DECENTRALIZED MULTI-CHAIN INFRASTRUCTURE FOR GS1 PRODUCT IDENTIFIER RESOLUTION

RICARDO BROWN, LUKMAN HAKIM, GOLD DAVID CHUKWUMAOBI, ABIDOYE O. ANTHONY,
JEFRESON B. MARAGGUN, SALMAN FARIZ, ISAAC A. SIMON, AND NAVOD ANUHAS
APRIL LABS

Abstract. Product identifier resolution is critical infrastructure for supply chain transparency, food safety systems, and digital product passports. Current resolution systems rely on centralized services operated by standards bodies or commercial entities, creating single points of failure and access restrictions. We present a decentralized, multi-chain resolver mesh for GS1 identifiers (such as GTINs) that provides open, non-discriminatory access to product information without requiring tokens or fees. Our architecture utilizes a hybrid on-chain/off-chain model of deterministic smart contracts that anchors the ownership and existence of product identifiers, while data resolution and retrieval occur via a distributed peer-to-peer mesh. This separates the Root of Trust from the Data Transport, ensuring scalability without gas costs for data updates. The system integrates with open product databases including Open Food Facts, enabling transparent access to community-maintained product information. The resolver operates independently of governance tokens and maintains strict neutrality regarding data sources and network preferences. We describe the technical architecture, security model, and operational guarantees of the resolver mesh. The system is designed as public digital infrastructure with explicit commitments to openness and non-discrimination, with economic sustainability enabled through emerging protocol public goods funding mechanisms rather than commercial extraction. We analyze the security properties, discuss deployment considerations across heterogeneous blockchain networks, and evaluate the system's resilience to various failure modes. This work contributes to the growing body of literature on decentralized public goods infrastructure and demonstrates practical approaches to multi-chain service deployment. The resolver mesh has been deployed and is operational across 10 blockchain networks.

Subject Areas: Cryptography and Security; Distributed Systems; Networking and Internet Architecture

Keywords: blockchain, supply chain, product identification, GS1, decentralized infrastructure, public goods, multi-chain systems

1 Introduction

1.1**Motivation** Product identifiers such as GTINs (Global Trade Item Numbers) are foundational to modern commerce, enabling unique identification of products across supply chains. Resolution of these identifiers—the process of converting an identifier into actionable product information—is critical infrastructure for:

- Supply chain transparency and traceability
- Consumer product information access
- Food safety and recall systems
- Sustainability and ethical sourcing verification
- Digital product passports (DPPs) for circular economy initiatives

Current resolution infrastructure exhibits several limitations: including centralization risks, access barriers, opacity, and geographic limitations. These limitations are particularly problematic for public-interest applications such as food transparency platforms and open-source supply chain tools.

1.2**Contributions**

1. **Architecture:** A multi-chain resolver mesh architecture providing redundant, open access across heterogeneous networks.
2. **Implementation:** Technical implementation details including deterministic contract deployment and resolution protocols.
3. **Security Analysis:** Analysis of the trust model, security properties, and failure modes.
4. **Public Goods Framework:** Design principles for blockchain-based public infrastructure prioritizing neutrality and sustainability.
5. **Deployment Evidence:** Documentation of a live deployment across 10 blockchain networks.

Date: January 2026.

1.3**Paper Organization** The remainder of this paper is organized as follows. Section 2 reviews related work in product identification standards, decentralized identifier systems, blockchain applications in supply chain management, multi-chain infrastructure, public goods theory and protocol public goods. Section 3 describes the system architecture, including design principles, core components, and the multi-chain deployment strategy. Section 4 presents a comprehensive security analysis, examining the trust model, attack vectors, privacy considerations, and economic security properties. Section 5 discusses implementation details and deployment experience. Section 6 provides critical discussion of limitations, identifies promising directions for future work, and explores broader implications for decentralized public infrastructure. Section 7 concludes.

2 Related Work

2.1**Product Identification Standards** The GS1 system of standards[1] defines how identification keys, data attributes, and barcodes must be used in business applications.[2] The foundational GS1 General Specifications, such as Version 25 released in January 2025, constitutes a 522-page framework governing the identification of items and logistics units.[1] These standards serve as a product's data "DNA," enabling organizations to identify, capture, and share information globally.[3] The identification keys utilized to categorize these items are summarized in Table 1.

Table 1. Primary GS1 Identification Keys and Use Cases.

ID Key	**Used to Identify**	**Example**
Global Trade Item Number (GTIN)	Products and services	Can of soup, chocolate bar
Global Location Number (GLN)	Parties and locations	Companies, warehouses, factories
Serial Shipping Container Code (SSCC)	Logistics units	Unit loads on pallets, parcels
Global Returnable Asset Identifier (GRAI)	Returnable assets	Pallet cases, crates, totes
Global Individual Asset Identifier (GIAI)	Assets	Medical or transport equipment

This framework is continuously developed through the GS1 Global Standards Management Process (GSMP) to address emerging trends such as Digital Product Passports (DPP), circular economy initiatives, and new regulatory requirements.[2]

GS1 Digital Link and URI Syntax The GS1 Digital Link standard provides a web-compatible representation of GS1 identifiers (GTINs, batch codes, serial numbers) as URIs.[4] Release 1.6.0, ratified in March 2025, defines the ABNF syntax for representing these keys within a web address.[5] This enables a single 2D barcode to function as a gateway to online information while remaining decodable by offline point-of-sale systems.[5] A typical URI structure follows the pattern `https://domain.com/01/GTIN-VALUE`, where "01" is the Application Identifier (AI) for a GTIN.[6] Central to this architecture is the resolver—intelligent middleware that acts as a router, directing the user to specific resources (e.g., sustainability data or recall status) based on their context.[5]

ISO and Regulatory Context International supply chain identification is further governed by ISO/IEC 15459, which establishes rules for unique identification assigned by authorized Issuing Agencies.[8] ISO/IEC 15459-1 specifically addresses individual transport units, ensuring identities remain unique across global borders.[9] This standardization is essential for the implementation of the EU's Ecodesign for Sustainable Products Regulation (ESPR) and the resulting Digital Product Passport (DPP), which require structured, web-enabled identification to empower green product choices.[10]

2.2**Decentralized Identifier Systems** Decentralized Identifiers (DIDs) represent a new type of globally unique identifier that enables verifiable, decentralized digital identity.[12] The W3C DID v1.0 Recommendation, published on July 19, 2022, specifies a common data model and syntax that decouples identifiers from centralized registries.[12] Table 2 details the core components of the DID architecture and their relevance to product resolution.

Table 2. DID Core Architecture Components.[12]

DID Component	Description	Relevance to Product Resolution
DID Syntax	`did:method:id`	Globally unique product addressing
DID Document	Metadata describing verification methods	Link to decentralized product registries
Service Endpoint	Interaction points for the DID subject	Directs to DPP or traceability data
Verification Method	Cryptographic material for authentication	Ensures product data authenticity

DID Resolution and Provenance DID resolution is the algorithm-driven process of retrieving a DID document for a given identifier.[14] By applying these principles to physical products, manufacturers can create persistent, stable identifiers with updatable service endpoints that are not controlled by a central host. This architecture supports immutable provenance, as the product's identity and its associated lifecycle records are cryptographically bound to a decentralized ledger.[15]

Ethereum Name Service (ENS) The Ethereum Name Service (ENS) serves as a distributed naming system that maps human-readable names (e.g., `brand.eth`) to machine-readable identifiers.[16] ENS utilizes a dot-separated hierarchical architecture where top-level domains like `.eth` are managed by registrar smart contracts.[16] Recent studies indicate ENS's growing popularity, with over 67% of record settings linked to blockchain addresses, despite challenges such as domain squatting.[18] The integration of ENS names as DIDs (via the `did:ens` method) allows for decentralized service discovery and verification of brand-controlled product information.[16]

2.3**Blockchain in Supply Chain** The integration of blockchain technology into supply chain management (SCM) addresses critical needs for traceability, transparency, and data integrity. Systematic reviews of blockchain-based frameworks highlight the use of immutable ledgers and smart contracts to automate trust and reduce administrative fraud.[19] The specific traceability challenges addressed by blockchain across various sectors are summarized in Table 3.

Table 3. Supply Chain Challenges and Blockchain Solutions by Sector.[3]

Supply Chain Sector	Traceability Challenge	Blockchain Solution
Food & Agriculture	Contamination and spoilage recalls	"Farm to table" immutable traceability
Pharmaceuticals	Counterfeit drug infiltration	End-to-end visibility and recall management
Electronics	Infiltration of untrusted hardware	Chip tracking via PUFs and Hyperledger
Logistics	Inefficient manual paperwork	Automated status updates via smart contracts

Anti-Counterfeiting and Authentication Counterfeiting costs the global economy approximately $169 billion annually, particularly impacting critical infrastructure in aerospace and defense.[22] Traditional anti-counterfeiting mechanisms, such as holograms or legacy barcodes, are vulnerable to replication due to their centralized nature.[23] Modern blockchain solutions like "Digi Seal" utilize unique digital fingerprints encoded in QR codes or NFC tags that link directly to immutable on-chain records.[24] These systems enable instant consumer verification through mobile applications, which retrieve data directly from the blockchain to check for duplicate identifiers or unauthorized batch numbers.[25] Advanced systems further utilize AES encryption or protected QR codes to block cloning at the packaging level.[23]

2.4**Multi-Chain Systems** The proliferation of diverse blockchain networks has created a fragmented ecosystem, leading to siloed dApp deployments.[26] Cross-chain interoperability is essential for achieving a unified digital economy where assets and data can move seamlessly.[28]

Cross-Chain Bridges and Protocols Cross-chain bridges allow independent blockchains to communicate and exchange data.[27] Leading protocols are compared in Table 4.

Table 4. Comparative Analysis of Cross-Chain Bridge Protocols.[27]

Bridge Protocol	**Security Model**	**Chain Support**
Wormhole	19-node Guardian network (2/3+ majority)	30+ chains including EVM and Solana
LayerZero	Independent oracles and relayers	Agnostic messaging across various L1s
Cosmos IBC	Light client verification	Native to the Cosmos ecosystem
Polkadot XCMP	Shared security via Relay Chain	Polkadot and Kusama parachains

Universal messaging networks like Wormhole utilize a decentralized network of 19 "Guardians" to produce Verifiable Action Approvals (VAAs), enabling arbitrary data transfer across 30+ chains.[30] LayerZero employs Ultra-Light Nodes and executors to achieve trust-minimized messaging with reduced latency. However, these bridges introduce significant security risks, with over $10 billion in assets currently locked in protocols that have historically been targets for major hacks.

Layer 2 Scaling Solutions Scalability remains a primary bottleneck for global blockchain infrastructure. Layer 2 (L2) solutions, such as Optimistic and Zero-Knowledge (ZK) Rollups, increase transaction throughput by offloading computation from the mainchain (Layer 1).[32] Optimistic rollups assume transaction validity while utilizing a challenge period for fraud proofs, whereas ZK-rollups use cryptographic proofs (zk-SNARKs) to provide immediate finality.[32] These solutions significantly reduce gas fees and congestion, making decentralized resolution economically viable for high-volume consumer interactions.[34]

2.5Public Goods Infrastructure and Sustainability

2.5.1*Theoretical Foundations* Digital Public Goods (DPGs) are open-source software and standards that adhere to privacy laws and help attain Sustainable Development Goals (SDGs).[35] The Digital Public Goods Alliance (DPGA) advocates for a future-oriented model of digital cooperation that prioritizes inclusivity and digital sovereignty.[37] In an interdependent world, DPGs serve as a "rallying cry" for addressing global challenges such as climate action and pandemic preparedness through shared digital infrastructure.[38]

Commons-Based Peer Production (CBPP)

The theoretical foundation for the resolver mesh is Yochai Benkler's theory of Commons-Based Peer Production (CBPP).[39] CBPP describes non-proprietary production models where individuals collaborate to create information materials without market-based hierarchies.[39] Open-source development is a primary example of CBPP, utilizing modular architectures and peer-reviewed codebases to ensure long-term infrastructure resilience.[39] By designing the resolver mesh as a non-proprietary public good, we ensure that the resolution of global product identifiers remains open, neutral, and protected from commercial capture.[39]

2.5.2*Protocol Public Goods Funding Mechanisms* The resolver mesh's long-term sustainability aligns with emerging protocol public goods funding paradigms developed in blockchain ecosystems. Unlike traditional public infrastructure requiring government funding or philanthropic support, protocol public goods utilize mechanism design to align incentives between infrastructure providers and beneficiaries.

Quadratic Funding (QF), formalized by Buterin, Hitzig, and Weyl (2018), demonstrates that properly designed matching mechanisms can enable sustainable funding for non-excludable goods. The QF formula:

$$\Phi_{QF}(c) = \left(\sum_{i=1}^{n} \sqrt{c_i} \right)^2 \tag{1}$$

Figure 1. Quadratic Funding Formula

creates superlinear matching for broadly-supported projects, enabling community-driven resource allocation without centralized discretion.

Complementary mechanisms include Retroactive Public Goods Funding (RetroPGF), pioneered by Optimism Collective, which rewards demonstrated impact rather than predicted value. As of 2025, RetroPGF has allocated $100M+ to public infrastructure, establishing the "impact = profit" principle where contribution to network health directly correlates with reward.

The resolver mesh's architecture as neutral, non-extractive infrastructure positions it as a candidate for protocol-level public goods funding across multiple ecosystems (Ethereum, Filecoin, etc.), creating a diversified sustainability model independent of any single funding source.

3 System Architecture

3.1**Design Principles**

- **P1: Openness:** No authentication, tokens, or fees required for resolution.
- **P2: Neutrality:** Non-discriminatory access regardless of jurisdiction or network.
- **P3: Redundancy:** Multi-chain deployment eliminates single points of failure.
- **P4: Transparency:** All anchoring and identity claims are on-chain and publicly auditable.
- **P5: Immutability:** Contracts are frozen (non-upgradeable).

3.2**Components**

3.2.1*HolocronRouter Contract* The core smart contract, HolocronRouter, serves as an immutable registry of existence and ownership. It provides a ***stock*** function to claim a coordinate derived from the product GTIN hash and a ***checkExistence*** function to verify its existence. It does not store mutable URLs, ensuring the contract remains gas-efficient and maintenance-free. It is deterministically deployed at address **0xeFaAB5Ec699d8c3Bd63d783025268c545357d45F** with identical bytecode across all networks.

3.2.2*Resolution Protocol* The HolocronRouter implements a minimalist interface designed for gas efficiency and broad client compatibility. The resolution protocol follows an Algorithmic Derivation model rather than a Lookup Table model.

(1) Derivation: The Client calculates the deterministic Agent ID by hashing the GS1 GTIN (Hash(GTIN + 'agent_id')).
(2) Verification (L1): The Client queries the HolocronRouter contract (checkExistence) to verify that the derived ID has been officially stocked (registered) by a valid owner.
(3) Discovery (Mesh): The Client connects to the P2P Mesh (Hyperswarm) using the derived Agent ID as the discovery topic.
(4) Retrieval: The Client locates the specific Gateway hosting that Agent's Hyperbee database and retrieves the signed EPCIS event log.

The architecture shifts resolution logic to the Edge (Client), removing the need to update the Smart Contract whenever an endpoint URL changes. This ensures that while the Root of Trust (Existence) is anchored on-chain, the Data Transport occurs entirely peer-to-peer,[45] eliminating gas costs for data retrieval.

3.2.3*Data Integration Layer* The resolver mesh integrates with multiple data sources to provide comprehensive product information while maintaining protocol neutrality. Data integration is handled by Gateway Nodes. A Gateway ingests data from external sources such as Open Food Facts,[47] CSVs, Legacy ERPs, or Brand Databases and converts them into Agent Digital Twins stored in immutable append-only logs.[46]

When a Client queries the mesh, the Gateway acts as a Proxy Oracle by retrieving the raw data and wrapping it in a compliant EPCIS 2.0[49] Document, and cryptographically signing the response via its wallet. This ensures that even 'dead' data from legacy web databases becomes live, verifiable, and signed when entering the mesh.

Primary Data Source: Open Food Facts

Open Food Facts (OFF)[47] serves as the primary data source for consumer packaged goods (CPG) in the food and beverage sector. OFF is a collaborative, open database containing over 3 million products worldwide, maintained by a global community of contributors. Product records in OFF include nutritional information, ingredient lists, allergen warnings, packaging materials, and sustainability scores (Nutri-Score, Eco-Score).

Integration with OFF operates through a standardized endpoint pattern:

```
https://world.openfoodfacts.org/api/v2/product/{GTIN}
```

When the resolver returns this endpoint for a given GTIN, client applications can retrieve the complete product record in JSON format. The API supports content negotiation for localized data based on the Accept-Language header, enabling multilingual consumer experiences.

Distributed Append-Only Storage

Unlike static content-addressing systems (such as IPFS) which require changing the reference identifier whenever data is updated, the resolver mesh utilizes Hypercore, a distributed append-only log protocol.

This architecture offers specific advantages for supply chain digital twins:

- Mutable Tips with Immutable History: The Agent ID (Public Key) remains constant, acting as a stable reference for the QR code. However, new EPCIS events can be appended to the log, allowing the product's history to grow over time without breaking the digital link.
- Sparse Replication: Client devices (such as mobile phones) can query and verify specific records within a massive database without downloading the entire dataset, a critical feature for enabling lightweight resolution on edge devices.
- Peer-to-Peer Swarming: Data availability is maintained via a Distributed Hash Table (DHT), allowing popular products to be cached by multiple nodes, increasing resilience and reducing latency compared to centralized gateways.

Federated Discovery Strategy

Since resolution is algorithmic rather than static, redundancy is achieved through the Federated Discovery mechanism in the client SDK. When a product is scanned, the client does not query a single endpoint. Instead, it:

(1) Derives the Topic: Hash(GTIN)
(2) Queries the Mesh: Connecting to multiple peers simultaneously via the DHT.
(3) Aggregates Responses: Collecting signed Digital Twins from multiple Gateways (e.g., Brand Gateway, Open Food Facts Gateway, Retailer Gateway).

This eliminates the need for an on-chain array of URLs. The “Fallback” is inherent to the mesh: if one Gateway is offline, the client automatically retrieves the Digital Twin from the next available peer serving that topic.

Data Source Registration

Data sources (brands, NGOs, or aggregators) register their data by Seeding a Hypercore Database.

(1) Ingestion: The source converts their product data (CSV/JSON) into the standardized Digital Twin AsciiDoc format.
(2) Indexing: The seeder builds local search indexes (GTIN, Brand, Keyword, etc) and Vector Symbolic Architecture (VSA) embeddings for semantic search.
(3) Announcement: The Gateway announces the existence of this new registry to the Backbone via the P2P discovery channel (cis_announce).

This process is permissionless. Anyone can seed a registry. Trust is established not by gatekeeping the registry, but by verifying the Cryptographic Signatures of the data provider against the on-chain Holocron anchor.

Data Quality and Trust Model

The resolver mesh makes no guarantees about the accuracy, completeness, or freshness of data returned from integrated sources. This design choice reflects the protocol's role as neutral infrastructure rather than a data validation authority. Responsibility for data quality lies with the respective data sources. The resolver mesh separates Transport from Truth. The protocol ensures data is delivered intact, but the veracity of the data is derived from the signer.

- Proxy Attestation: When data comes from a third-party aggregator (e.g., a Gateway mirroring Open Food Facts), the Gateway signs the response as an Oracle, explicitly marking the data as a “Digital Twin” rather than a direct brand claim.
- Direct Brand Claim: When a brand runs their own node, they sign with their corporate key.
- Client Verification: The Client application checks the signature against the Holocron Registry. If the signer is not the registered owner of the GTIN, the UI displays a “Proxy/Unverified” warning, allowing users to distinguish between official brand data and community-sourced information.

Client applications are encouraged to implement source-aware trust models, potentially weighting or filtering results based on data provenance. Future work may explore integration with cryptographic attestation systems (e.g., verifiable credentials) to enable data source authentication without compromising protocol neutrality.

Integration with GS1 Digital Link Resolvers

While the decentralized resolver mesh operates independently, it is designed for interoperability with existing GS1 Digital Link resolver infrastructure. Clients scanning a 2D barcode can query both centralized GS1 resolvers and the decentralized mesh in parallel, comparing results or falling back to the mesh when centralized services are unavailable.

This hybrid approach enables progressive decentralization: brand owners currently using centralized resolvers can add their products to the mesh as a redundancy layer, gradually shifting traffic as confidence in the decentralized infrastructure grows. The protocol imposes no restrictions on data source participation, allowing commercial, non-profit, and community-driven databases to coexist within the resolver's endpoint registry.

The system is also designed for backward compatibility with existing GS1 Digital Link infrastructure. A HTTP Bridge component allows standard Web2 clients (such as legacy barcode scanners or mobile browsers) to query the decentralized mesh via standard HTTPS requests.

When a legacy client requests `https://resolver.domain/01/{GTIN}`, the Bridge translates the request into a P2P mesh lookup, retrieves the signed Digital Twin, and returns it as a standard GS1 Linkset (JSON-LD). This allows the decentralized mesh to act as a "Drop-in Replacement" or "Redundancy Layer" for existing centralized resolvers like Scanbuy or Kezzler without requiring changes to the physical packaging.

3.3**Multi-Chain Deployment Strategy** The resolver mesh achieves high availability by deploying the same HolocronRouter across multiple independent blockchain networks. Table 5 compares the deployment approaches considered for this infrastructure.

TABLE 5. Multi-Chain Service Deployment Strategies.[40]

Deployment Strategy	Mechanism	Benefit for Resolver Mesh
IntegrateX Protocol	Logic-State Decoupling (LSD)	Efficient multi-chain logic execution
SCaaS Paradigm	Modular "LEGO-brick" assembly	Reusability of mature, secure modules
Deterministic Deployment	CREATE2 bytecode matching	Identical contract addresses across networks

Deterministic Deployment Mechanism

To maintain a unified interface for client applications, we utilize the CREATE2 opcode (EIP-1014) for deterministic contract deployment. Unlike the standard CREATE opcode, which calculates addresses based on the deployer's nonce, CREATE2 uses a formula:

$$\text{address} = \text{keccak256}(\text{0xff} \mathbin{\|} \text{deployer} \mathbin{\|} \text{salt} \mathbin{\|} \text{keccak256}(\text{init\_code})) \tag{2}$$

FIGURE 2. CREATE2 Address Generation Formula

where `0xff` is a constant prefix, `deployer` is the address deploying the contract, `salt` is a chosen 32-byte value, and `init_code` is the contract creation bytecode.

This approach enables "counterfactual" deployments, where a contract address is pre-determined and can be interacted with before the actual deployment transaction occurs.

By pinning the Solidity compiler version and keeping the bytecode identical (ensuring a matching `init_code` hash), the HolocronRouter is accessible at the fixed address:

0xeFaAB5Ec699d8c3Bd63d783025268c545357d45F

across all supported EVM networks. This eliminates the need for complex multi-chain address registries and simplifies client integration.

Network Selection and Viability Rationale

The selection of blockchains for the resolver mesh is based on four primary viability criteria:

- **EVM Compatibility:** Essential for utilizing the CREATE2 opcode and ensuring the deterministic execution of the HolocronRouter logic across all target environments.[40]
- **Geographic Distribution:** We prioritize networks with global node infrastructure to prevent "geospatial centralization". By distributing the resolver across chains with diverse physical footprints, we enhance resilience against regional connectivity issues or jurisdictional data silos.

- **Community and Governance Viability:** We assess the active developer ecosystem and long-term governance stability of each network to ensure the infrastructure remains persistent as a public good.
- **Scalability and Performance:** Networks must support high-volume "view" queries without incurring transaction latency that would degrade the consumer scanning experience.

4 Security Analysis

4.1**Trust Model** The security of the decentralized resolver mesh is based on a defined boundary between trusted and untrusted system components.

Trusted Components:

- **Cryptographic Attestation:** The system relies on Ed25519 signatures generated by Gateway nodes via a wallet. The Resolver verifies that the signer of the data matches the owner anchored on the Holocron contract.
- **Blockchain Network Consensus:** The resolver inherits the security and liveness guarantees of each underlying blockchain network.[44]
- **Execution Environment:** We rely on the correctness of the EVM implementation for deterministic logic execution.[44]
- **Data Source Integrity:** The system trusts the integrity of curated open databases, such as Open Food Facts, to provide accurate base product metadata.[47]

Untrusted Components:

- **Gateway Data Transmission:** While the Gateway is a transport node, it cannot tamper with the data without breaking the signature hash, which is cross-referenced against the on-chain registry.
- **Individual RPC Endpoints:** Single RPC providers are not inherently trusted and are mitigated through multi-endpoint client failover strategies.[45]
- **Client-Side Resolution Logic:** While clients execute the final resolution path, this logic is open-source and subject to public verification.[44]

Threats Outside Scope:

- **Data Source Compromise:** The resolver acts as a routing mesh and does not independently authenticate the veracity of source data.[44]
- **Network-Level Partitioning:** Global Internet routing failures or massive, coordinated network partitions are considered external systemic risks.[44]

4.2**Attack Vectors and Mitigations** The multi-chain architecture is designed to proactively address several classes of security threats common to resolution infrastructure. Our threat model and associated mitigations are analyzed in Table 6.

Table 6. Threat Model Mitigation Analysis.[44]

Security Threat	**Mitigation Strategy**	**Residual Risk**
Contract Compromise	Frozen bytecode (non-upgradeable)	Initial deployment audit gaps
Network Failure	Multi-chain redundant paths	Coordinated global network outage
Data Source Outage	Integration with multiple sources (IPFS/Arweave)	Multi-source synchronized failure
RPC Censorship	Support for diverse, decentralized RPCs	Centralized provider bottlenecks

4.3**Privacy Considerations** Privacy is maintained through a combination of view-first queries and peer-to-peer routing.

Unlike centralized resolvers where a single server sees every scan, the Reap Protocol utilizes a Distributed Hash Table (DHT)[45] for discovery. A client looking up a product connects to a dynamic swarm of peers. Because the connection is encrypted and peer-to-peer, no central server sees the specific product data being transferred.

Metadata leakage is minimized by splitting the query:

(1) The RPC Provider (Blockchain) only sees a `checkExistence` call for a hashed coordinate. It does not see the product data or the user's location.
(2) The Mesh Peer (Gateway) sees the data request but does not necessarily know the user's identity, as connections are ephemeral.

Crucially, the protocol is built on a no-user-tracking principle, with no persistent user identifiers or analytics cookies built into the core resolution logic.

4.4**Economic Security** The system is intentionally designed without a native token or staking mechanism. This choice eliminates entire categories of attack vectors related to tokenomics manipulation, price volatility, or economic incentive gaming.

Economic sustainability is achieved through a public goods funding model:

- **Cost Amortization:** The costs of initial deterministic deployment are one-time and amortized over the infrastructure's lifetime.
- **Zero-Fee Resolution:** Because data retrieval occurs over the peer-to-peer mesh[45], resolution queries incur zero gas costs for end-users, ensuring that access to critical product safety and sustainability information is never gated by financial barriers.
- **Neutrality:** By removing commercial relationships and fee structures, the mesh remains a neutral piece of digital infrastructure serving the broad public interest rather than narrow institutional goals.

5 Implementation and Deployment

The reference implementation is built using Hyperswarm for networking and Hyperbee for decentralized database storage. The client SDK (ReapClient) handles the cryptographic derivation of GS1 keys into mesh topics. The Gateway utilizes AsciiDoc[45] as the canonical format for Digital Twin metadata, ensuring human-readability while maintaining machine-parseable structure for EPCIS compliance.

5.1**Core Smart Contract Interface** The on-chain anchor is implemented as a minimalist Solidity contract designed for consistent CREATE2 deployment addresses across all EVM chains.

Listing 1. Core Holocron Router Interface (Solidity)

```
interface IHolocron {
    struct HolocronEntry {
        bool exists;
        address provider;
        uint40 timestamp;
    }

    event HolocronStocked(uint256 indexed coordinate, address indexed provider);

    function stock(uint256 coordinate) external;

    function checkExistence(uint256 coordinate) external view returns (bool);

    function registry(uint256 coordinate) external view returns (
        bool exists,
        address provider,
        uint40 timestamp
    );
}
```

5.2**Deterministic Identifier Derivation** To ensure the resolver mesh remains stateless and algorithmic, client software derives network topics directly from GS1 standard identifiers using a standard hashing scheme (SHA-256). This eliminates the need for a centralized lookup table.

Listing 2. Deterministic Identity Derivation (TypeScript)

```
import { createHash } from ’crypto’;

function deriveIdentity(gtin: string) {
    const agentId = createHash(’sha256’)
        .update(‘${gtin}:agent_id‘)
        .digest(’hex’);

    const coordinate = BigInt("0x" + agentId).toString();
    return { agentId, coordinate };
}
```

6 Discussion

6.1**Limitations** While the decentralized resolver mesh addresses critical vulnerabilities in centralized resolution infrastructure, the system exhibits several inherent limitations that warrant careful consideration.

L1: Data Source Dependencies

The resolver mesh operates as routing infrastructure rather than a data validation authority. Consequently, the system inherits reliability and accuracy characteristics from its underlying data sources. If Open Food Facts experiences data quality issues—such as outdated product information, incomplete records, or community-contributed errors—these deficiencies propagate to applications relying on the resolver. While the multi-source endpoint strategy provides redundancy, it does not guarantee data correctness.

This limitation is fundamental to the resolver's design philosophy of neutrality. Any attempt to implement protocol-level data validation would require the resolver to make subjective trust decisions, potentially favoring certain data sources over others and undermining the system's commitment to non-discrimination. We accept this trade-off explicitly: the resolver ensures *access* to information but delegates *assessment* of that information to client applications and end users.

A partial mitigation exists through cryptographic content addressing and potential integration with verifiable credential systems, but these approaches do not solve the underlying challenge of data source trustworthiness. Applications requiring high-assurance data must implement source-aware validation logic or multi-source consensus mechanisms at the client layer.

L2: Network Heterogeneity

The multi-chain deployment strategy introduces operational complexity arising from heterogeneous blockchain network characteristics. Different networks exhibit significant variation in:

- **Anchoring Latency vs. Resolution Speed:** While resolution (reading) occurs instantly over the peer-to-peer mesh, anchoring (writing new identities) is subject to the block times and finality of the underlying L1/L2 network. This decoupling ensures that consumer scanning experiences remain sub-second even during blockchain congestion events.
- **RPC endpoint reliability**: Public RPC infrastructure varies widely in uptime, rate limiting, and geographic distribution.
- **Economic costs**: Although resolution queries use view calls (zero gas cost), the initial contract deployment incurs different costs across networks based on their fee markets.
- **Governance models**: Networks have diverse governance structures that may impact long-term protocol stability.

These variations mean that user experience may differ depending on which network a client queries. A client connected to a congested or poorly-maintained network may experience slower resolution times or higher failure rates compared to clients using well-provisioned networks. While the multi-chain architecture provides failover capability, it does not eliminate performance variability entirely.

Furthermore, maintaining awareness of network status requires clients to implement sophisticated monitoring and selection logic. Simple client implementations that always query the same network may not fully benefit from the mesh's redundancy guarantees. This tension between simplicity and resilience represents an ongoing design challenge.

L3: Adoption Barriers

Despite the resolver's design as open, token-free infrastructure, practical adoption faces several barriers:

- **Blockchain infrastructure requirements**: Client applications must integrate Web3 libraries or RPC endpoint connectivity, which adds technical complexity compared to traditional REST API integration. For resource-constrained IoT devices or legacy systems, this requirement may be prohibitive.
- **Developer familiarity**: Blockchain-based systems remain unfamiliar to many supply chain and e-commerce developers. The learning curve for understanding contract interfaces, RPC providers, and multi-chain failover strategies may slow adoption compared to conventional centralized APIs.
- **Perception challenges**: Association with cryptocurrency volatility and blockchain hype cycles may create institutional skepticism, even though the resolver operates without tokens or speculative elements.
- **Network effect dynamics**: The resolver's value increases with the number of registered products and integrated data sources. During early deployment, the product catalog may be sparse compared to mature centralized databases, creating a "cold start" problem that could hinder adoption.

These barriers suggest that widespread adoption will require not only technical excellence but also ecosystem-building efforts, developer education, and demonstration of operational reliability over extended periods. The resolver may initially serve as a redundancy layer for existing systems rather than a primary resolution mechanism.

To mitigate these barriers, the architecture includes a HTTP Bridge functionality. This allows legacy Web2 clients (standard web browsers and mobile cameras) to query the mesh via standard HTTPS requests (`https://resolver.domain/01/{GTIN}`), with the Bridge node handling the P2P complexity and returning standard JSON-LD Linksets.

6.2**Future Work** The current implementation establishes a foundation for decentralized product resolution, but several research and development directions could significantly enhance the system's capabilities and robustness.

F1: Enhanced Data Verification

Future iterations of the resolver could integrate cryptographic attestation mechanisms to enable verifiable data provenance without compromising protocol neutrality. Potential approaches include:

- **Verifiable Credentials (VCs)**: Data sources could issue W3C Verifiable Credentials signed by the data provider's cryptographic key. The resolver would return both the endpoint and an associated credential, allowing clients to verify the data's origin and integrity without trusting the resolver itself.
- **Merkle proofs for data integrity**: For large product databases, root hashes could be stored on-chain with client-verifiable Merkle proofs demonstrating that retrieved data belongs to the attested dataset. This approach would enable detection of data tampering without requiring on-chain storage of full product records.
- **Multi-party computation for consensus**: When multiple data sources provide conflicting information for the same product, cryptographic protocols could enable privacy-preserving consensus computation, identifying majority opinions or weighted trust scores without revealing individual source data.

These enhancements would shift the resolver from a purely routing function toward a verifiable information infrastructure while maintaining the core principle of not imposing a single "truth" authority.

F2: Cross-Chain Communication

The current architecture deploys isolated contract instances across networks with no inter-chain communication. Future work could explore coordination mechanisms to maintain consistency or enable federated governance:

- **Cross-chain state synchronization**: Bridge protocols like Wormhole or LayerZero could enable the propagation of endpoint registry updates across all deployed networks, ensuring that product additions on one chain become visible mesh-wide within bounded time.
- **Decentralized governance of registry updates**: A cross-chain governance protocol could allow community voting on disputed product entries or data source additions, with results committed to all networks atomically.
- **Query result aggregation**: Advanced client protocols could query multiple networks in parallel and use cross-chain verification to detect inconsistencies, potentially indicating network-specific attacks or data corruption.

However, introducing cross-chain dependencies would increase system complexity and potentially reintroduce centralization risks through bridge governance. Any such enhancements must carefully balance functionality gains against the mesh's current property of complete network independence.

F3: Performance Optimization

Several optimizations could improve resolution speed and resource efficiency:

- Gateway Indexing: While the mesh is distributed, Gateways maintain local Hyperbee B-tree indexes (GTIN, Brand, Keyword). This allows $O(\log n)$ lookup times even as the registry scales to billions of items, preventing the need to scan linear logs.
- Sparse Replication: Leveraging the underlying Hypercore protocol, client devices can request and verify specific blocks of the Digital Twin metadata (e.g., just the sustainability score) without downloading the full history or binary attachments. This is critical for bandwidth-constrained IoT devices.
- Tiered caching strategies: Client libraries can implement intelligent caching based on product type, with longer cache TTLs for stable grocery items (where data rarely changes) versus shorter TTLs for high-velocity supply chains.

These optimizations would be particularly valuable as the resolver scales to support global product catalogs potentially numbering in the billions of unique identifiers.

6.3**Broader Implications** The successful deployment of a decentralized resolver mesh for product identifiers demonstrates principles applicable to a wider class of digital infrastructure needs. This work contributes to an emerging understanding of how blockchain technology can serve public goods functions beyond financial applications.

6.3.1*Applicability to Other Identifier Systems* The architectural patterns developed for GS1 product resolution translate directly to other critical identifier systems:

- **Open data registries**: Scientific datasets, government records, and cultural heritage collections currently rely on centralized DOI (Digital Object Identifier) systems or institutional repositories. A decentralized registry mesh could provide redundant access paths resilient to institutional failures or funding disruptions.
- **Public certificate authorities**: TLS certificate transparency logs and code signing infrastructure represent critical trust infrastructure currently operated by centralized entities. Multi-chain deployment of certificate verification logic could reduce dependence on any single authority.
- **Decentralized naming systems**: Building on the ENS integration discussed in Section 2.2, the resolver's principles could extend to general-purpose naming infrastructure, providing alternatives to DNS with enhanced censorship resistance.
- **Scientific identifier resolution**: ORCID for researchers, ROR for institutions, and RID for research resources could benefit from decentralized resolution infrastructure, particularly for cross-border research collaboration where access to centralized systems may be restricted.

6.3.2*Public Goods Infrastructure Design Patterns* This work establishes several reusable patterns for blockchain-based public infrastructure:

(1) **Deterministic multi-chain deployment**: The CREATE2-based approach to identical contract addresses enables transparent, verifiable service availability without cross-chain communication overhead.
(2) **Token-free operation**: Demonstrating that blockchain infrastructure can provide public utility without introducing economic game theory complexity or speculative elements.
(3) **Progressive decentralization**: Enabling coexistence with centralized systems during transitional periods, rather than requiring immediate wholesale replacement.
(4) **Separation of routing and storage**: The resolver's architecture cleanly separates identifier resolution (on-chain) from data storage (off-chain/decentralized), providing a template for other information systems.
(5) **Neutrality as a core principle**: Explicitly designing infrastructure to avoid discrimination or preferential treatment, even when such neutrality introduces limitations.

These patterns may inform future research on blockchain-based public goods, particularly in contexts where openness, resilience, and long-term sustainability are prioritized over performance optimization or feature richness.

6.3.3*Implications for Supply Chain Transparency* Beyond the technical contributions, this work engages with broader societal questions about information access and supply chain accountability. As regulatory mandates like the EU's Digital Product Passport create requirements for product traceability, the choice between centralized and decentralized resolution infrastructure has profound implications:

- **Regulatory compliance without vendor lock-in**: Decentralized infrastructure enables compliance with transparency mandates while preventing single vendors from capturing regulatory infrastructure.
- **Global access to safety information**: During product recalls or safety incidents, centralized systems may be unavailable due to traffic surges, maintenance windows, or geographic restrictions. Decentralized resolution ensures that critical safety information remains accessible even during crisis conditions.
- **Empowering consumer choice**: By ensuring open access to product information, the resolver supports informed consumer decision-making regarding sustainability, ethics, and health—aligning with the digital public goods principle of empowering individuals.

The resolver mesh represents one component of a broader ecosystem needed to realize truly transparent supply chains. Its success depends not only on technical robustness but also on adoption by brands, retailers, consumer applications, and regulatory frameworks.

6.3.4*Economic Sustainability and Public Goods Funding* The resolver mesh's viability as permanent public infrastructure depends on sustainable funding mechanisms that do not compromise its neutrality or openness. Traditional models—venture capital, advertising, data monetization—create incentive misalignments that would undermine the system's public goods characteristics.

Protocol public goods funding mechanisms provide an alternative sustainability path. The resolver's characteristics align with criteria for protocol-level funding:

- Non-excludability: Resolution is free and open to all.
- Non-rivalry: One user's query does not diminish availability.
- Network effects: Value increases with broader adoption.
- Neutrality: No single party controls or profits from the infrastructure.

Potential funding pathways include:

- **Gitcoin Grants:** Quadratic funding rounds have distributed $67M to Ethereum ecosystem public goods. The resolver's multi-chain deployment and regulatory compliance value proposition position it as a candidate for recurring grants.
- **Optimism RetroPGF:** Demonstrated infrastructure improvements qualify for retroactive rewards. The resolver's integration with supply chain transparency and food safety systems creates measurable impact suitable for RetroPGF evaluation.
- **Protocol treasury allocation:** Networks like Filecoin have established dedicated public goods funding ($3.68M+ in 2025), with resolver infrastructure potentially qualifying as foundational to network utility.
- **Public Goods Network (PGN):** As a Layer 2 channeling sequencer fees directly to public goods, PGN demonstrates protocol-level integration of funding mechanisms. The resolver could deploy on PGN, automatically receiving fee-share support.

Critical to these mechanisms is that they preserve the resolver's neutrality—funding is allocated based on demonstrated public value rather than commercial extraction. This contrasts with platform business models where monetization incentives drive feature decisions and access restrictions.

The emergence of composable funding infrastructure (Allo Protocol, Grant Ships, Flow State) enables the resolver to participate in multiple funding mechanisms simultaneously, creating resilience through diversification. This multi-source model mirrors the resolver's multi-chain deployment strategy: no single point of failure in either technical infrastructure or economic sustainability.

Long-term viability depends not only on initial deployment but on sustained maintenance, security audits, and adaptation to evolving standards. Protocol public goods funding represents the first mechanism design framework capable of sustaining neutral infrastructure indefinitely without requiring philanthropic dependence or commercial compromise.

7 Conclusion

We have presented a decentralized, multi-chain resolver mesh for GS1 product identifiers that addresses critical limitations in current centralized resolution infrastructure. By combining deterministic smart contract anchors with a distributed peer-to-peer transport layer, the system provides redundant, open access to product information without requiring tokens, fees, or authentication.

The resolver mesh demonstrates that blockchain technology can serve public goods functions, providing critical infrastructure with explicit commitments to openness, neutrality, and long-term sustainability. The system integrates with community-maintained databases such as Open Food Facts and supports distributed append-only storage via Hypercore, enabling diverse trust models while maintaining protocol neutrality.

Our security analysis establishes the trust model and evaluates resilience to various failure modes, showing that multi-chain redundancy effectively eliminates single points of failure inherent in centralized systems. The economic model, based on one-time deployment costs and zero-fee resolution, ensures that access to product information remains permanently open regardless of changing commercial or institutional priorities.

The resolver mesh is operational across 10 blockchain networks, serving as both functional infrastructure and proof-of-concept for broader applications of decentralized public goods. The architectural patterns developed—deterministic multi-chain deployment, token-free operation, and separation of routing from storage—provide reusable templates for other critical identifier systems including scientific data registries, public certificate authorities, and decentralized naming services.

While limitations exist, particularly regarding data source dependencies and adoption barriers, the system establishes a foundation for progressive decentralization of supply chain information infrastructure. Future work on cryptographic data verification, cross-chain governance, and performance optimization can build upon this foundation while preserving the core principles of openness and neutrality.

As regulatory frameworks increasingly mandate product transparency and traceability, the choice of resolution infrastructure becomes a choice about the future structure of global information systems. This work demonstrates that decentralized alternatives to centralized services are not only technically feasible but operationally viable, offering a path toward digital infrastructure that serves broad public interest rather than narrow institutional control.

Availability: Specification and code at `https://github.com/apriloracle/unified-gs1-resolver-mesh`

References

[1] GS1, "GS1 Barcode Standards - Bar Code Graphics," [Online]. Available: `https://www.barcode.graphics/tools-gs1-general-specs/`

[2] GS1, "GS1 General Specifications," Jan. 2025. [Online]. Available: `https://www.gs1uk.org/insights/news/GS1-General-Specifications-updated-for-2025`.

[3] GS1 US, "GS1 Standards." [Online]. Available: `https://www.gs1us.org/industries-and-insights/standards`.

[4] GS1, "GS1 Digital Link URI Syntax," Release 1.6.0. [Online]. Available: `https://developer.digitalhealth.gov.au/standards/gs1-digital-link-uri-syntax`.

[5] GS1, "GS1 Digital Link." [Online]. Available: `https://www.gs1.org/standards/gs1-digital-link`.

[6] GS1, "GS1 Digital Link URI Syntax - Fact Sheet." [Online]. Available: `https://assets.ctfassets.net/9uypwcnuzbqi/5ngHfDU3Edexh1ofdnSqvt/266e75d721697fab001e918b6dcdcac3/GS1au-fact-sheet-dl-syntax.pdf`.

[7] Bar Code Graphics, "Common Mistakes When Implementing GS1 Digital Link QR Codes." [Online]. Available: `https://www.barcode.graphics/common-mistakes-companies-make-when-implementing-gs1-digital-link-qr-codes/`.

[8] AIM Global, "Registration Authority | ISO/IEC 15459." [Online]. Available: `https://www.aimglobal.org/registration-authority-iso-iec-15459/`.

[9] ISO/IEC, "ISO/IEC 15459-1:2014 Information technology — Automatic identification and data capture techniques — Unique identification — Part 1: Individual transport units," 2014. [Online]. Available: `https://cdn.standards.iteh.ai/samples/54779/e0e908fa163e43e0a4d9a6ec663bb498/ISO-IEC-15459-1-2014.pdf`.

[10] GS1 in Europe, "Web enabled, structured path identification," 2024. [Online]. Available: `https://gs1.eu/wp-content/uploads/2024/07/GS1-Enabling-DPP-White-Paper.pdf`.

[11] GS1, "GS1 Standards enabling the EU digital product passport," 2025. [Online]. Available: `https://gs1.eu/wp-content/uploads/2025/04/GS1-Standards-Enabling-DPP-V2.2.pdf`.

[12] M. Sporny et al., "Decentralized Identifiers (DIDs) v1.0," W3C Recommendation, Jul. 2022. [Online]. Available: `https://www.w3.org/TR/did-1.0/`.

[13] W3C, "Decentralized Identifiers (DIDs) v1.0," 2022. [Online]. Available: `https://www.w3.org/news/2022/decentralized-identifiers-dids-v1-0-is-a-w3c-recommendation/`.

[14] W3C Credentials Community Group, "Decentralized Identifier Resolution (DID Resolution) v0.3." [Online]. Available: `https://w3c.github.io/did-resolution/`.

[15] P. K. S. et al., "Fake Product Detection through QR Code using Blockchain," *Int. J. of Adv. Research in Science, Comm. and Tech.*, Jan. 2026. [Online]. Available: `https://ijarsct.co.in/Paper25075.pdf`.

[16] Veramo Labs, "did:ens Method Specification." [Online]. Available: `https://github.com/veramolabs/did-ens-spec`.

[17] ENS, “What is the Ethereum Name Service?” [Online]. Available: https://docs.ens.domains/learn/protocol/.

[18] N. Johnson et al., “Ethereum Name Service: the Good, the Bad, and the Ugly,” *arXiv preprint*, 2021. [Online]. Available: https://arxiv.org/pdf/2104.05185.

[19] “Blockchain-Driven Food Supply Chains: A Systematic Review for Unexplored Opportunities,” *Applied Sciences*, vol. 14, no. 19, 2026. [Online]. Available: https://www.mdpi.com/2076-3417/14/19/8944.

[20] “Blockchain-Enabled Supply Chain Management: A Review of Security, Traceability, and Data Integrity,” *MDPI*, 2026. [Online]. Available: https://www.mdpi.com/2076-3417/15/9/5168.

[21] “Blockchain implementation in pharmaceutical supply chains: A review and conceptual framework,” *Taylor & Francis Online*, 2026. [Online]. Available: https://www.tandfonline.com/doi/full/10.1080/00207543.2022.2125595.

[22] “A Blockchain-Based Framework for Supply Chain Provenance,” *IEEE Xplore*, 2026. [Online]. Available: https://ieeexplore.ieee.org/ielaam/6287639/8600701/8884089-aam.pdf.

[23] “Blockchain-Based Counterfeit Product Detection,” *IARJSET*, 2026. [Online]. Available: https://iarjset.com/wp-content/uploads/2025/08/IARJSET.2025.12745.pdf.

[24] “Identification of fraudulent products using blockchain technology,” *GJETA*, 2026. [Online]. Available: https://gjeta.com/sites/default/files/GJETA-2025-0222.pdf.

[25] “FAKE PRODUCT IDENTIFICATION BY QR CODE USING BLOCKCHAIN,” *IRJMETS*, 2026. [Online]. Available: https://www.irjmets.com/upload_newfiles/irjmets70800070871/paper_file/irjmets70800070871.pdf.

[26] Chainlink, “Cross-chain vs Multi-chain.” [Online]. Available: https://chain.link/education-hub/cross-chain-vs-multi-chain.

[27] CryptoEQ, “Blockchain Interoperability: Challenges, Solutions, and the Future of a Connected Multi-Chain Ecosystem,” 2026. [Online]. Available: https://www.cryptoeq.io/articles/blockchain-interoperability-solutions.

[28] “Cross-Chain Interoperability in DeFi: Architectures for Seamless and Secure Multi-Blockchain Transactions,” *ResearchGate*, 2026. [Online]. Available: https://www.researchgate.net/publication/392967924_Cross-Chain_Interoperability_in_DeFi_Architectures_for_Seamless_and_Secure_Multi-Blockchain_Transactions.

[29] Syndika, “Research: Cross-chain interoperability,” 2026. [Online]. Available: https://syndika.co/blog/research-cross-chain-interoperability/.

[30] BizThon, “Wormhole — Cross-chain Bridge Protocol for Assets and Messages,” *Medium*, 2026. [Online]. Available: https://medium.com/@BizthonOfficial/wormhole-cross-chain-bridge-protocol-for-assets-and-messages-c0900e0d50f3.

[31] Wormhole, “Architecture.” [Online]. Available: https://wormhole.com/docs/protocol/architecture/.

[32] “Layer 2 Scaling Solutions for Blockchain Networks: An Analysis of Rollups, State Channels, Sidechains and Sharding,” *IJIRT*, 2026. [Online]. Available: https://ijirt.org/publishedpaper/IJIRT188921_PAPER.pdf.

[33] “Blockchain Layer 2 Rollups, Optimistic vs zero knowledge,” *ResearchGate*, 2026. [Online]. Available: https://www.researchgate.net/publication/382319087_Blockchain_Layer_2_Rollups_Optimistic_vs_zero_knowledge.

[34] Gemini, “Layer-2 Scaling: zk-Rollups and Optimistic Rollups,” 2026. [Online]. Available: https://www.gemini.com/cryptopedia/layer-2-scaling-zk-rollup-optimistic-rollup.

[35] UNICEF, “2024 State of the Digital Public Goods Ecosystem,” 2026. [Online]. Available: https://www.unicef.org/digitalimpact/media/826/file/DPG-Ecosystem-2024.pdf.pdf.

[36] BMZ, “2024 State of the Digital Public Goods Ecosystem,” 2026. [Online]. Available: https://www.bmz-digital.global/wp-content/uploads/2025/02/DPG-Ecosystem-2024.pdf.

[37] Digital Public Goods Alliance, “DPG Ecosystem 2023,” 2026. [Online]. Available: https://www.digitalpublicgoods.net/DPG-Ecosystem-2023.pdf.

[38] UNDP, “Human Development Report 2023/2024,” 2026. [Online]. Available: https://www.undp.org/sites/g/files/zskgke326/files/2024-03/human_development_report_2023-24.pdf.

[39] Y. Benkler, “Coase’s Penguin, or, Linux and The Nature of the Firm,” *The Yale Law Journal*, vol. 112, no. 3, 2002. [Online]. Available: https://www.researchgate.net/publication/1955545_Coase’s_Penguin_Or_Linux_and_The_Nature_of_the_Firm.

[40] “Atomic Smart Contract Interoperability with High Efficiency via Cross-Chain Integrated Execution,” *arXiv*, 2026. [Online]. Available: https://arxiv.org/html/2502.12820v2.

[41] “Atomic Smart Contract Interoperability With High Efficiency via Cross-Chain Integrated Execution,” *IEEE Computer Society*, 2026. [Online]. Available: https://www.computer.org/csdl/journal/td/2025/12/11180137/2aiMtCYDEuk.

[42] University of Washington, “A Multidimensional Contract Design for Smart Contract-as-a-Service,” 2026. [Online]. Available: https://faculty.washington.edu/weicaics/paper/papers/JinghanSKFESC2025.pdf.

[43] “A Multidimensional Contract Design for Smart Contract-as-a-Service,” *ResearchGate*, 2026. [Online]. Available: https://www.researchgate.net/publication/388050717_A_Multidimensional_Contract_Design_for_Smart_Contract-as-a-Service.

[44] G. Wood, “Ethereum: A Secure Decentralized Generalised Transaction Ledger,” *Ethereum Project Yellow Paper*, vol. 151, pp. 1-32, 2014. [Online]. Available: https://ethereum.github.io/yellowpaper/paper.pdf.

[45] P. Maymounkov and D. Mazieres, “Kademlia: A Peer-to-Peer Information System Based on the XOR Metric,” in *International Workshop on Peer-to-Peer Systems*, 2002. [Online]. Available: https://pdos.csail.mit.edu/~petar/papers/maymounkov-kademlia-lncs.pdf.

[46] The Holepunch Team, “Hypercore Protocol: A Distributed Append-Only Log.” [Online]. Available: https://hypercore-protocol.org/.

[47] S. Gigandet et al., “Open Food Facts: A collaborative, free and open database of food products,” 2012. [Online]. Available: https://world.openfoodfacts.org/.

[48] The Eclipse Foundation, “AsciiDoc Language Specification.” [Online]. Available: https://asciidoc.org/.

[49] GS1, “EPCIS 2.0 JSON/JSON-LD Binding,” 2022. [Online]. Available: https://ref.gs1.org/standards/epcis/2.0.0/epcis-context.jsonld.

[50] V. Buterin, Z. Hitzig, and E. G. Weyl, “Liberal Radicalism: A Flexible Design for Philanthropic Matching Funds,” *Management Science*, 2018. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3243656.

[51] Optimism Collective, “Retroactive Public Goods Funding: Impact = Profit.” [Online]. Available: https://app.optimism.io/retropgf.

[52] K. Owocki et al., “Scaling Funding for Blockchain-era Public Goods,” *Gitcoin Governance*, 2023. [Online]. Available: https://gov.gitcoin.co/t/scaling-funding-for-blockchain-era-public-goods/9797.