https://www.cuencaverde.org/informes-de-gestion/2023.

Dasgupta P. 2021. *The Economics of Biodiversity: The Dasgupta Review*. London: HM Treasury. https://www.gov.uk/government/publications/final-report-the-economics-of-biodiversity-the-dasgupta-review.

Dimitrova M, Senderov VE, Georgiev T, Zhelezov G, Penev L. 2021. Infrastructure and population of the OpenBiodiv biodiversity knowledge graph. *Biodiversity Data Journal* 9:e67671. https://doi.org/10.3897/BDJ.9.e67671.

Dunne JA, Williams RJ, Martinez ND. 2002. Network structure and biodiversity loss in food webs: robustness increases with connectance. *Ecology Letters* 5:558-567. https://doi.org/10.1046/j.1461-0248.2002.00354.x.

Empresas Publicas de Medellin. 2023. *Informe de Gestion 2023*. Medellin, Colombia: EPM Group. https://www.epm.com.co/content/dam/epm/institucional/transparencia/rendicion-de-cuentas/rendicion-de-cuentas-2023/informe-de-gestion-version-final.pdf.

Estes JA, Terborgh J, Brashares JS, Power ME, Berger J, Bond WJ, Carpenter SR, Essington TE, Holt RD, Jackson JBC, et al. 2011. Trophic downgrading of planet Earth. *Science* 333:301-306. https://doi.org/10.1126/science.1205106.

Fragoso C, Lavelle P. 1992. Earthworm communities of tropical rain forests. *Soil Biology and Biochemistry* 24:1397-1408. https://doi.org/10.1016/0038-0717(92)90124-G.

GBIF Secretariat. 2026. GBIF.org. Accessed 29 April 2026. https://www.gbif.org.

Hawkins H-J, Cargill RIM, Van Nuland ME, Hagen SC, Field KJ, Sheldrake M, Soudzilovskaia NA, Kiers ET. 2023. Mycorrhizal mycelium as a global carbon pool. *Current Biology* 33:R560-R573. https://doi.org/10.1016/j.cub.2023.02.027.

Helmer EH, Gerson EA, Baggett LS, Bird BJ, Ruzycki TS, Voggesser SM. 2019. Neotropical cloud forests and paramo to contract and dry from declines in cloud immersion and frost. *PLOS ONE* 14:e0213155. https://doi.org/10.1371/journal.pone.0213155.

Hobbs RJ, Gutierrez Betancur JA. 2024. *The Biological Quark: Structural Risk Intelligence for Nature-Economy Systems*. Graph of Life Working Paper. Unpublished framework source; not peer reviewed.

Jaramillo A, Sharpe CJ. 2020. Antioquia Brushfinch (*Atlapetes blancae*), version 1.0. In: del Hoyo J, Elliott A, Sargatal J, Christie DA, de Juana E, eds. *Birds of the World*. Ithaca, NY: Cornell Lab of Ornithology. https://doi.org/10.2173/bow.antbrf1.01.

Kappelle M, ed. 2006. *Ecology and Conservation of Neotropical Montane Oak Forests*. Berlin: Springer-Verlag. Chapter 11: Pulido MT, Cavelier J, Cortes-S SP. Structure and Composition of Colombian Montane Oak Forests, pp. 141-152. https://doi.org/10.1007/3-540-28909-7_11.

Klein A-M, Vaissiere BE, Cane JH, Steffan-Dewenter I, Cunningham SA, Kremen C, Tscharntke T. 2007. Importance of pollinators in changing landscapes for world crops. *Proceedings of the Royal Society B: Biological Sciences* 274:303-313. https://doi.org/10.1098/rspb.2006.3721.

Kushwaha N, Singh A, Sheikh HA. 2025. NatureKG: an ontology and knowledge graph for nature finance with a Text2Cypher application. *Frontiers in Artificial Intelligence* 8:1693843. https://doi.org/10.3389/frai.2025.1693843.

Lindberg AB, Olesen JM. 2001. The fragility of extreme specialization: *Passiflora mixta* and its pollinating hummingbird *Ensifera ensifera*. *Journal of Tropical Ecology* 17:323-329. https://doi.org/10.1017/S0266467401001213.

Meyer ST, Leal IR, Tabarelli M, Wirth R. 2011. Leaf-cutting ants as ecosystem engineers: topsoil and litter perturbations around nests. *Ecological Entomology* 36:14-24. https://doi.org/10.1111/j.1365-2311.2010.01241.x.

Natural Capital Finance Alliance. 2024. ENCORE: Exploring Natural Capital Opportunities, Risks and Exposure. Global Canopy, UNEP FI and UNEP-WCMC. https://doi.org/10.34892/dz3x-y059.

Nichols E, Spector S, Louzada J, Larsen T, Amezquita S, Favila ME, The Scarabaeinae Research Network. 2008. Ecological functions and ecosystem services provided by Scarabaeinae dung beetles. *Biological Conservation* 141:1461-1474. https://doi.org/10.1016/j.biocon.2008.04.011.

Penev L, Dimitrova M, Senderov V, Zhelezov G, Georgiev T, Stoev P, Simov K. 2019. OpenBiodiv: a knowledge graph for literature-extracted linked open data in biodiversity science. *Publications* 7:38. https://doi.org/10.3390/publications7020038.

Poelen JH, Simons JD, Mungall CJ. 2014. Global biotic interactions: an open infrastructure to share and analyze species-interaction datasets. *Ecological Informatics* 24:148-159. https://doi.org/10.1016/j.ecoinf.2014.08.005.

Still CJ, Foster PN, Schneider SH. 1999. Simulating the effects of climate change on tropical montane cloud forests. *Nature* 398:608-610. https://doi.org/10.1038/19293.

Stuart SN, Chanson JS, Cox NA, Young BE, Rodrigues ASL, Fischman DL, Waller RW. 2004. Status and trends of amphibian declines and extinctions worldwide. *Science* 306:1783-1786. https://doi.org/10.1126/science.1103538.

Swiss Re Institute, The Nature Conservancy. 2022. *Insuring Nature: Parametric Insurance for Ecosystem Services*. Zurich: Swiss Re Institute. https://www.swissre.com/institute/research/topics-and-risk-dialogues/climate-and-natural-catastrophe-risk/insuring-nature.html.

Taskforce on Nature-related Financial Disclosures. 2023. *Recommendations of the Taskforce on Nature-related Financial Disclosures*. London: TNFD. https://tnfd.global/publication/recommendations-of-the-taskforce-on-nature-related-financial-disclosures/.

United Nations, European Union, Food and Agriculture Organization, International Monetary Fund, OECD, World Bank Group. 2021. *System of Environmental-Economic Accounting–Ecosystem Accounting*. New York: United Nations Statistical Division. https://seea.un.org/ecosystem-accounting.

van der Heijden MGA, Martin FM, Selosse M-A, Sanders IR. 2015. Mycorrhizal ecology and evolution: the past, the present, and the future. *New Phytologist* 205:1406-1423. https://doi.org/10.1111/nph.13288.

Vanbergen AJ, Woodcock BA, Heard MS, Chapman DS. 2017. Network size, structure and mutualism dependence affect the propensity for plant-pollinator extinction cascades. *Functional Ecology* 31:1285-1293. https://doi.org/10.1111/1365-2435.12823.

Villa F, Bagstad KJ, Voigt B, Johnson GW, Portela R, Honzak M, Batker D. 2014. A methodology for adaptable and robust ecosystem services assessment. *PLOS ONE* 9:e91001. https://doi.org/10.1371/journal.pone.0091001.

Villegas JC, Tobon C, Breshears DD. 2008. Fog interception by non-vascular epiphytes in tropical montane cloud forests: dependencies on gauge type and meteorological conditions. *Hydrological Processes* 22:2484-2492. https://doi.org/10.1002/hyp.6844.

Walls RL, Deck J, Guralnick R, Baskauf S, Beaman R, Blum S, Bowers S, Buttigieg PL, Davies N, Endresen D, Gandolfo MA, Hanner R, Janning A, Krishtalka L, Matsunaga A, Midford P, Morrison N, O Tuama E, Schildhauer M, Smith B, Stucky BJ, Thomer A, Wieczorek J, Whitacre J, Wooley J. 2014. Semantics in support of biodiversity knowledge discovery: an introduction to the Biological Collections Ontology and related ontologies. *PLOS ONE* 9:e89606. https://doi.org/10.1371/journal.pone.0089606.