NIST Internal Report
NIST IR 8536

Supply Chain Traceability Principles:

A Manufacturing Meta-Framework

Final

Michael Pease
Evan Wallace
Harvey Reed
Robert Martin
Dr. Vivian L. Martin
Steve Granata

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8536

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Internal Report
NIST IR 8536

Supply Chain Traceability Principles:

A Manufacturing Meta-Framework

Final

Michael Pease
Applied Cybersecurity Division
Information Technology Laboratory

Evan Wallace
System Integration Division
Engineering Laboratory

Harvey Reed
Robert Martin
Vivian L. Martin
Steve Granata
MITRE

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8536

September 2026

U.S. Department of Commerce
Howard Lutnick, Secretary

National Institute of Standards and Technology
Arvind Raman, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at https://csrc.nist.gov/publications.

NIST Technical Series Policies
Copyright, Use, and Licensing Statements
NIST Technical Series Publication Identifier Syntax

Publication History
Approved by the NIST Editorial Review Board on 2026-08-18

How to Cite this NIST Technical Series Publication
Pease M, Wallace E, Reed H, Martin R, Martin VL, Granata S (2026) Supply Chain Traceability Principles: A Manufacturing Meta-Framework. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) NIST IR 8536. https://doi.org/10.6028/NIST.IR.8536

Author ORCID IDs
Michael Pease: 0000-0002-6489-2621
Evan Wallace: 0000-0001-9368-5616
Harvey Reed: 0000-0002-4589-2677
Vivian L. Martin: 0009-0000-8698-4730
Robert A. Martin: 0000-0002-7373-2275

Additional Information
Additional information about this publication is available at https://csrc.nist.gov/pubs/ir/8536/final, including related content, potential updates, and document history.

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof-of-concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of information other than national security-related information in federal information systems. The Special Publication 800-series reports on ITL's research, guidelines, and outreach efforts in information system security, as well as its collaborative activities with industry, government, and academic organizations.

Patent Disclosure Notice

NOTICE: The Information Technology Laboratory (ITL) has requested that holders of patent claims whose use may be required for compliance with the guidance or requirements of this publication disclose such patent claims to ITL. However, holders of patents are not obligated to respond to ITL calls for patents, and ITL has not undertaken a patent search in order to identify which, if any, patents may apply to this publication.

Following the ITL call for the identification of patent claims whose use may be required for compliance with the guidance or requirements of this publication, notice of one or more such claims has been received.

By publication, no position is taken by ITL with respect to the validity or scope of any patent claim or of any rights in connection therewith. The known patent holder has, however, provided to NIST a letter of assurance stating either (1) a general disclaimer to the effect that it does not hold and does not currently intend to hold any essential patent claim(s), or (2) that it will negotiate royalty-free or royalty-bearing licenses with other parties on a demonstrably nondiscriminatory basis with reasonable terms and conditions.

Details may be obtained from meta-framework_nccoe@nist.gov.

No representation is made or implied that this is the only license that may be required to avoid patent infringement in the use of this publication.

Abstract

Manufacturing supply chains are vital to national security and economic resilience. As these networks become increasingly complex, tracing products back to their origins becomes more difficult and exposes acquirers to greater risks from logistical disruptions, fraud, and counterfeit materials. To address the critical need for verifiable supply chain traceability data, a Manufacturing Supply Chain Traceability Meta-Framework is introduced in this report. The Meta-Framework is presented as a practical, conceptual approach for organizing, linking, and querying traceability data across diverse manufacturing supply chain ecosystems. Released alongside an open-source Python reference implementation, the Meta-Framework presents practical technical considerations for organizing, linking, and querying traceability data across disparate manufacturing ecosystems.

By leveraging common structural data patterns, such as encapsulation, linking, and interoperable interfaces, the Meta-Framework is designed to illustrate how to build a continuous, temporally ordered provenance chain spanning industry sectors and geographic locations. The Meta-Framework is an extension of previous NIST research (NIST IR 8419) and reflects input from industry, standards organizations, and academic collaborations. Templates and tailoring are used to preserve flexibility, facilitating the alignment of payload pedigree data with specific industry standards and manufacturing needs. By abstracting internal operations data into standardized, shareable supply chain event data and using cryptographically verifiable links, the framework is used to establish a secure timeline of events that reflects product provenance without requiring a centralized repository. Additionally, a set of supply chain traceability principles is presented to guide broader, risk-informed, and interoperable approaches to traceability.

This structured approach helps stakeholders independently verify product history. It emphasizes the principle of selective disclosure, allowing organizations to share necessary supply chain traceability data while protecting sensitive intellectual property. Ultimately, the foundational principles and conceptual mechanisms required to support verifiable traceability are presented within this report, demonstrating how stakeholders can enhance visibility and support supply chain risk management while preserving flexibility and protecting proprietary information.

Keywords

cryptographic linkage; integrity; interoperability; meta-framework; pedigree; provenance; selective disclosure; supply chain risk management; supply chain traceability; traceability chain; verifiable traceability.

Acknowledgments

The authors would like to acknowledge the contributions from the Community of Interest (COI) and their valuable suggestions and comments, which helped improve the paper. They would

also like to thank Robert Harder (MITRE) for his technical support for the Meta-Framework's reference implementation and for reviewing the data model and use cases.

Disclosure

This manuscript was edited with the assistance of NIST Enterprise subscription to Google's Gemini Platform and Grammarly, which includes Generative AI tools developed by Google and Grammarly Inc. Gemini and Grammarly were used to refine language, improve clarity, and enhance readability in accordance with the authors' instructions. All content, scientific claims, and conclusions have been reviewed and verified by the authors to ensure accuracy and originality.

Table of Contents

List of Tables

List of Figures

Executive Summary

Supply chains are the bedrock of national security and economic resilience. However, as these global networks become increasingly complex and opaque, they present a highly lucrative attack surface for sophisticated adversaries. The inability to definitively trace a product back to its origins exposes acquiring organizations to significant risks, including counterfeit hardware, unauthorized firmware tampering, and hidden dependencies on compromised upstream suppliers. Beyond national security implications, these supply chain blind spots directly translate into critical business impacts, including significant financial losses, prolonged operational disruptions, regulatory penalties, and lasting reputational harm.

To address the critical need for verifiable supply chain traceability data, a Manufacturing Supply Chain Traceability Meta-Framework is introduced within this report. The Meta-Framework presents a conceptual outline of the principles and structural data elements necessary to harmonize traceability information across organizations. Designed as a flexible architectural guide, this structure complements existing operations by enabling organizations to share inter-organizational data while maintaining their current manufacturing systems, data management structures, and proprietary business processes.

By leveraging common structural data patterns, including encapsulation, linking, and interoperable interfaces, the Meta-Framework is designed to integrate data expressed in established industry standards, creating a continuous, ordered provenance chain spanning industry sectors and geographic locations. At the same time, flexibility is preserved so that internal pedigree data aligns with specific industry standards and manufacturing needs. Guidance is provided on how stakeholders can structure inter-organization supply chain events (e.g., Make, Assemble, Store, Ship, Receive, and Employ) as traceability records and connect them across the supply chain network using cryptographically verifiable traceability links. To address industry concerns about intellectual property and operational security, the principle of selective disclosure is integrated into the framework. Organizations share only the minimum structural event data needed to construct a continuous traceability chain, while keeping sensitive, proprietary manufacturing evidence gated within their own controlled repositories.

To guide the design and implementation of interoperable federated ecosystems, a core set of organizing principles, categorized into three fundamental themes is presented:

- **Value-Driven Participation:** Traceability efforts should directly support an acquirer's operational risk decisions, scaling proportionally to the product's risk and using market incentives to drive supplier adoption.
- **Pedigree and Provenance Resilience:** Real-world supply chains contain missing data, analog paper, and tightly guarded trade secrets. Traceability systems must account for the lack of perfect transparency, gracefully measuring and communicating known data gaps, and allow for alternative evidence.
- **Decentralized Trust:** Accountability is maintained through verifiable organizational attribution, strong cyber-physical linkage mechanisms that anchor digital records to products, and prioritization of traceability record interoperability over uniformity.

By outlining a conceptual approach to structuring and exchanging supply chain event data, the Meta-Framework is used to illustrate how complex supply chain information can be organized to support enterprise risk managers engaged in Supply Chain Risk Management (SCRM). Ultimately, the foundational principles and conceptual mechanisms are presented to support verifiable traceability and to provide senior leaders with the actionable intelligence required for cost-effective, defensible risk decisions while preserving market flexibility and protecting proprietary corporate data.

1. Introduction

Organizations obtain and provide products through complex global supply chains and are increasingly concerned about supply chain risks, including malicious tampering, counterfeit components, regulatory noncompliance, and other potential supply chain vulnerabilities [1][2]. Because these supply chains are critical to economic and national security, organizations need visibility into how the products and technologies they procure are sourced, developed, manufactured, assembled, and deployed. Without this, an organization's exposure to supply chain risks is difficult to measure. To help mitigate these risks, cybersecurity requirements and guidance, including the latest versions of the NIST Special Publication (SP) 800-161 [3], NIST SP 800-53 [4], and the NIST Cybersecurity Framework [5], establish Supply Chain Risk Management (SCRM) control objectives for organizations to perform independent due diligence to ensure the security, resilience, and quality of the products they procure and deploy within their own distribution channels or environments [6][7].

To effectively support these control objectives, organizations must trace and obtain product information from suppliers and manufacturers. However, a significant information sharing asymmetry exists. While acquirers need supply chain traceability to understand and manage their risk exposure, manufacturers often guard the composition of their supply chains as proprietary business information. Suppliers and manufacturers are understandably hesitant to provide information that could expose intellectual property or trade secrets. Therefore, a conceptual approach is required to balance this tension, satisfying the acquirer's SCRM data needs, safeguarding the supplier's business interests, and ensuring that necessary information is accessible for effective and efficient due diligence [6][7].

Specifically, to perform this due diligence effectively, the essential information organizations require centers on product **pedigree** (data validating the product's internal composition) and **provenance** (the chronological history of its origins, changes, and transfers). **Traceability** serves as the operational enabler for tracking and verifying this information across the product's supply chain. By bridging the current visibility gap, supply chain traceability facilitates the secure capture and access of the SCRM data necessary to validate a product's origins and composition, eliminating the requirement for suppliers to disclose sensitive internal processes.

Obtaining and analyzing pedigree and provenance information across complex supply chain networks is essential yet challenging. Supply chains are rarely straightforward or linear; instead, they tend to be complex, multi-tiered structures. Furthermore, information sharing seldom extends beyond typical inter-organization Business-to-Business (B2B) bilateral exchanges, often remaining isolated within individual organizations or shared only through disjointed, point-to-point agreements. This fragmentation creates semantic gaps and broken chains of trust between tiers, making it challenging for acquirers to trace a product back to its origin or verify its composition within broader supply chain networks [3][4].

Historically, attempts to overcome these traceability challenges have relied on highly centralized data repositories to collect and distribute information [3][6]. This approach poses an unacceptable risk to organizations by requiring the release of sensitive operational data. However, enabling traceability does not mean that all internal product information must be

made publicly available. To protect intellectual property, the industry requires a shift in focus away from monolithic central repositories and all-or-nothing data publishing. Organizations require an approach in which the targeted traceability data records that support pedigree and provenance information are independently verifiable and readily accessible to authorized downstream acquirers. The core traceability principles that vendors and manufacturers can consider to generate and publish, including SCRM-relevant records that are accessible and can be safely accessed across federated ecosystems, empowering acquirers with the pedigree and provenance needed to support SCRM processes, are explored herein.

1.1. Authority and Statutory Background

The National Institute of Standards and Technology (NIST) developed this document in furtherance of its statutory responsibilities under the Federal Information Security Modernization Act (FISMA) of 2014, 44 U.S.C. § 3551 et seq., Public Law (P.L.) 113-283. NIST is responsible for developing information security standards and guidelines, including minimum requirements for federal information systems. However, such standards and guidelines shall not apply to national security systems without the express approval of appropriate federal officials exercising policy authority over such systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130.

Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and binding on federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other federal official. This publication is provided as a reference and may be used by federal agencies and nongovernmental organizations on a strictly voluntary basis. It is not subject to copyright in the United States. Attribution would, however, be appreciated by NIST.

1.2. Purpose

The aim of this report is to outline key principles for generating, linking, and validating acquirer-relevant supply chain data to map a product's supply chain timeline (provenance). By proposing a conceptual, consistent method for connecting digital traceability records, the objective is to facilitate independent tracing and verification of pedigree and provenance information while also addressing and mitigating historical challenges in information sharing. It emphasizes the importance of data record verifiability to support SCRM requirements and to protect supplier intellectual property and proprietary business information.

1.3. Scope

This report is designed to facilitate dialogue and establish a conceptual foundation for acquirer SCRM use cases. While numerous industries and standards (e.g., [8],[9],[10],[11]) currently support elements of supply chain visibility, the concepts presented here are not intended to replace them. Instead, this document is structured to complement existing frameworks.

Furthermore, this report does not constitute a formal regulatory standard, a compliance mandate, or an exhaustive analysis of prior traceability efforts.

Terminology and Cross-Sector Applicability

This report distinguishes between *industry*, *sector*, and *domain*, while recognizing that these terms may overlap in common usage. *Industry* refers broadly to communities of organizations engaged in related economic or production activities. *Sector* refers to a grouping of related economic, operational, governmental, or regulatory activities and may also reflect an established classification or context. The compound term *industry sector* is used broadly to refer to a recognized grouping or segment of industry and does not imply a separate level of classification from *industry* or *sector*. *Domain* describes an area of manufacturing, technology, or application that may span multiple industries or sectors (e.g., discrete manufacturing or software). Manufacturing supply chains frequently cross these boundaries; therefore, the Meta-Framework is intended to support traceability across industry sectors, domains, and geographic locations. Where *Critical Infrastructure (CI)* is used, it refers to the U.S. critical infrastructure context and may also be used as shorthand to identify an organization, acquirer, environment, or ecosystem operating within that context; it is not used as a synonym for industry, industry sector, or domain.

The Scope of Traceability

Within this document, the scope of traceability is limited to the inter-organizational flow and handoffs of specific products between entities, covering the supply chain lifecycle up to the point of initial deployment or acquisition (e.g., the “Employ” event). The objective is to enable acquirers to verify the attribution of products to their producing organizations. It explicitly does not address the intra-organizational traceability mechanisms manufacturers use for internal process optimization, inventory control, or routine B2B logistics management.

Primary Focus: Discrete Manufacturing and Software

The structural data patterns outlined in this report primarily focus on the discrete manufacturing and software domains. To help visualize complex supply chains, multiple federated ecosystems, and the application of traceability data records, a primary, running use case centered on a discrete manufacturing process (e.g., an enterprise assembling components from multiple suppliers to produce a security camera) is utilized. Specifically, this primary focus includes:

- **Discrete part manufacturing:** The production of distinct, countable items.
- **Discrete assembly manufacturing:** The combination of discrete parts into a more complex, discrete product.
- **Applicable software:** The development, compilation, and integration of specific software components and applications.

Broader Applicability: Batch, Lot, and Continuous Flow

While the primary illustrative model focuses on discrete manufacturing, the foundational traceability principles are conceptually applicable across a wider range of manufacturing methodologies. To address industry needs, high-level narrative examples are provided for batch, lot, and continuous-flow processing (e.g., processes involving raw materials, liquids, or bulk commodities).

However, the exhaustive application of these traceability mechanisms to non-discrete domains and to post-deployment lifecycle phases (e.g., maintenance, refurbishment, disposal, and retirement occurring after the Employ event) introduces distinct complexities. These areas are reserved for broader discussion and future research, as outlined in Appendix E.

1.4. Target Audience

This document is intended for a broad audience of stakeholders responsible for assessing and managing supply chain risk, as well as those who design, operate, or participate in supply chain traceability ecosystems. The primary target audiences include:

- **Acquiring Enterprises, and SCRM Professionals:** Consumers, Government agencies, acquisition officials, and risk managers who require verifiable pedigree and provenance data to perform independent due diligence, ensure product integrity, and manage supply chain risk.
- **Manufacturers and Component Suppliers:** Large manufacturing primes that act as ecosystem anchors, encouraging adoption across their supply networks, as well as small and medium-sized manufacturers that need interoperable approaches to securely share traceability records without exposing proprietary business information.
- **Industry Consortia and Sector-Based Working Groups:** Organizations and governance bodies that define common data requirements, establish standards, and provide oversight for traceability ecosystem operations.
- **Technology Providers and System Integrators:** Entities tasked with building technical infrastructure to support traceability records, verifiable data exchanges, and participant interfaces.

1.5. Document Structure

The remainder of this document is organized as follows:

- **Section 2: Background** - Contains an overview of the current manufacturing supply chain landscape and the imperatives driving the need for enhanced traceability. The challenges posed by information asymmetry, the visibility gap, and the need for verifiable pedigree and provenance data to support acquirers' due diligence and SCRM are explored.

- **Section 3: Traceability Principles** - Contains an analysis of the foundational principles required to achieve supply chain traceability. This section outlines the conceptual requirements for enabling verifiable visibility and secure data exchange across multi-tiered supplier ecosystems.
- **Section 4: A Manufacturing Meta-Framework** - Contains a notional data structure for capturing, linking, and validating supply chain traceability data. This section details the specific traceability metadata and cryptographic mechanisms that allow acquirers to independently verify traceability records without requiring access to a supplier's proprietary internal data.
- **Section 5: Conclusion** - Contains a summary of the key findings of the report and reiterates the value of the traceability principles and the Manufacturing Meta-Framework in addressing SCRM requirements.
- **References** - Contains a list of all sources and literature cited within this document.
- **Appendices** - Contain supplementary material and supporting documentation.

1.6. Linguistic Conventions

The terms "shall" and "shall not" indicate requirements to be followed strictly to conform to the publication, and from which no deviation is permitted. The terms "should" and "should not" indicate that among several possibilities one is recommended as particularly suitable, without mentioning or excluding others, or that a certain course of action is preferred but not necessarily required, or that (in the negative form) a certain possibility or course of action is discouraged but not prohibited. The terms "may" and "need not" indicate a course of action permissible within the limits of the publication.

Note that the use of these terms (particularly "shall") within this document to describe the correct technical implementation of a concept or framework does not imply that this publication itself is a mandatory or legally binding federal requirement.

2. Background

To understand why a new approach for verifiable traceability records is required, it is necessary to examine how the supply chain is currently managed and where those legacy methods fall short. Today, product lifecycle and manufacturing data are largely siloed within individual organizations. When information is shared, it typically occurs through linear, "one-up, one-down" transactions between direct business partners.

Because there is no standardized mechanism to securely link these isolated transactions across multiple tiers or disparate industries, acquirers are left with blind spots regarding the origins and pedigree of their assembled products. The following subsections illustrate the complexity of a multi-tiered manufacturing supply chain, the specific data characteristics that acquirers require for SCRM, and the friction involved in extracting this data across organizations to establish true end-to-end provenance.

The supply chain visibility challenges are summed up in the most recent Quadrennial Supply Chain Review, where the National Economic Council & National Security Council stated that "limited supply chain visibility remains a challenge, meaning that companies still lack sufficient insight into the materials, components, [products], and processes across their networks. Many have been investing in technologies to enhance real-time tracking, but more work is needed." [12].

2.1. The Operational Reality: Internal vs. Supply Chain Event Data

A supply chain is the connected network of processes and partners that moves raw materials through component production and product assembly, managing the flow of products, including goods and materials, until the finished item reaches the acquiring end organization. To illustrate this, Fig. 1 depicts a notional example of an Original Equipment Manufacturer (OEM), ACME, that produces security cameras for Critical Infrastructure (CI). A networked security camera serves as an example because it is a discrete, multi-component cyber-physical product. Given its crucial role in CI environments, acquiring organizations must rely on SCRM to ensure the device's integrity is not compromised prior to deployment.

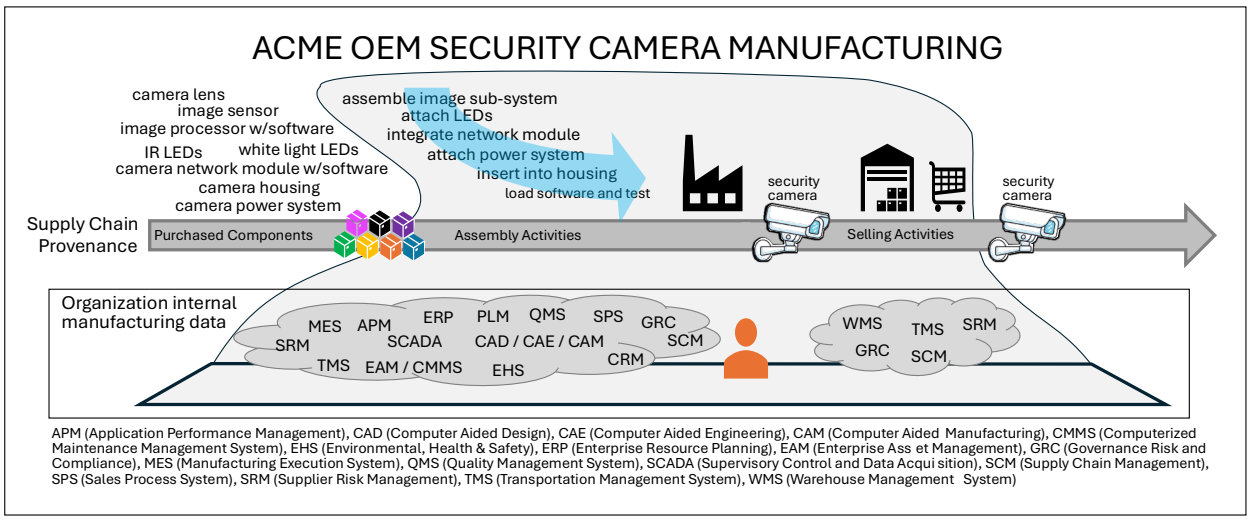

Fig. 1. OEM Manufacturer for Security Cameras

To manufacture these security cameras, ACME purchases required components such as sensors, lenses, and network modules. These components are typically purchased from other manufacturers in the supply chain, some of whom may be OEMs themselves. The progression of inter-organizational B2B steps (e.g., purchasing components, assembling them, and selling the final product) is shown in the gray supply chain timeline in Figure 1. Importantly, the information required to describe these supply chain events is captured in internal systems (e.g., Manufacturing Execution Systems (MES), Enterprise Resource Planning (ERP), etc.) that contain highly sensitive privacy and intellectual property information. Therefore, a distinction must be made between internal manufacturing operations and shareable supply chain event data. For the purposes of supporting supply chain risk management, the minimum event data extracted from internal systems and shared is sufficient to identify the manufactured product, the organization that performed the supply chain activity, and the timestamps of the activity and record creation.

Figure 2 expands this view to show ACME within a larger supply chain context, situated between upstream component suppliers, distributors, and downstream critical infrastructure acquirers.

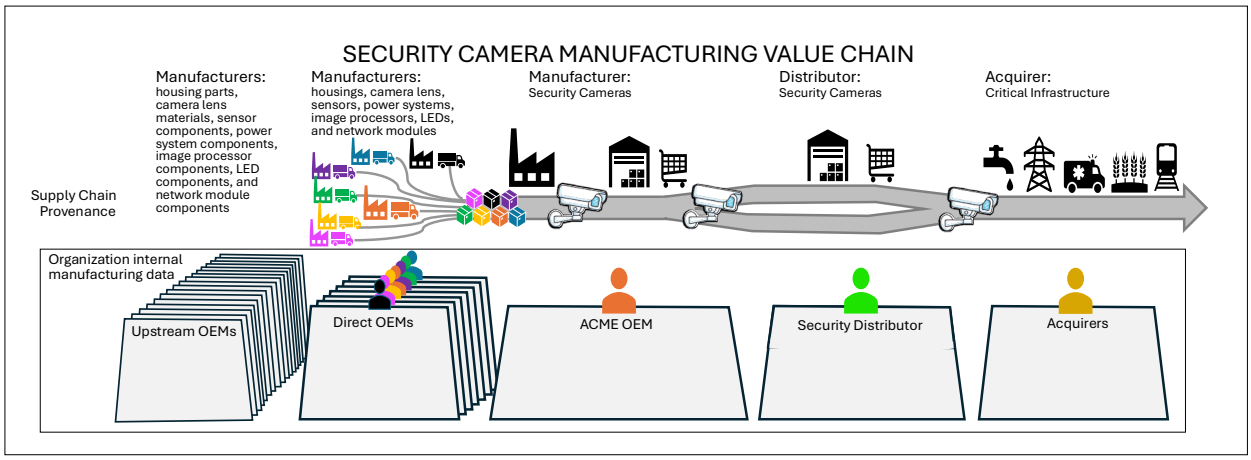

Fig. 2. OEM Manufacturer in Supply Chain Context

From the acquirer's perspective, the purpose of the supply chain is to produce security cameras for installation in their environments. However, manufacturing supply chains are vast and self-organized, forming massive, distinct ecosystems around commonalities in product types, raw-material processing, and more. These ecosystems tend to have their own conversational and data languages. Despite these varying languages, organizations performing SCRM require consistent, high-level provenance data (e.g., product identification, movement through attributed organizations, etc.) to validate product integrity, without requiring access to upstream suppliers' internal manufacturing data, intellectual property, or privacy-protecting systems.

2.2. The Nature of SCRM Information

SCRM analysis of manufactured products is increasingly important for reducing risk. Without clear visibility into a product's origins, manufacturers and acquirers face significant blind spots that expose them to operational and security threats. These risks include the introduction of counterfeit or substandard components, unauthorized hardware or firmware substitutions (tampering), and hidden dependencies on compromised or sanctioned upstream suppliers. [3],[14],[15],[16] By establishing an unbroken, verifiable chain of custody, effective SCRM directly reduces these risks. It enables organizations to proactively verify component authenticity and verify critical cyber-physical systems, such as the ACME security camera, before deployment.

To perform this analysis successfully, acquirers must obtain accurate and trustworthy provenance and pedigree information. Modern manufacturing supply chains form a complex network of globally distributed stakeholders, processes, and systems, with a variety of technologies, varying levels of digitization of manufacturing records, diverse digital record repositories, and decentralized ownership and control of data.

Establishing reliable SCRM visibility requires satisfying several distinct aspects of provenance:

- **Identification:** Establishing and maintaining unique identification of the supply chain elements, processes, and organizations associated with a system and critical system

components. This binds the specific instance of a manufactured product to the organization's attribution in a single digital record. For example, uniquely identifying a specific image sensor produced by an upstream supplier before it arrives at the ACME facility.

- **Traceability (Tracking and Tracing):** Capturing and chronologically linking key supply chain events to form a continuous history. Building upon identification, traceability mechanisms are used to identify a product's component parts and access metadata captured throughout its lifecycle. In our ACME example, this allows an acquirer to reliably track the journey of the camera assembly components, such as the image sensor, through its transformation events to the material components' points of origin.
- **Validation:** Determining whether a component is genuine and unaltered. Through this capability, the acquirer can confirm that the product received is the correct component and has not been compromised since its initial manufacture. For example, validating that the camera's network module received by ACME matches the exact specification claimed by the original manufacturer and is free from tampering.
- **Pedigree and Integrity Assurance:** Verifying the comprehensive composition of a product and the cryptographic wholeness of its history. This requires linking the pedigree evidence for components and assemblies to establish an unbroken, trusted provenance chain. In our ACME example, this ensures the critical infrastructure operator can verify the pedigree of the security camera system informed by data from multiple manufacturers.

While these four capabilities support effective SCRM, the specific data languages and standards used to record them vary widely across industries. Each industry supply chain network typically relies on its own established data models and ontologies, compliance frameworks, and communication protocols. For instance, the aerospace sector may use fundamentally different pedigree terminology and data formats than the microelectronics or chemical sectors.

Furthermore, fragmented data languages can create an interoperability challenge. To successfully trace a product from its raw materials to its final assembly, acquiring organizations need a mechanism that can bridge these varying data standards without forcing all suppliers to abandon their established, industry-specific formats.

2.3. The Visibility Challenge and Traceability Chains

When an acquirer requests SCRM provenance and pedigree information from suppliers, the visibility and accessibility decrease the further back in the supply chain an acquirer goes. The increasing opacity illustrates this challenge in Fig. 3 [1], where the greater the distance in the number of tiers from the acquiring enterprise, the more difficult it is to obtain SCRM data.

Because of this opacity, the traversal path of a critical component, from raw material to the origin of the component, and then to the assembly where the component is used, may be difficult to establish with certainty. If SCRM pedigree data for raw materials, components, or

[1] Adapted from [NIST SP 800-161r1, Fig 2]

assemblies is unavailable, validating that a component is genuine and unaltered is difficult, if not impossible. For example, to validate a security camera sensor, an acquirer must piece together pedigree data into a provenance timeline for the raw materials, sensor components, and final assembly. However, because each discrete pedigree record is generated by a different manufacturer and captured in its own isolated, native repository, this information is often difficult or impractical for the acquirer to obtain, access, and reconstruct.

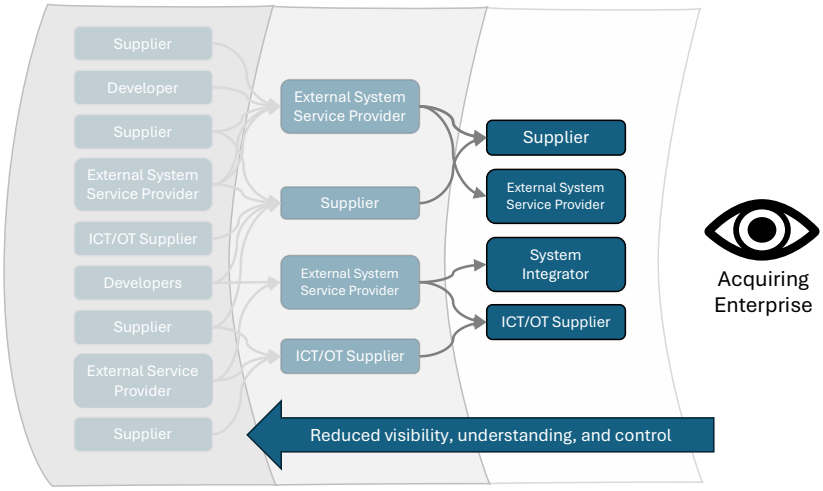

Fig. 3. Supply Chains Exhibit Reduced Visibility Over Tiers

To overcome this fundamental visibility challenge without forcing the entire supply chain into a single, centralized repository, acquirers need a mechanism to connect these distributed records. By relating previously disconnected pieces of information to one another in chronological order and causal dependencies across their respective repositories, organizations can form a traceability chain, as illustrated in Fig. 4. A traceability chain serves as the structural foundation for provenance by providing an ordered set of traceability records, each representing a supply chain event with provable data integrity.

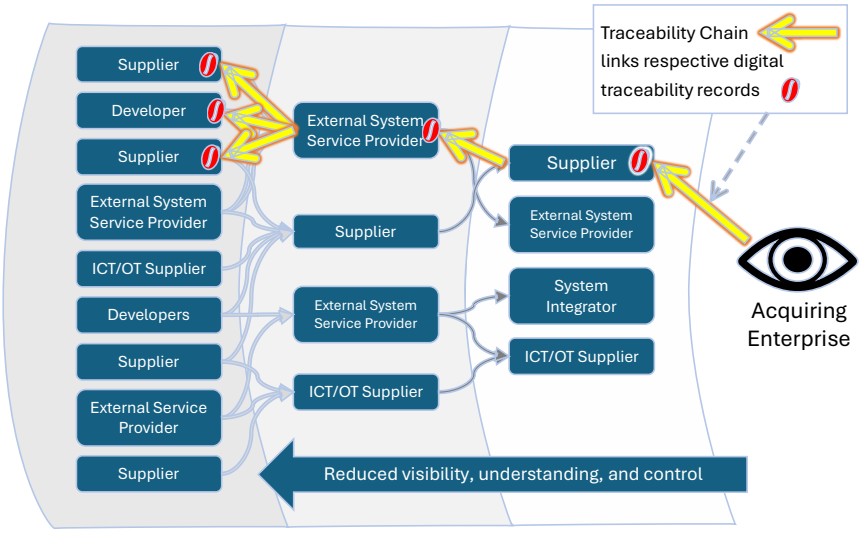

Fig. 4. SCRM Traceback

SCRM is supported by traceability chains through the provision of a verifiable chronology of key supply chain events and possesses the following core characteristics:

- **Chronological Order and Causal Dependencies:** The chain contains a continuous timeline of individual supply chain event records, accurately reflecting the chronological sequence of operations and handoffs between manufacturers.
- **Scalable Scope:** The depth of the chain is scalable; for critical components or raw materials, the chain can extend from the acquirers back to the point of origin.
- **Cross-Ecosystem Continuity:** The chain remains unbroken even when spanning disparate federated ecosystems, relying on structural interoperability to bridge the varying standards and data languages used across the broader network.

However, missing traceability records (e.g., due to failure to record an event, lack of visibility, or lack of access) create traceability gaps. These traceability gaps can hinder SCRM analysis and may prevent validation of a component's genuineness and integrity.

2.4. Summary

Supply chains are networks of manufacturers and vendors that provide raw materials, components (both physical and digital), and assemblies to an acquirer. Acquirers perform SCRM to verify that constituent components and assemblies are genuine and unaltered. This requires provenance information, written by the applicable manufacturer at the time of the supply chain event, without exposing detailed internal manufacturing data. Together, these event records form a traceability chain. However, modern supply chains face traceability gaps caused by missing, isolated, or inaccessible records. The following section outlines principles to guide efforts to implement supply chain traceability, close traceability gaps, and improve traceability.

3. Supply Chain Traceability General Principles

To maintain alignment and interoperability among independent traceability efforts across disparate industries, a core set of guiding principles is established within this framework. These principles are not intended to replace specific, documented technical requirements; instead, they provide the overarching context for decentralized decision-making. By establishing a shared vocabulary and communicating the "what" and "why" of supply chain traceability, they enable individual ecosystems to build out their own systems while maintaining structural compatibility with the broader global network.

In alignment with the National Cybersecurity Center of Excellence (NCCoE) mission of applied, practical cybersecurity, these principles were developed through stakeholder engagement and collaborative. NIST initially developed a reference implementation of the Supply Chain Traceability Meta-Framework (detailed in Section 4). Through iterative development, real-world testing, and extensive industry collaboration, NIST abstracted the following generalized principles from direct community input and lessons learned.

As established in the previous section, manufacturing supply chains are highly complex and reflect the variety of industries, industry sectors, and domains of the global economy. While many organizations possess tracing capabilities designed to support their specific business community, most have not addressed the acquirer's (acquiring enterprises/organizations) need to trace the supply chain timeline (provenance) of an acquired item when that timeline crosses multiple distinct industries, industry sectors, or domains.

This specific need, together with the operational insights from the reference implementation, has informed the following general principles. To guide traceability efforts toward a pragmatic, resilient, and interoperable approach, these principles are organized into three fundamental themes:

- **Value-Driven Participation (Principles 1–3):** Traceability should directly support acquirer risk decisions and scale proportionally to risk, relying on economic incentives instead of mandates to drive adoption.
- **Pedigree and Provenance Resilience (Principles 4–6):** Traceability approaches should accommodate the realities of missing data, protected trade secrets, analog boundaries, and continuous lifecycle updates.
- **Decentralized Trust (Principles 7–9):** Trust is anchored in the individual traceability records through verifiable organizational identity, robust cyber-physical and digital anchors, consistent data encapsulation, and cryptographic links that support data integrity and secure record connections, maintaining unbroken integrity without dictating uniform data formats or centralized storage mechanisms.

These principles, enumerated in Tables 1, 2, and 3, are intended to guide supply chain traceability efforts so that independent manufacturers can write and link traceability records, allowing acquirers to make defensible, resilient procurement and operational SCRM decisions. They are designed to support verifiable traceability even when parts of the supply chain are opaque, without excluding key suppliers or disrupting existing industrial supply chain networks.

Because global supply chains vary significantly in technical maturity, regulatory burden, and operational scale, these principles may not apply in whole or in part to every industry or sector. They are designed as an architectural compass, not a rigid compliance checklist.

Theme 1: Value-Driven Participation

The principles within this theme focus on aligning traceability efforts with business value and operational risk. Table 1 outlines how organizations can scale their traceability requirements in proportion and leverage market incentives to drive meaningful adoption across the supply chain.

Table 1. Principles for Value-Driven Participation

Principle	Description	Primary Impact
1. Decision-Centric Traceability	External traceability data sharing focuses on supporting specific organizational SCRM decisions for the acquiring organization.	Traceability data maps to the acquirer's procurement requirements and regulatory obligations, avoiding unnecessary data hoarding.
2. Proportional, Risk-Scaled Traceability	The depth and rigor of traceability records are scaled to match the mission impact and systemic risk posed by the acquired item.	Traceability is graduated and proportional to operational consequences, preventing unnecessary reporting burdens on manufacturers and suppliers for lower-criticality components.
3. Incentive-Aligned Participation	Economic and operational incentives are aligned to motivate upstream participants to generate and share records without requiring uniform capabilities.	Traceability maturity evolves incrementally; critical suppliers are not excluded from the framework due to temporary capability gaps.

Theme 2: Pedigree and Provenance Resilience

This theme addresses the practical challenges of capturing supply chain data in complex, real-world environments. Table 2 outlines the principles necessary for handling missing data, protecting sensitive trade secrets, and maintaining an accurate historical ledger as products evolve throughout their lifecycle.

Table 2. Principles for Pedigree and Provenance Resilience

Principle	Description	Primary Impact
4. Evidence-Based Provenance, Not Perfect Transparency	Traceability approaches utilize verifiable alternative evidence (e.g., certifications, audits) when direct or continuous digital event data is unavailable.	Suppliers can provide bounded disclosures that support validation while protecting intellectual property and sensitive trade secrets, ensuring the overall provenance chain remains continuous and verifiable even when granular event data cannot be shared.
5. Measure and Communicate Traceability Gaps	Traceability chains expose data gaps and trust boundaries by design, ensuring limitations in supply chain visibility are clearly identified and measured.	Procuring organizations can make informed, defensible risk decisions even when traceability data is incomplete or withheld.
6. Lifecycle-Oriented Provenance and Accountability	Traceability extends across the product's lifecycle to reflect modifications, firmware updates, and operational events.	By functioning as a cumulative historical ledger, the traceability chain supports the maintenance of an accurate, up-to-date view of the item's composition and status throughout its operational lifespan.

Theme 3: Decentralized Trust

To scale globally, traceability frameworks must operate across independent trust boundaries without relying on a single central authority or monolithic database. Table 3 enumerates the principles that enable verifiable organizational attribution, secure cyber-physical anchoring, and consistent interoperability across disparate ecosystems.

Table 3. Principles for Decentralized Trust

Principle	Description	Primary Impact
7. Verifiable Organization Attribution	Associating records with verifiable identifiers (e.g., GLN, UEI, LEI) establishes clear organizational accountability for supply chain events.	Facilitates the reliable attribution of every traceability record to a specific entity, thereby supporting decentralized trust across the supply chain network.
8. Cyber-Physical and Digital Linkage	Reliable digital anchors associate traceability records with the specific physical product, digital object, or batch they describe.	Mitigation of the substitution of unverified products or software, establishing a verifiable link between the physical item and its digital record.
9. Interoperability Over Uniformity	Addresses the capability to retrieve and link records across industry standards, formats, and storage mechanisms.	Facilitates data exchange and linking across ecosystem-specific repositories without requiring a single, monolithic data or technology stack.

Each of these principles is explored in greater detail within the following subsections. To illustrate their practical application, each section contains notional examples across various supply chain contexts. Furthermore, each principle includes a discussion of its operational and architectural implications, demonstrating how they collectively shape a pragmatic, resilient, and interoperable approach to supply chain risk management.

3.1. Principle 1. Decision-Centric Traceability

While internal data collection (such as within Manufacturing Execution Systems) focuses on detailed operational efficiency and telemetry, external traceability data sharing should focus on utility for the acquiring organization. This principle establishes that the value of traceability efforts is not measured by the sheer volume of data collected or the number of participating suppliers, but by how effectively the data supports specific organizational decisions, regulatory compliance, and operational resilience.

> **Notional Example: Continuous Process Sourcing.** Consider a city water authority acquiring treatment chemicals under Environmental Protection Agency (EPA) Safe Drinking Water Act obligations. Historically, verifying these shipments required a manual review of unstructured, bulk compliance documentation. Using the Meta-Framework, the authority instead queries the traceability records linked to the specific shipment. As they trace the links back

through the supply chain, they obtain pedigree information such as purity thresholds and unbroken cold-chain handling. By providing this historical data in a structured, linked format, the acquiring entity's engineering, legal, and public health teams can efficiently validate the provenance of the chemicals against their approval thresholds.

Notional Example: Discrete Manufacturing Procurement. Consider a state Department of Transportation (DOT) procuring steel girders for a bridge rehabilitation project that must simultaneously satisfy Buy America requirements, environmental compliance standards, and structural-grade certifications. Traditionally, the DOT would have to manage three separate verification workstreams. The DOT can use aggregated traceability data from the traceability chain to map each prospective supplier's provenance across all three obligation sets in a single, consolidated view. This allows the project engineer, procurement officer, and agency counsel to confirm that their respective requirements are met through a single shared decision-support framework.

Aligning to this principle shifts the focus from data hoarding to targeted intelligence. Individual traceability records should be structured to support SCRM decisions, including procurement requirements and regulatory obligations. Additionally, supply chain specialists should not be the sole consumers of this information. Instead, this data must be presented in ways that multiple organizational functions (e.g., engineering, legal, and operational teams) can easily interpret and apply. This approach also encourages data minimization and helps prevent the over-collection of proprietary data, ensuring that suppliers share only the external traceability data necessary to support the specific decision-making process. By centering the framework on SCRM processes and related decisions required by the acquiring organization (whether an intermediate manufacturer or the final acquirer), traceability systems avoid becoming burdensome reporting exercises and instead function as targeted decision-support capabilities.

3.2. Principle 2. Proportional, Risk-Scaled Traceability

The depth, rigor, and resolution of traceability data should be scaled in proportion to a product’s mission impact, safety criticality, and systemic risk. Applying a uniform level of scrutiny to every component is neither practical nor efficient. This helps ensure traceability investments and supplier burdens are proportional to the operational consequences of failure. Traceability should be graduated, recognizing that a critical industrial control system requires a different level of provenance validation than a commercial-off-the-shelf office peripheral.

Notional Example: Healthcare Supply Chain. Consider a regional hospital network evaluating three competing bulk saline solution suppliers against FDA drug shortage guidance and internal continuity-of-care thresholds. Because bulk saline is a critical commodity but structurally simple, the hospital does not need granular, component-level manufacturing logs. Instead, the required

traceability records provide verified organizational identifiers and geographic sourcing origins. This foundational data can then be cross-checked against external intelligence to flag suppliers with active FDA warning letters or to identify single-point-of-failure vulnerabilities in raw material sourcing (e.g., discovering that all three direct suppliers rely on the same upstream regional facility). This enables the hospital to validate supply chain resilience and make informed, risk-based decisions regarding ongoing sourcing and product acceptance.

Notional Example: Critical Infrastructure Procurement. Consider a regional electric utility procuring substation automation equipment that must meet North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP) compliance requirements. Because of the high systemic risk, the traceability records for acquired substation products are structured to provide deep, component-level answers (e.g., the country of origin for covered sub-components and component inventories. This data can then be mapped against known vulnerability databases and supplier ownership registries for foreign adversary screening. The utility's compliance teams can use these enriched results immediately. The rigorous level of tracing applied here is directly proportional to the catastrophic risk posed by a compromised electrical substation.

Aligning with this principle ensures traceability operates as a proportional, risk-driven capability. Acquirers should focus on provenance information that directly affects operational availability, integrity, safety, security, resilience, and recoverability. By scoping required traceability records strictly to their relevance to decisions and the component’s criticality, organizations avoid imposing unnecessary data collection and reporting burdens on suppliers while still securing their most vital operational assets.

3.3. Principle 3. Incentive-Aligned Participation

A minimum viable traceability implementation requires sufficient incentives to motivate upstream supply chain participants to write and share records. The traceability framework should align economic, operational, and risk-based incentives so that all participating tiers of the supply chain are motivated to generate, maintain, and share traceability data. This alignment helps ensure the data is usable by downstream decision-makers without imposing a disproportionate reporting burden or requiring uniform technological capabilities from participants.

Notional Example: Government Procurement. Consider a prime contractor acquiring component modules. A small, specialized subcontractor may not initially possess the Information Technology (IT) infrastructure to participate in a fully automated traceability framework. However, they can provide a standardized digital "Ship" event record accompanied by a signed attestation

of origin for the components. The acquirer's traceability system accepts this foundational traceability record, but ties access to preferred subcontractor tiers and accelerated payment terms to the subcontractor's eventual ability to provide linked records that extend visibility down to their own tier-2 suppliers. This creates a clear incentive for incremental maturity without excluding a critical supplier.

Notional Example: Industrial Chemical Sourcing. Consider a chemical supplier producing industrial solvents that may provide a basic traceability record for their finished batch, but cannot digitally link that record back to the specific "Receive" events for their incoming raw materials. Downstream procuring organizations can establish voluntary, recognized accreditations for suppliers who link their event records to their upstream supply network. Achieving this "trusted supplier" status allows the chemical manufacturer to elevate their market standing, streamline compliance checks, and access preferred procurement tiers. This encourages a gradual expansion of traceability data tied directly to positive market incentives and brand differentiation, without ever requiring the supplier to expose their proprietary manufacturing formulas or internal equipment logs.

Notional Example: Software Supply Chain and Build Provenance. Consider a critical infrastructure acquirer procuring operational software. A software supplier provides a component inventory without verifiable build provenance that aligns with industry-recognized secure software development frameworks. The acquirer accepts the product based on the component inventories provided but assigns the software a higher operational risk score and requires additional compensating controls for deployment. However, suppliers that mature their capabilities to provide digitally linked traceability records, connecting software artifacts to attested, secure build environments and upstream open-source components using interoperable standards, are granted preferred vendor status and accelerated accreditation. This creates deeper cybersecurity transparency with a procurement advantage, incentivizing the suppliers to improve their software traceability without excluding them from the initial bid.

Aligning with this principle recognizes that voluntary participation in traceability efforts increases significantly when suppliers receive tangible economic or operational benefits, such as preferred sourcing status, reduced audit frequency, or accelerated payments. Traceability maturity should be encouraged to evolve incrementally across the supply chain. Demanding immediate, uniform compliance could stifle progress. Scaling participation based on business incentives, capacity, and risk helps ensure that critical suppliers are not excluded due to temporary capability gaps, fostering a more resilient and transparent supply base over time.

3.4. Principle 4. Evidence-Based Provenance, Not Perfect Transparency

When direct, continuous supplier data is unavailable, either due to a lack of upstream digital records or to protect sensitive trade secrets, traceability should rely on verifiable alternative evidence rather than on absolute transparency. Demanding uninterrupted digital event data across every tier of a supply chain is often impractical and can disincentivize supplier participation by threatening their intellectual property.

Instead, traceability approaches should utilize bounded disclosures (sharing only the specific data necessary to satisfy a risk or compliance requirement rather than the underlying intellectual property). In practice, this means suppliers can encapsulate alternative proof of provenance (e.g., digital attestations, certifications, scanned third-party audits) as part of the data "payload" of a standard traceability record. This allows the overall traceability chain to remain unbroken, supporting acquirer validation without exposing the supplier's proprietary details.

Notional Example: Commodity Sourcing and Non-Digital Origins. When the automotive industry attempts to establish a vehicle's provenance, bulk materials such as leather often lack a continuous digital trail back to their origins. Forcing the creation of a fully digitized, end-to-end chain of events back to the individual ranch is often technologically or economically prohibitive. Consequently, the traceability system should accept alternative verifiable evidence. A bulk leather supplier that provides a verified attestation of origin, along with references to physical records or third-party sourcing certifications on which that attestation is based, serves as acceptable evidence of provenance instead of perfect digital continuity.

Notional Example: Intellectual Property Protection. Consider a defense acquirer procuring specialized microelectronics. The supplier's exact manufacturing techniques, chemical formulations, and specific tier-3 component sources may constitute highly guarded trade secrets. By allowing the supplier to provide a bounded disclosure, the traceability approach eliminates the need for absolute transparency that would otherwise expose their IP to downstream competitors. The supplier writes a traceability record containing a conformance attestation verifying that all underlying components meet specific regulatory requirements (e.g., sourcing only from approved countries), and a third-party audit certificate may support the record. This evidence helps provide the acquirer with the necessary assurance without compromising the supplier's proprietary business intelligence.

Notional Example: Analog and Paper Records. Consider a tier-1 heavy equipment manufacturer sourcing specialized metal castings from a small, legacy foundry. The foundry lacks the IT infrastructure to generate digital

traceability event records; instead, it ships the physical castings with handwritten Material Test Reports (MTRs) and paper bills of lading. This traceability principle accommodates such analog evidence, ensuring that capable legacy suppliers are not excluded and that the lack of digital data is not treated as an insurmountable break in the provenance chain. When the tier-1 manufacturer physically receives the castings, they generate a digital "Receive" event and attach a scanned, verifiable copy (or cryptographic hash) of the paper MTRs as the alternative evidence payload. This preserves the historical provenance for downstream acquirers without forcing the legacy foundry to adopt expensive digital infrastructure.

Aligning with this principle shifts the industry mindset away from an unrealistic standard of omniscient supply chain visibility. It recognizes that digital traceability chains will occasionally reach points where continuous digital event data ceases due to technological limitations or proprietary boundaries. When this occurs, multiple forms of alternative evidence can be acceptable and valid within a traceability chain, including digitized legacy paper records, scanned physical manifests, physical certifications, audits, conformance attestations, and regulatory filings. By deliberately designing traceability records to accommodate bounded disclosures, acquirers enable their suppliers to protect trade secrets, national or regional interests, and upstream contractual commitments.

However, relying on alternative evidence requires the acquirer to carefully assess the level of assurance in the claims. Because bounded disclosures and attestations abstract the underlying data, they may warrant additional scrutiny. There is no universally fixed minimum threshold for acceptable alternative evidence; rather, the appropriate threshold is determined by the acquiring organization based on the item's operational risk. For example, to mitigate the risk of fraudulent claims (e.g., a supplier falsely attesting to a "Made in the USA" country of origin while masking overseas sourcing of components), an acquirer may require third-party verification, conduct random physical audits, or assign a higher risk score to the acquired item. Ultimately, the traceability chain can be used to securely carry alternative evidence while leaving the trust and verification decisions to the acquirer's SCRM processes.

3.5. Principle 5. Measure and Communicate Traceability Gaps and Trust Boundaries

In real-world supply chains, data will occasionally be missing, incomplete, or inaccessible. Traceability chains should allow for the explicit identification of data gaps and trust boundaries. While specific data standards will define their own exact vocabularies, systems should distinguish between states such as "unknown," "withheld," "not applicable," and "unreachable." Whether these gaps are proactively recorded by a supplier or discovered dynamically as a broken link during an acquirer's analysis, the overarching approach should gracefully handle them. Missing information should not be omitted, falsely interpolated, or treated as a catastrophic failure. Deliberately illuminating these distinct states provides acquirers with an honest assessment of supply chain opacity. This enables acquirers to make informed, defensible risk decisions even when conditions are imperfect.

Notional Example: Legacy Data Gaps in Aviation. Consider a Maintenance, Repair, and Overhaul (MRO) facility acquiring a used aircraft part whose complete operational history and original pedigree lack full digital documentation. The traceability system does not attempt to present an artificially clean record or outright reject the part due to this missing information. Instead, the historical maintenance events are explicitly marked within the traceability chain as "not available—pre-digital archive." Because this trust boundary is actively communicated, the MRO facility knows exactly where digital provenance begins and ends. They can then implement appropriate compensating physical inspection protocols, preventing them from unknowingly assuming the part's history is complete or discarding a usable component.

Notional Example: Operational Data Gaps in Healthcare. Consider a shipment of cold-chain pharmaceuticals in which a temperature sensor goes offline for 4 hours during a cross-border transfer. The traceability system avoids mathematically interpolating the missing data to artificially complete the chart or automatically rejecting the entire shipment. Instead, the specific event window is explicitly recorded in the traceability chain as "unknown—sensor offline." A hospital pharmacist reviewing the provenance data can clearly see the bounded gap and apply their own clinical risk thresholds before dispensing the medication.

Notional Example: Structural Gaps and Broken Chains. Consider an industrial equipment manufacturer attempting to trace the provenance of a critical microchip sourced through a third-party distributor. The distributor provides a valid "Receive" record, but the reference to the original manufacturer's "Ship" event points to an inaccessible repository. This could occur because the original manufacturer went out of business, suffered a technical outage, or purged the record in accordance with their data retention policies. The traceability chain does not fail silently or break the acquirer's SCRM tool; it deliberately illuminates this structural break. The acquirer can factor this trust boundary into their analysis, perhaps assigning a higher risk score to the component and requiring independent testing before integration.

Aligning with this principle helps ensure that supply chain risk management becomes more transparent and defensible. By carrying these distinct risk signals forward, a known data gap, whether declared by a participant or identified as a broken link, becomes highly valuable intelligence. It allows the acquirer to decide whether to accept the uncertainty, quarantine the affected units, trigger a supplier audit, or implement compensating controls. Conversely, approaches that hide gaps or force a binary "compliant/non-compliant" status foster false confidence and brittle supply chains.

3.6. Principle 6. Lifecycle-Oriented Provenance and Accountability

A manufactured product is rarely static; it undergoes modifications, firmware updates, repackaging, and component substitutions as it moves through distributors and integrators and ultimately into the hands of the acquirer. Traceability should extend across the product's lifecycle to reflect these updates and operational events. This helps ensure that the traceability chain remains a living timeline of the product's actual composition, allowing acquirers to accurately map external threat intelligence against the current state of their assets.

> **Notional Example: Firmware Updates and Continuous SCRM.** Consider an industrial controller that sits in a distributor's inventory for six months before being acquired by a regional power grid operator. During that time, the distributor applies a critical firmware update. The traceability framework is used to capture this "Update" event and appends the new component inventories to the controller's existing traceability chain. Months later, when a zero-day vulnerability is announced on the National Vulnerability Database (NVD), the grid operator's SCRM system queries the controller's traceability chain. Because the chain extends beyond the original factory state to reflect the updated lifecycle composition, the operator can accurately assess their exposure to the vulnerability from NVD without physically inspecting the deployed hardware.

> **Notional Example: Material Repackaging and Recall Response.** Consider a bulk chemical preservative manufactured, shipped to a regional distribution hub, and split and repackaged into smaller drums for delivery to various industrial facilities. The traceability chain records the consolidation and splitting events, linking the smaller containers back to the original bulk lot. If the original manufacturer later issues a recall due to contamination, the acquirers' facilities managers can use the traceability data to cross-reference the recalled lot number against their active inventory. Because the chain tracks the physical lifecycle of the material, not just the initial point of origin, they can immediately identify which specific sub-batches contain the affected preservative and intercept them before use.

Aligning with this principle shifts traceability from a one-time compliance check into a continuous accountability mechanism. Traceability records should be updated to reflect significant changes in the physical or digital makeup of the acquired item. By maintaining a continuous, lifecycle-oriented record, organizations help ensure that their SCRM tools are querying an accurate representation of the product. This provenance data can also remain actionable for vulnerability response, product recalls, and incident investigations long after the initial procurement has concluded.

3.7. Principle 7. Verifiable Organization Attribution

Identity is the foundation of accountability in any supply chain network. Acquirers should be able to definitively associate a specific supply chain event record with the organization that authored it. To support this, traceability records should include a positive, structural declaration of the accountable organization utilizing standard, verifiable identifiers (e.g., Global Location Number (GLN), Legal Entity Identifier (LEI), Data Universal Numbering System (DUNS) number, or Unique Entity Identifier (UEI)). Furthermore, to support effective risk management, this attribution should extend beyond the enterprise level to identify the specific organizational unit or facility where the event occurred, while intentionally avoiding the targeting of individual employees.

> **Notional Example: Enterprise-Level Attribution.** Consider a commercial software vendor releasing a new application. The vendor publishes a traceability event containing component inventories. Within the core structure of that record, the vendor includes their globally registered UEI. This simple, positive association enables downstream acquirers' automated systems to attribute software composition claims correctly to the specific corporate entity responsible for the product, distinguishing it from open-source forks or unauthorized redistributions.

> **Notional Example: Granular Unit-Level Attribution.** Consider a global electronics manufacturer that produces the same router model at three facilities. A traceability record that only identifies the parent enterprise is insufficient for supply chain risk management. Instead, the traceability record is structurally designed to capture both the parent organization identifier and a specific organizational unit identifier (e.g., "Enterprise A, Factory 3"). If sanctions are later imposed on the region where Factory 3 is located, or if a localized quality defect is discovered there, the acquirer can use these identifiers to precisely isolate the affected routers in their inventory without unnecessarily quarantining the routers produced at Factories 1 and 2.

Aligning with this principle helps ensure that every traceability record serves as a definitive statement of accountability. By embedding standardized organizational and unit-level identifiers directly into the traceability data model, traceability frameworks enable acquirers to map supply chain events to the specific corporate entities and physical locations responsible for them. This creates a highly functional, scalable foundation for supporting decentralized trust without overcomplicating the core data structure with complex cryptographic mandates.

3.8. Principle 8. Cyber-Physical and Digital Linkage

The integrity of any traceability system relies on the strength of the connection between the provenance record and the actual item it represents. If a legitimate digital provenance record can be easily decoupled and reassigned to a counterfeit physical item or an altered software

binary, the traceability data loses its value for SCRM. Traceability efforts should ensure that records can be definitively and verifiably associated with the specific instance of a physical product, digital asset, or batch they describe. To maintain this anchor of trust, a persistent identifier should be reliably associated with the tracked item. For physical products, this involves leveraging markers like barcodes, QR codes, or Physically Unclonable Functions (PUFs); for purely digital products, this linkage should rely on cryptographic hashes or digital signatures to securely bind the digital asset to its historical record.

Notional Example: Discrete Component Anti-Counterfeiting. Consider a defense contractor acquiring specialized networking switches. A malicious actor in the distribution chain might attempt to intercept a legitimate shipment, steal the authentic switches, and replace them with visually identical counterfeits while passing along the authentic digital traceability records. By laser-etching a unique and verifiable 2D data matrix directly onto the chassis of each instance of an authentic switch, the manufacturer establishes a secure anchor. When the acquirer receives the physical hardware, they scan the matrix, which verifiably bridges to the corresponding digital "Ship" and "Receive" records. Because the physical anchor cannot be easily removed or cloned onto a counterfeit chassis, the acquirer gains greater assurance to support their SCRM trust decisions.

Notional Example: Batch Integrity in Continuous Process Sourcing. Consider a pharmaceutical manufacturer acquiring bulk active pharmaceutical ingredients shipped in large industrial drums. Because liquids cannot be individually tagged, the cyber-physical linkage relies on securing the container. The supplier applies a tamper-evident Radio Frequency Identification (RFID) seal to the drum, linking the seal's unique digital signature to the specific digital batch record in the traceability chain. Upon arrival, the acquirer scans the physical seal. If the digital record matches but the physical seal shows signs of tampering, or the RFID signature fails validation, the acquirer gains actionable evidence that the cyber-physical link has been compromised, indicating that the physical contents may have been adulterated or substituted in transit.

Aligning with this principle helps ensure that the physical product and digital assets moving through the supply chain are the true subjects of the shared provenance data. The rigor of this linkage should be scaled with the item's operational risk (as outlined in Principle 2). While a simple adhesive QR code or standard file checksum may suffice for commercial office supplies, high-risk operational technology may require more robust, tamper-evident physical identifiers or cryptographically signed binaries. Ultimately, establishing this firm anchor deters and detects the substitution of unverified products or altered software masquerading as legitimate digital records, allowing acquirers to gain greater assurance in the authenticity of what they are procuring.

3.9. Principle 9. Interoperability Over Uniformity

Global supply chains are highly heterogeneous, spanning different industries, industry sectors, domains, geographic locations, and levels of technological maturity. Attempting to enforce a single, uniform data model, centralized platform, or proprietary standard across all participants is practically impossible and may lead to vendor lock-in. Instead, traceability efforts should prioritize interoperability. A principled approach relies on standardized traceability links to securely connect independent supply chain events (traceability records), while allowing the associated pedigree data within those records to be expressed in industry-specific formats.

Notional Example: Encapsulating Disparate Data Standards. Consider a defense prime contractor integrating Commercial Off-The-Shelf (COTS) networking equipment into a military communications system. The COTS supplier formats their traceability data using the Global Standards One (GS1) Electronic Product Code Information Services (EPCIS) standard, which is heavily adopted in retail logistics. In contrast, the defense contractor operates on a specialized defense-sector data exchange standard. Forcing the commercial supplier to completely retool its IT infrastructure would be a massive barrier to entry. Instead, the supplier can encapsulate their GS1 EPCIS data within a supply chain event record and explicitly declare this data type as part of the payload format. When the acquiring defense contractor's system receives it, it reads the format identifier, applies the appropriate parsing rules for GS1 EPCIS data, and seamlessly integrates the information.

Notional Example: Navigating Heterogeneous Storage and Ensuring Integrity. Consider a medical device manufacturer building a surgical robotic arm. A specialized microchip supplier tracks its provenance on a decentralized, distributed ledger (blockchain); the device manufacturer tracks assembly in a traditional, centralized relational database; and a lower-tier packaging supplier hosts secure flat files. Because the traceability framework relies on resolvable linkage pointers, the necessity for a consistent storage requirement is eliminated, allowing the acquiring hospital to navigate the full chain. These traceability links also incorporate cryptographic references (such as hashes) to prior records, ensuring that the sequence of records maintains data integrity. The responsibility for proving integrity shifts from the underlying storage repositories to the records. Consequently, the acquirer can independently validate that a predecessor record has not been altered since the successor linked to it, regardless of where the file is hosted.

Aligning with this principle helps ensure the overarching traceability approach scales globally. Acquirers and suppliers retain the flexibility to choose the technical architecture, storage mechanisms, and internal data formats that best align with their specific business models and regulatory environments. When sharing pedigree information, entities can abstract the necessary data from these internal systems to publish traceability records. By leveraging the

data encapsulation pattern, traceability records can accommodate different data payloads without mandating uniformity across every repository, while still maintaining standardized, interoperable links to predecessor events. By using interoperable traceability records and links, organizations can support the formation of continuous, end-to-end provenance chains that cross the natural boundaries of supply chain networks.

3.10. Synthesis of Principles

The nine principles outlined in this section collectively establish the foundation for a resilient, distributed, and decentralized architecture that reflects practical business realities. By acknowledging the complex, multi-ecosystem nature of modern manufacturing, these principles support global traceability without requiring a universally centralized data repository. Aligning with these guiding principles requires a consistent, structured approach that uses foundational event templates to represent key supply chain events.

These templates serve as a starting point for encapsulating provenance data, ensuring systemic interoperability, and enabling discrete records to link and form traceability chains that accurately mirror the timeline of supply chain events. Using these consistent data templates helps avoid dictating the internal contents of every supply chain payload. Ultimately, industry consortia, standards organizations, and manufacturers themselves bear the responsibility for deciding which specific pedigree information, data standards, and formats best serve their industries, industry sectors, or domains, and for tailoring the templates accordingly into specific traceability records.

To bridge theory and practice, Section 4 introduces the conceptual Supply Chain Traceability Meta-Framework, which outlines a set of high-level event data templates (e.g., Make, Assemble, Store, Ship, Receive, and Employ). The Meta-Framework contains the structural fields to support data encapsulation, interoperable traceability links, and independent integrity verification. Operating without a prescriptive centralized technology stack, the framework is structured to allow manufacturers across supply chain sectors to independently generate traceability records within their respective data repositories and link them across distributed boundaries to form the verifiable, end-to-end traceability chains that acquirers need to inform SCRM analysis.

4. A Manufacturing Meta-Framework

Within the Meta-Framework, the guiding principles established in Sec. 3 are translated into a data architecture. This architecture consists of high-level supply chain event traceability data record templates that organizations can use as blueprints to encapsulate pedigree data and link across records to support product provenance

While the successful adoption of traceability within the global supply chain networks relies heavily on non-technical drivers (e.g., industry governance, economic incentives, and proportional risk decisions (Principles 1–3)), the Meta-Framework is designed to operationalize technical requirements to inform standards, supply chain automation, and industry best practices. The framework is structured to account for the complex operational realities of manufacturing (Principles 4–6), allowing organizations to establish evidence-based provenance that protects intellectual property, account for inevitable traceability gaps, and support continuous updates to the product lifecycle. To achieve this, structural mechanics are detailed within the framework to support verifiable organization attribution (Principle 7), cyber-physical linkage (Principle 8), and consistent interoperability (Principle 9).

By using a consistent, interoperable approach to recording, linking, and retrieving supply chain event data, discrete records can be linked across organizational boundaries. This results in a causally and chronologically ordered chain of traceability records that are cryptographically verifiable and allow stakeholders to validate product provenance and make informed SCRM decisions.

What is a "Meta" Framework?

The prefix “meta” indicates that the framework is designed as a high-level conceptual structure to support supply chain pedigree and provenance information without requiring a specific technology or data standard. As established in Principle 9 (Interoperability Over Uniformity), attempting to force every global supplier to adopt a single, uniform data standard is impractical. Instead of replacing existing, highly functional industry standards (such as GS1 EPCIS in retail or specialized defense schemas), common structural data patterns, specifically linking, interoperable interfaces, and encapsulation of payload data expressed in data standards are used to support the Meta-Framework.

These interoperable interfaces serve as communication pathways, allowing downstream acquirers to navigate cryptographic links and retrieve relevant records seamlessly during traceback analysis. By using these patterns, independent, federated ecosystems can link data across trust boundaries.

> Common data architecture patterns are used for the Meta-Framework to build a continuous, causal, and temporally ordered chain of pedigree information that spans industry sectors and geographic locations, while incorporating the necessary flexibility in pedigree data to align with industry standards and manufacturing tailoring.

To validate the conceptual architecture of this Meta-Framework, NIST and MITRE jointly developed an open-source Reference Implementation (RI). This implementation serves as a minimum viable product to demonstrate and investigate supply chain traceability across multiple independently governed ecosystems. A high-level summary of the reference implementation's architecture and goals is provided in Appendix F.

4.1. Core Architectural Concepts

A decentralized, interoperable approach is used within the Meta-Framework across the supply chain network to support verifiable, accessible, and traceable records, while mechanisms are in place to protect the intellectual property and operational security of participating suppliers. This architectural model is built upon four core components, as shown in Table 4 below:

Table 4. Core Architectural Concepts

Component	Description
Structured Traceability Records	The foundational components of the framework are structured traceability records containing essential information about supply chain events at specific stages. Using these records, organizations can extract and summarize a subset of their internal manufacturing data to share as product pedigree. This approach facilitates the capture of critical data needed for verifiable visibility without exposing sensitive, proprietary manufacturing details.
Federated Data Repositories and Ecosystems	A flexible storage model in which traceability records are managed in data repositories hosted by individual manufacturers, industry coalitions, or regional consortia (e.g., a semiconductor fabrication ecosystem or a mining industry collective). This federated approach accommodates various data and technology implementation models: while larger organizations may host their own infrastructure, smaller manufacturers might utilize shared industry platforms or third-party service providers to maintain their records. Through these arrangements, participants can benefit from economies of scale and professionalized data management without needing to sustain independent infrastructure. Regardless of the hosting model, trust is decoupled from the underlying repository and is cryptographically embedded in the records themselves through secure, verifiable links to predecessor events. Additionally, these repositories are managed under governance frameworks, defined by the individual manufacturer or the collective ecosystem, in which the access controls and data retention policies necessary to maintain data accessibility for authorized stakeholders are established.
Traceability Links	Traceability links are utilized within the framework to connect individual records, creating a continuous chain of traceability across complex networks. These verifiable links enable stakeholders to trace the lineage of a product or component over time and across data repositories managed by different organizations.

Component	Description
Common Data Models and Ontologies	Aligning traceability record internal payloads with established standards organizations helps maintain syntactic and semantic consistency across intersecting supply chains. Through this alignment, traceability data remains interoperable, comprehensible, and actionable across supply chain networks.

4.1.1. Supply Chain Event Categories

Complex operational milestones are generalized within the Meta-Framework into supply chain event categories (Principle 6). These initial categories provide a common language for documenting a product's progression throughout its lifecycle. Table 5 describes an initial set of event categories for supporting a manufacturing meta-framework from initial “make” events to the “employ” event, deploying a product within an acquiring enterprise environment:

Table 5. Supply Chain Event Categories

Event Category	Description
Make	The initial creation or fabrication of a product or component.
Assemble	The integration of multiple tracked components into a new, more complex system.
Store	The staging of materials, including physical location and, where relevant, environmental conditions.
Ship	The outbound transfer of materials to another entity.
Receive	The inbound receipt of an item from another entity.
Employ	The deployment or activation of the product within its final operational environment.

Supply chain events do not occur in isolation; they represent specific milestones in a product's physical lifecycle. While event records are captured chronologically as products move through the supply chain, mechanisms are included within the framework for the secure verification of this history using backward-pointing cryptographic linkages. For example, a verifiable thread across a trust boundary is created when a Receive record is linked to the supplier's predecessor Ship record. These structural data relationships are shown in Figure 5, where Unified Modeling Language (UML) cardinality indicates how each successor event category is connected to its logical predecessors, forming a continuous provenance chain.

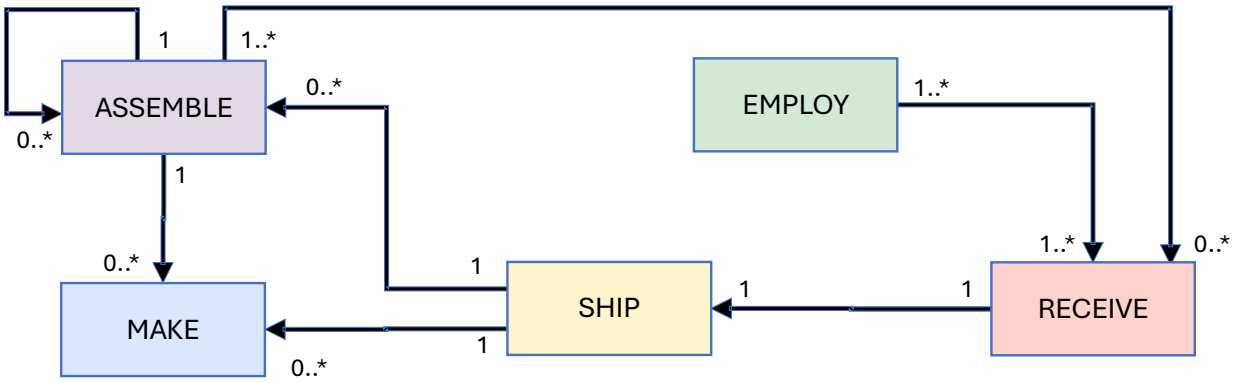

Fig. 5. Supply Chain Event Relationships

Note: The Store event represents an internal intermediary state and is omitted from the main diagram paths to preserve readability. Architecturally, a Store event must trace back to a prior internal event that established custody or creation (Receive, Make, Assemble, or a prior Store). Conversely, successor events (Ship, Assemble, Employ, or a subsequent Store) may trace back to a Store event. To preserve trust boundaries, a Store event should not trace directly back to an external Ship event without a preceding Receive.

4.1.2. From Internal Operations to Shareable Provenance

Differentiating between an organization's internal operational data and shareable supply chain event data written to a traceability record clarifies how the core components interact in practice. This abstraction is illustrated in Fig. 6 using a notional use case: the ACME OEM Security Camera Manufacturing process. Within the facility, the manufacturer utilizes highly complex, proprietary systems (e.g., ERP, MES, and Quality Management Systems (QMS)) to manage the intricate steps involved in assembling the camera's subsystems. This granular internal data is typically not shared with external parties.

Instead, organizations can selectively extract and transform these internal activities into shareable supply chain event data, filtering out proprietary manufacturing details to protect intellectual property. The supply chain event information encapsulated inside these individual traceability records (e.g., Make, Assemble, Store, Ship, Receive, and Employ records) represents pedigree information. As the traceability records are cryptographically linked together, across organizational and geographic locations, they form the product's verifiable provenance.

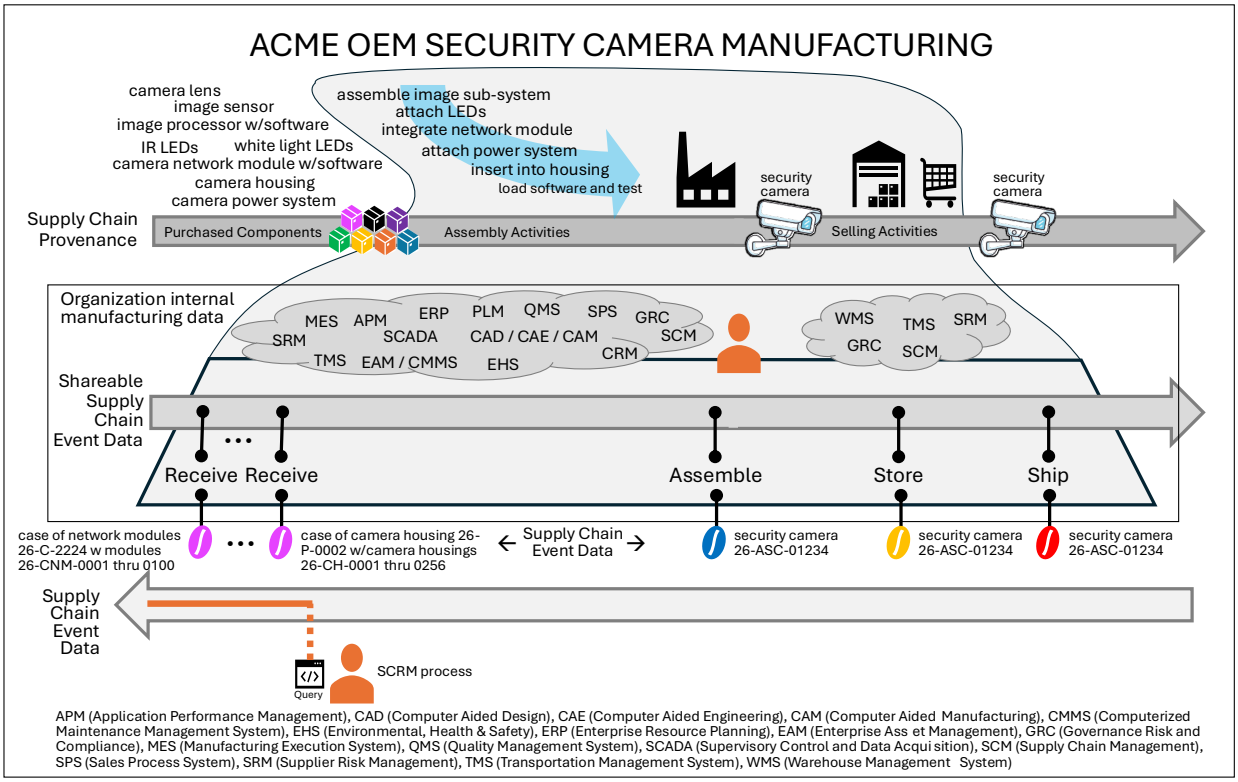

Fig. 6. Abstracting Internal Manufacturing Data into Shareable Supply Chain Events

For example, once an organization actively curates and transforms its internal operational data into traceability records, these records are made available within the supply chain network for downstream entities to link.

This linkage occurs through the exchange of traceability links between the acquirer and the supplier manufacturer, which serve as verifiable pointers to specific records. When a supplier provides a traceability link (for example, pointing to a "Ship" record) to their customer, the customer can embed that link in their corresponding record (such as a "Receive" record). This process creates a verifiable connection between the two organizations without requiring either party to grant access to their internal manufacturing systems.

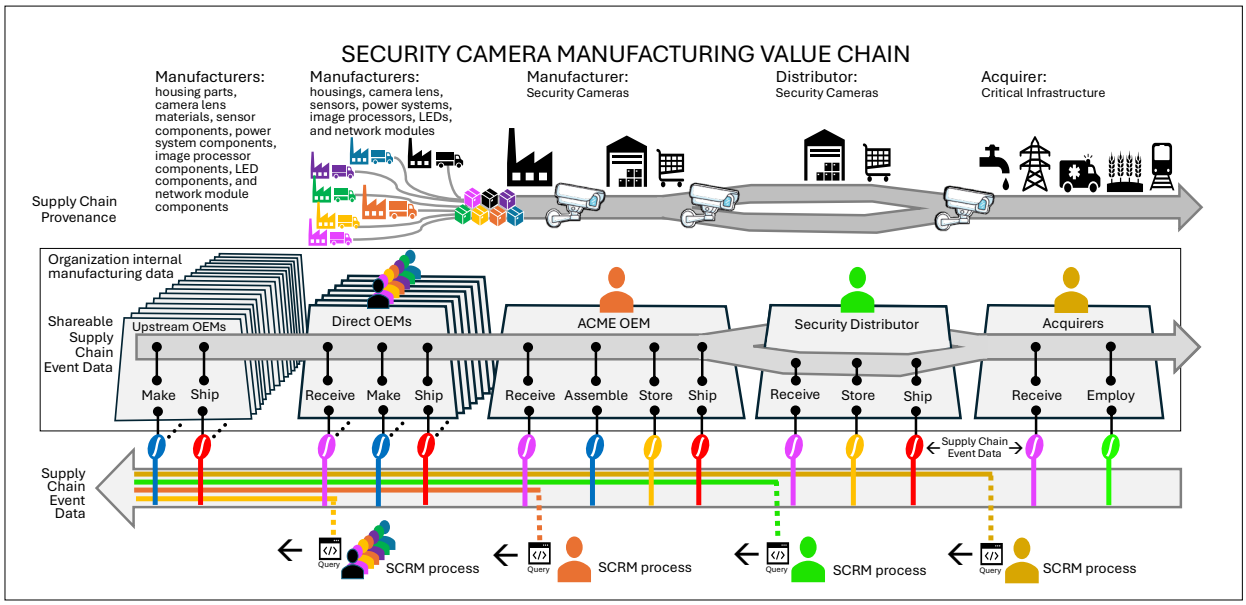

Fig. 7. End-to-End Traceability Chain Across the Manufacturing Value Chain

As shown in Fig. 7, the traceability chain begins with upstream component manufacturers generating "Make" and "Ship" records for items like lenses and sensors. The records include a unique traceability link for that specific component instance or batch. For example, when ACME receives these items, they generate a "Receive" record that references and cryptographically links back to the supplier's "Ship" record using the provided pointer.

This linking process continues through the ACME assembly process, down through security distributors, and ultimately to the acquirer. The traceability link for the current supply chain event traceability record is passed forward at each supply chain event to the acquirer, such that the traceability record for the successor supply chain event can link backward to the traceability record of the predecessor supply chain event. The final traceability chain enables an acquirer to query it and independently trace the product's provenance back to its origins, even across disparate data repositories.

4.2. Templated Traceability Record Structure (Metadata and Pedigree)

Traceability records are the logical foundation of the Meta-Framework. Each record contains the details of a discrete supply chain event and includes the structural consistency needed for data to be shared and interpreted across supply chain networks. With the encapsulation and linking design patterns mentioned earlier, a rigid, universal schema is not required within the framework. This approach is based on a consistent, templated container for pedigree information for each type of supply chain event. Consequently, the metadata required for navigation via links remains interoperable, while the internal data remains flexible and tailorable.

4.2.1. Core Traceability Record Elements and Principle Alignment

To support consistency across the supply chain network while preserving flexibility for industry-specific tailoring, a set of core data elements is defined. Table 6 lists these core components and details the application of the fundamental principles defined in Section 3 to establish a consistent structure for encapsulating pedigree information and linking records.

Table 6. Core Traceability Record Elements

Element Category	Description and Architectural Purpose	Principle Alignment
Record Identifier	A unique identifier for the traceability record, establishing a distinct reference for the specific event entry within an ecosystem.	**Principles 5 & 9:** The precise referencing mechanism needed to link across ecosystems and locate missing nodes (gaps).
Event Category	Indicates the supply chain event lifecycle milestone (e.g., Make, Assemble, Store, Ship, Receive, Employ) contained within the record.	**Principle 6:** Represent the item's operational status and progress throughout its lifecycle.
Event Timestamps	Precise chronological data representing the distinction between when the event physically occurred (Occurrence) and when it was logged (Recorded).	**Principles 5 & 6:** Definition of the product lifecycle timeline, facilitating the detection of temporal anomalies or data gaps by acquirers.
Organization Identifier & Subunit [2]	Represents the legally identifiable business entity responsible for the event. May include a subunit identifier for facility-level granularity, without capturing individual employee data.	**Principle 7:** Facilitates the definitive verification of the provenance claim's author by an acquirer.
Tracked Entity Identifier	A persistent, unique identifier associated with the specific physical component, digital object, or batch affected by the event.	**Principle 8:** The primary digital anchor between the record and the item.
Traceability Link(s)	Contains verifiable pointers (e.g., query parameters, cryptographic hashes) referencing preceding event record(s) to support backward traversal	**Principles 8 & 9:** Support the formation of a continuous, tamper-evident chain.
Data Type Identifier	An industry- or regulatory-defined label indicating the specific industry schema, standard, or shared data ontology with which the internal payload is aligned.	**Principle 9:** Supports syntactic and semantic consistency across intersecting supply chains.

[2] The Meta-Framework is designed to track organizational accountability, not individuals.

Element Category	Description and Architectural Purpose	Principle Alignment
Tracked Entity Data (Payload)	A flexible data block containing detailed, event-specific pedigree information (e.g., chemical or material composition, serial numbers, batch numbers, or product grade). Alignment of this internal data structure with established common data models and ontologies facilitates universal comprehensibility of the specific pedigree across downstream systems.	**Principle 9:** Accommodation of the operational nuances of different sectors without the imposition of a rigid, universal schema, utilizing shared ontologies to support semantic interoperability.
Supplemental Data References	Optional pointers to external validation sources (e.g., certifications, test reports) are maintained in gated or paywalled stakeholder systems separate from the immutable chain.	**Principle 4:** Support for data minimization and the protection of intellectual property through the separation of sensitive evidence from routing data.

4.2.2. Data Type Identifier and Tracked Entity Data (Payload)

To maintain structural interoperability across supply chains, an encapsulation strategy is utilized within the Meta-Framework to separate the mechanics of traceability (the “envelope”) from the highly specific, domain-relevant pedigree data (the “letter”). Routing, cryptographic linking, and event sequencing are handled externally, while the industry-specific details are secured inside the Tracked Entity Data Payload.

Because the payload is flexible, a machine-readable method is required for parsing—a function fulfilled by the Data Type Identifier. Similar to how structured markup data requires paired processing templates or stylesheets to establish how information is interpreted and rendered, the Data Type Identifier acts as the translation key for payload information. It specifies the schema or ontology required by the receiving system to parse the internal payload. If an acquirer only needs to verify structural provenance, the outer traceability links can be authenticated without requiring the specific methods needed to decode a specialized internal payload.

While this decoupled approach yields a flexible environment, it also necessitates industry-level decisions about how to balance encapsulating legacy standards with evolving them for modern traceability. In the short term, manufacturers can use this structure to capture existing, unmodified standards-based records into the payload, enabling ecosystems to establish cryptographically verifiable traceability without requiring organizations to abandon current data formats. Over time, as supply chain networks mature, industry standards organizations can evolve and update their specifications to incorporate payload information and potentially support key Meta-Framework parameters natively, such as cryptographic traceability links and verifiable organization/entity identifiers.

4.2.3. Supplemental Data References

While traceability links constitute the immutable backbone of product traceability record provenance lineage, organizations often require additional evidence to conduct comprehensive SCRM. To accommodate this, references to supplemental data can be incorporated into records within the Meta-Framework (Principle 4).

Unlike traceability links, which connect supply chain event records to form the provenance chain, supplemental data references designate external validation sources. These sources may include test reports, quality inspection results, audit summaries, or compliance attestations. The separation of the "predecessor traceability record link" from the "supplemental evidence" facilitates data minimization and the protection of intellectual property.

The traceability link is utilized to support the routing infrastructure, whereas the supplemental data contains additional, possibly sensitive evidence. Because these external references may contain proprietary business information, they might not be universally accessible. Instead, they can be stored in access-restricted systems where suppliers retain strict control over who can view detailed evidence. This architecture prevents the exposure of trade secrets while maintaining the accessibility of shareable pedigree information.

4.3. Traceability Links and Traceability Chains (Provenance Chronology Information)

While individual traceability records capture localized supply chain events, their primary value lies in their connection to predecessor records. Traceability links are used within the Meta-Framework to connect discrete records into a continuous, verifiable sequence known as a traceability chain. These links enable stakeholders to follow the lineage of a product or component over time, even when that data is distributed across different organizations and repositories.

4.3.1. Integrity and Trust

The Meta-Framework is designed to function without a central authority, establishing a basis of trust within and across individual records so that trust requirements are not imposed on the data repository. This decentralized trust model is categorized into two functional areas:

- **Record Integrity:** This property verifies that the record has not been altered since its creation. To achieve this, a cryptographic hash representing the predecessor record is embedded within the traceability link. This structure yields a tamper-evident timeline; if any historical record in the chain is modified, the cryptographic hashes in all subsequent records will no longer match, allowing the acquirer to detect an integrity failure.
- **Data Attribution:** This helps ensure the record accurately references the tracked item and is authored by the claimed organization. This is supported by the Verifiable Organization Identifiers (Principle 7) and Tracked Entity Identifiers (Principle 8) discussed in previous sections.

By directly incorporating these markers into the record structure, a basis for trust in the data is established regardless of the repository or ecosystem from which the record is retrieved.

4.3.2. The Traceback Mechanism and Unidirectional Links

Traceability chains are utilized by acquiring and other organizations to conduct traceback investigations to support SCRM or related risk analysis. Through this backward-linking logic, a specific record can be retrieved, followed by the tracing of the event chain. To support verifiable component validation while protecting operational privacy and intellectual property, these links are unidirectional. By referencing only predecessor records, backward tracing is achieved without creating a "track-forward" capability that might expose a supplier's downstream supply chain.

This unidirectional logic mirrors the physical flow of products and creates a predictable causal order for traceability records:

- An Employ record points to the Assemble, Store, or Receive event that preceded its deployment.
- An Assemble record points to the Receive, Store, Assemble, or Make events of its underlying components.
- A Receive record points to the corresponding Ship record from the upstream supplier.
- A Ship record points to the Make, Assemble, or Store record that created or housed the product.

4.4. Federated Traceability Ecosystems (Assurance)

A single centralized database is not required within the Meta-Framework; instead, organizations can implement it across decentralized, federated data repositories. Because of this distributed architecture, participants must prioritize governance and access management. Ecosystem stakeholders establish the governance policies and rules that organizations use to regulate how traceability records are created, stored, and accessed. By establishing localized rules tailored to specific industries, organizations within these ecosystems can protect sensitive supply chain data while maintaining interoperability with the broader global supply network.

4.4.1. Governance

Each participating ecosystem is governed by its own framework, wherein membership criteria, compliance obligations, and participant responsibilities are established so that traceability principles are consistently applied across the network. Through the delegation of ecosystem-level governance, highly flexible implementation models are facilitated in alignment with Principle 3 (Incentive-Aligned Participation).

> **Notional Example**: A large prime contractor might establish an ecosystem with specific governance rules for its immediate subcontractors, while a group of

small- to medium-sized enterprises might rely on a shared industry consortium or service provider to govern and host their records. Through this federated structure, a major operational concern for lower-tier suppliers is mitigated: the fatigue of reporting across multiple customer portals. For example, in traditional supply chain models, a component manufacturer might be required to log in to Prime Contractor A's proprietary system to record the production and shipment of Batch A, and then log in to Prime Contractor B's entirely separate system to record Batch B. Through the Meta-Framework, that same supplier can generate and maintain all their traceability records within a single ecosystem or shared service repository of their choosing. When shipping Batch A, the supplier provides Prime A with a traceability link corresponding to that specific record. When shipping Batch B, they provide Prime B with a distinct traceability link for that batch.

Because interoperable interfaces are utilized, downstream acquirers use those links to retrieve the records directly from the supplier's chosen repository. Through this mechanism, small and medium manufacturers can provide pedigree information and actively contribute to the provenance chain.

4.4.2. Controlled Access and Data Retention

To support the security and availability of supply chain event data, ecosystem data repositories should implement technical and procedural safeguards to ensure traceability records remain protected, verifiable, and accessible throughout their lifecycle. Key ecosystem capabilities include:

- **Authentication Mechanisms:** Systems that verify the identity of stakeholders when accessing or submitting traceability records.
- **Access Control Mechanisms:** Policies that define and enforce who can read, write, or manage traceability records, protecting supply chain data from unauthorized use.
- **Data Retention Policies:** Rules that specify how long traceability records must be stored to meet operational, contractual, and compliance obligations, including contingency plans for repository dissolution or business continuity to maintain long-term data availability.

Together, these practices help maintain the integrity, availability, and long-term reliability of traceability records.

4.4.3. Link-Based Querying and Interoperability

A fundamental challenge in supply chain visibility is crossing trust boundaries. An acquiring enterprise may have an established relationship with its prime contractor, but it likely does not have direct accounts or authentication credentials configured for the data repositories of Tier-3 or Tier-4 suppliers.

To overcome this and support Principle 9 (Interoperability Over Uniformity), access to traceability records needs to be intelligently managed. Traceability Link-Based Querying is supported by the Meta-Framework, enabling users to retrieve data without direct authentication to upstream ecosystems. A traceability link is a data object that contains three critical pieces of information:

- **Location:** Where to obtain the predecessor traceability record.
- **Query Parameters:** The specific parameters required to retrieve the requested record.
- **Cryptographic Hash:** The expected hash value of the returned data, which helps ensure data integrity and allows the requester to independently validate that the information has not been altered since it was originally recorded.

When a downstream acquirer possesses this link, it acts as a targeted retrieval mechanism. This allows the acquirer to follow the traceability chain across ecosystem borders, retrieving only the specific historical records necessary to validate product provenance, without requiring overarching access to the upstream supplier's entire data repository. **This targeted, link-based retrieval allows acquirers to obtain a continuous, verifiable traceability chain across complex, decentralized supply networks.**

4.5. Operational Execution: Enabling Provenance Query and SCRM

To understand how these components interact conceptually in practice, we examine a generalized operational scenario. An acquirer utilizing the Meta-Framework's decentralized architecture to perform SCRM validation is illustrated in Fig. 8.

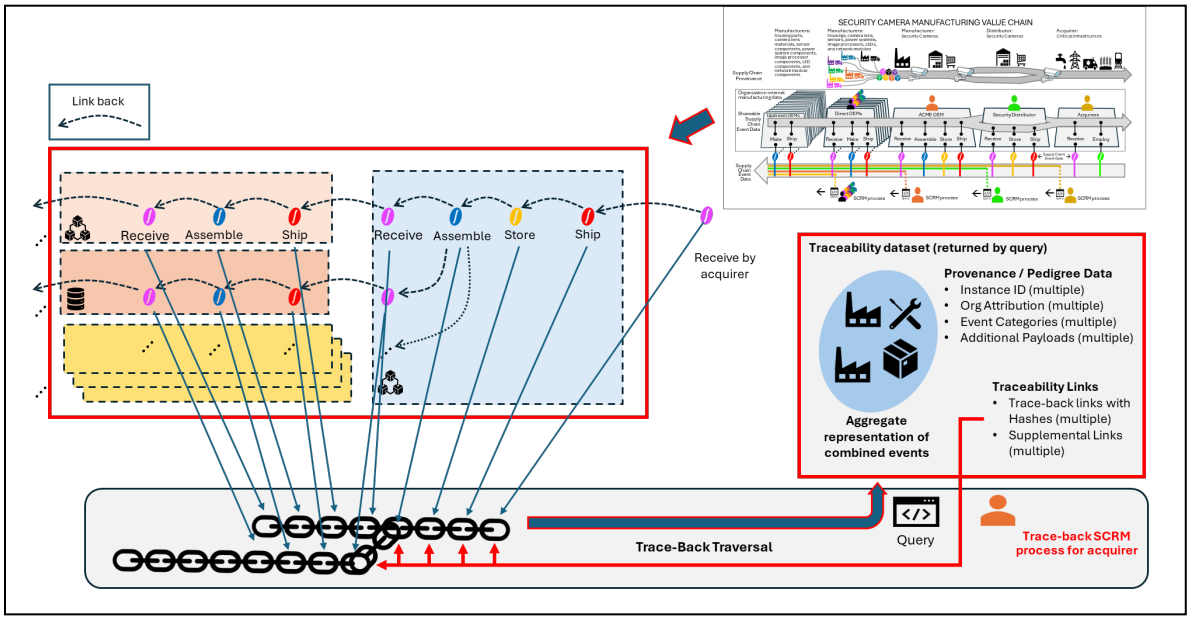

Fig. 8. Traceability Chain Spanning Multiple Supply Chain Ecosystems

In this scenario, an Acquiring Organization (acquirer) is deploying a critical asset (e.g., a security camera) and initiates a Provenance Query for risk analysis.

- **Entry Point:** The query begins digitally, starting from the acquirer's initial "Receive" or "Employ" record.
- **Walk the Chain:** From that entry point, the user's SCRM tool automatically follows the cryptographically verifiable traceability links (the "link back" arrows) backward through the chronological timeline. **Because the Meta-Framework is fundamentally designed to connect successor events to predecessor events, this link-based query is intentionally optimized for traceback rather than forward tracking.** The tool retrieves these historical records across multiple distributed repositories and ecosystem boundaries.
- **Validate the Traceability Record:** Upon retrieving the records, the acquirer can perform integrity verification:
    - Compare and validate the Cryptographic Integrity Hash of the record with the record itself; note that validating the record also validates the integrity of the contained links pointing back to the preceding records.
    - Use the records' Digital Anchors to (a) confirm the Cyber-Physical Linkage to the specific item and to (b) establish the Organization Attribution for who authored the record.
    - If necessary, use the Supplemental Links (optional references) to retrieve gated evidence (e.g., component inventories or test reports) for deeper technical analysis.

Once this integrity verification is complete, the acquirer can "Accept" that specific traceability record and continue tracing the chain back to the originating "Make" traceability record of the most critical sub-components. By verifying each linked traceability record in this manner, the acquirer builds a verifiable Provenance Timeline. This actionable intelligence is then fed directly back into the acquirer's SCRM process, enabling a defensible risk decision based on objective product evidence.

This operation can be performed at any supply chain step, including the "Receive" step in the acquirer enterprise. Any acquiring organization can use this traceability chain to assess the fitness-for-use of a product before integration or further manufacturing.

5. Conclusion

Traceability of products and components across the supply chain is essential to inform acquirer SCRM analysis, thereby verifying product integrity, building stakeholder trust, and supporting accountability throughout manufacturing ecosystems. However, collecting and verifying the associated product and component provenance data is a significant challenge, especially for complex, multi-tiered supply chains with fragmented and inconsistent data practices.

To address this challenge, the guiding principles established in Sec. 3 lay the foundation for a resilient, decentralized approach grounded in practical business realities. By translating these principles into a structural model, the Meta-Framework (Sec. 4) is designed to provide common data architecture patterns to build a continuous, ordered provenance chain spanning industry sectors and geographic locations. It achieves this while incorporating the necessary flexibility to keep the internal pedigree data aligned with existing industry standards and manufacturer-tailored specifications. By defining this structured, tailorable, and interoperable model for recording, linking, and retrieving supply chain event data, SCRM analysts and other stakeholders can:

- Retrieve and organize traceability records into ordered timelines of relevant supply chain events.
- Interpret retrieved provenance and pedigree information within each traceability record in its appropriate ecosystem and manufacturer-defined context.
- Use cryptographic proofs of data integrity and authenticity for each retrieved traceability record to validate product pedigree and provenance.

Traceability chains are formed by linking traceability records, from current to predecessor, created from supply chain events (e.g., Make, Assemble, Store, Ship, Receive, Employ) using cryptographically verifiable connections. These links allow stakeholders to construct a coherent timeline of supply chain events that reflect a product's movement and transformation across the manufacturing supply chain network. Trust is supported by these cryptographic validation mechanisms. Using hash-based traceability links, the record sequence is made tamper-evident and verifiable connections to predecessor records are established. This structure facilitates consistent validation of data integrity over time, even as traceability records are written to data repositories, where applicable, without relying on a single centralized authority.

Finally, this verifiability is supported within the Meta-Framework through data minimization and the selective disclosure of only the necessary information, promoting transparency without compromising operational security. Ecosystem data repositories are not intended to be centralized data lakes storing all manufacturing or organizational data; rather, they strictly house consistently templated and tailorable supply chain event traceability records. Organizations can extract and publish high-level pedigree data required for external provenance validation while keeping sensitive intellectual property, internal engineering models, and Personally Identifiable Information (PII) securely within their internal systems. When deeper validation is required, it can be facilitated through gated supplemental data references, allowing the supplier to retain ultimate control over their proprietary information.

When this curated data is shared, understanding and interoperability are further enhanced using flexible, ecosystem-specific data models and ontologies. The Meta-Framework's templated traceability records offer the flexibility to align with externally defined traceability requirements (e.g., from industry consortia or standards organizations) while remaining comprehensible and actionable across heterogeneous, cross-sector environments.

While a strong foundation for cross-ecosystem traceability is provided by this framework, several areas require further development. Future work should prioritize expanding interoperability models, refining data integrity validation methods, integrating advanced privacy-enhancing technologies, and adding new traceability event categories to reflect emerging operational needs. For additional discussion on future directions, see Appendix E.

Appendix A. References

[1] National Counterintelligence and Security Center (NCSC), NCSC Unveils the New National Counterintelligence Strategy. 1 Aug. 2024. https://www.dni.gov/files/NCSC/documents/features/FINAL-NCSC-National-CI-Strategy-Press-Release_2024.pdf Accessed 20 Aug. 2026.

[2] National Counterintelligence and Security Center (NCSC), Supply Chain Risk Management: Reducing Threats to Key U.S. Supply Chains. 25 Sept. 2020. https://www.dni.gov/files/NCSC/documents/supplychain/20200925-NCSC-Supply-Chain-Risk-Management-tri-fold.pdf Accessed 20 Aug. 2026.

[3] Boyens, Jon M, et al., Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations. May 2022, https://doi.org/10.6028/nist.sp.800-161r1-upd1. Accessed 20 Aug. 2026.

[4] Joint Task Force, Security and Privacy Controls for Information Systems and Organizations. September 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf. Accessed 20 Aug. 2026.

[5] The NIST Cybersecurity Framework (CSF) 2.0. February 2024, https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf. Accessed 20 Aug. 2026.

[6] Stouffer, Keith, et al., Blockchain and Related Technologies to Support Manufacturing Supply Chain Traceability: April 2022, https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8419.pdf. Accessed 20 Aug. 2026.

[7] Pease, Michael, et al., Project Description: Manufacturing Supply Chain Traceability with Blockchain Related Technology: Reference Implementation. August 2023, https://www.nccoe.nist.gov/sites/default/files/2023-08/mfg-sct-blkchn-project-description-final.pdf. Accessed 20 Aug. 2026.

[8] Bećirović Ramić, Š. C. (2024). Selective disclosure in digital credentials: A review. ICT Express, 10(4), 916-934. https://people.etf.unsa.ba/~smrdovic/publications/ICT_Express_Becirovic_Ramic_et_al.pdf Accessed 20 Aug. 2026.

[9] Birkholz, H. D.-L. (2025). An architecture for trustworthy and transparent digital supply chains (RFC 9943), Internet Engineering Task Force. https://www.rfc-editor.org/rfc/rfc9943.pdf Accessed 20 Aug. 2026.

[10] Object Management Group. (2025). System Package Data Exchange (SPDX) Specification Version 3.0. https://www.omg.org/spec/SPDX/3.0/. Accessed 20 Aug. 2026.

[11] Europe, U. N. (2026). United Nations Transparency Protocol (UNTP) specification, UNECE. https://untp.unece.org/docs/specification. Accessed 20 Aug. 2026.

[12] National Economic Council & National Security Council, 2021–2024 Quadrennial Supply Chain Review, December 2024. https://www.whitehouse.gov/wp-content/uploads/2024/12/20212024-Quadrennial-Supply-Chain-Review.pdf Accessed 20 Aug. 2026.

[13] NIST Privacy Framework, National Institute of Standards. Available: https://www.nist.gov/privacy-framework. Accessed 20 Aug. 2026.

[14] CISA, Supply Chain Risks for Information and Communication Technology (National Risk Management Center, December 2018). https://www.cisa.gov/sites/default/files/publications/19_0424_cisa_nrmc_supply-chain-risks-for-information-and-communication-technology.pdf. Accessed 20 Aug. 2026.

[15] U.S. Government Accountability Office. (2018). Information Security: Supply Chain Risks Affecting Federal Agencies (GAO-18-667T). https://www.gao.gov/assets/gao-18-667t.pdf. Accessed 20, Aug. 2026.

[16] Federal Acquisition Regulation (FAR) § 4.2102, Prohibition, implementing Section 889(a)(1)(A)–(B) of the John S. McCain National Defense Authorization Act for Fiscal Year 2019, Pub. L. No. 115-232. https://www.acquisition.gov/far/4.2102. Accessed 20 Aug. 2026.

Appendix B. List of Symbols, Abbreviations, and Acronyms

API
Application Programming Interface

CI

Critical Infrastructure

COTS
Commercial Off-The-Shelf

DUNS
Data Universal Numbering System

DOT
Department of Transportation

EPA
Environmental Protection Agency

EPCIS
Electronic Product Code Information Services

ICAM

Identity, Credential, and Access Management

IT
Information Technology

ITL
Information Technology Laboratory

LEI
Legal Entity Identifier

MES
Manufacturing Execution Systems

MRO
Maintenance, Repair, and Overhaul

MTRs
Material Test Reports

NCCoE
NIST National Cybersecurity Center of Excellence

NERC CIP
North American Electric Reliability Corporation Critical Infrastructure Protection

NIST
National Institute of Standards and Technology

NIST IR
National Institute of Standards and Technology Internal Report

NIST SP
National Institute of Standards and Technology Special Publication

NVD
National Vulnerability Database

OEM
Original Equipment Manufacturer

PII
Personally Identifiable Information

PUFs
Physically Unclonable Functions

PIAs
Privacy Impact Assessments

PRAM
Privacy Risk Assessment Methodology

QMS
Quality Management Systems

RFID
Radio Frequency Identification

SCRM
Supply Chain Risk Management

SMEs
Small- to Medium-sized Enterprises

SP
Special Publication

SPDX
System Package Data Exchange

UEI
Unique Entity Identifier

Appendix C. Glossary

Acquirer

An organization or entity that acquires or procures a product or service. [IEEE15288, adapted)

Critical Infrastructure (CI)

Systems and assets, whether physical or virtual, so vital to the United States that their incapacity or destruction would have a debilitating impact on security, national economic security, or national public health or safety. Note: Within this publication, this term is also used as a shorthand identifier for an organization, acquirer, environment, or ecosystem operating within that context, and is not synonymous with industry, industry sector, or domain.

Cyber-Physical Link

A unique identifier that digitally associates a traceability record with a physical or virtual item to ensure that the item can be tracked and verified throughout its lifecycle.

Data Repository

A data storage system or service designated within an ecosystem to store traceability records and governed by policies that control access, enforce data retention requirements, and support data retrieval mechanisms for acquirers.

Domain

An area of manufacturing, technology, or application that may span multiple industries or sectors (e.g., discrete manufacturing or software).

Event Type

A classification that describes the kind of supply chain activity being recorded, such as make, assemble, receive, or employ. Event types define the structure of the associated variable data block.

Federated Traceability Ecosystem

A bounded network of supply chain participants operating under a shared governance framework and data models and ontologies. Ecosystems serve as localized administrative domains that manage access controls, storage repositories, and payload formats for their specific industry sector, while using the Meta-Framework to ensure interoperability with external ecosystems.

Governance

A set of policies, rules, and enforcement mechanisms that are defined by an ecosystem to ensure the integrity, security, and proper management of traceability records and participant interactions.

Industry

A broad community of organizations engaged in related economic or production activities.

Industry Sector

A recognized grouping or segment of an industry. Note: Within this publication, this compound term is used broadly and does not imply a level of classification separate from "industry" or "sector."

Paywalling

Paywalls are a method of restricting access to content or features on a website or app, requiring users to pay or subscribe to access them. They generate revenue for content creators.

Pedigree

The validation of the composition and provenance of technologies, products, and services is referred to as the pedigree. For microelectronics, this includes the material composition of components. For software, this includes the composition of open-source and proprietary code, including the version of the component at a given point in time. Pedigrees increase the assurance that the claims suppliers assert about the internal composition and provenance of the products, services, and technologies they provide are valid. [NIST SP 800-161 Rev. 1]

Personally Identifiable Information

Information that can be used to distinguish or trace an individual's identity—such as name, social security number, biometric data records—either alone or when combined with other personal or identifying information that is linked or linkable to a specific individual (e.g., date and place of birth, mother's maiden name, etc.). [FIPS 201-3]

Privacy

Assurance that the confidentiality of, and access to, certain information about an entity is protected. [NIST SP 800-130]

Product

Any physical, digital, or cyber-physical item manufactured, assembled, processed, or transferred within a supply chain network. In this framework, "product" is an inclusive term encompassing raw materials, discrete physical components, complex hardware assemblies, and purely intangible digital assets (such as software, firmware, and data models) that require traceability.

Provenance

The chronology of the origin, development, ownership, location, and changes to a system or system component and associated data. It may also include personnel and processes used to interact with or modify the system, component, or associated data. [NIST SP 800-53 Rev. 5]

Sector

A grouping of related economic, operational, governmental, or regulatory activities, which may also reflect an established classification or context.

Supplemental Link

A non-mandatory data reference that connects a traceability record to external data sources (e.g., certifications, test reports, quality inspection results, operational logs, audit summaries, compliance attestations) stored outside the traceability repository. These are used to support additional validation or compliance requirements.

Supply Chain Network

The overarching, interconnected network of stakeholders (e.g., businesses, vendors, suppliers, partners, people, processes, and resources) that collectively work to physically create, distribute, and manage products. [Derived from https://www.gep.com/blog/strategy/supply-chain-ecosystem-how-to-build-competitive-advantage].

Traceability

The ability to follow the lineage, application, or location of what is under consideration. [ISO 21931-2:2019, adapted]

Traceability Chain

An ordered set of traceability records (causally and temporally linked), each representing a supply chain event, with provable integrity at each step.

Traceability Link

A reference mechanism within a traceability record that connects it to a prior record, enabling the reconstruction of the product's history. Links typically include record identifiers, query parameters, and cryptographic hashes.

Traceability Record

A structured data object that contains a specific supply chain event (e.g., make, ship, receive) and includes metadata, traceability links, and optionally supplemental references. These records support the construction of traceability chains.

Trust Boundary

A logical or physical perimeter where control, ownership, or security policies change. In a supply chain traceability context, a trust boundary is the point at which physical products or digital records transition between independent organizations, ecosystems, or distinct IT systems. Because systems on opposite sides of a trust boundary operate under different governance, data crossing this perimeter cannot be implicitly trusted and requires explicit cryptographic or procedural verification.

Variable Data Block

A flexible portion of a traceability record used to store industry- or event-specific metadata, defined according to schemas referenced by the Event Type.

Appendix D. Security, Privacy, and Access Control Considerations

The Meta-Framework is designed to illustrate one way to apply the foundational traceability principles to address specific cybersecurity and privacy considerations, including linking traceability records across repositories in federated ecosystems, enforcing authentication, and managing potentially sensitive personal or business information. This appendix highlights key conceptual considerations to support risk management for confidentiality, integrity, availability, and privacy across the supply chain network.

This appendix is not intended to serve as a comprehensive security or implementation guide. Individual ecosystems should define detailed technical deployment practices tailored to their specific architecture, operational requirements, and risk assessments. Instead, this appendix contains high-level security and privacy guidelines. Alignment with broader NIST guidelines, such as NIST SP 800-53 [4] and the NIST Privacy Framework [12], can further enhance the trustworthiness of traceability systems.

D.1. Identity, Authentication, and Access Control

Identity, Credential, and Access Management (ICAM) is essential to securing traceability records [4]. Ecosystems should implement mechanisms to authenticate stakeholders before allowing them to access or manage records and authorize data access strictly on a need-to-know basis.

Given the multi-organizational nature of most ecosystems, access should be carefully scoped to reflect participant roles. While external auditors or ecosystem coordinators may require broad query capabilities, competitors and unauthorized parties should be technologically restricted from accessing or inferring proprietary information.

To support decentralized access without the administrative burden of cross-organizational account management, ecosystems can also leverage traceability links to support authorized access. By using unguessable, cryptographically secure links or query parameters, an acquirer in possession of a product can retrieve its specific traceability record, containing pedigree information. Importantly, because these links cannot be sequentially guessed or iterated, this mechanism protects the supplier's repository from automated data mining or bulk extraction. This ensures that downstream entities can retrieve only the records for which they possess the requisite links, preserving the vendor's broader operational security and proprietary data without requiring formal, authenticated user accounts in every upstream repository.

D.2. Privacy Measures

Although traceability records primarily support product pedigree and provenance, they may reference sensitive organizational or personal information (e.g., contact details for shipping events). To address privacy risks, ecosystems should adopt core privacy engineering principles:

- **Data Minimization:** Collect and share only the minimum data necessary to fulfill traceability use cases.

- **Redaction and Pseudonymization:** Ensure sensitive fields are masked when queried across ecosystems.
- **Scoped Retention and Purpose Limitation:** Retain data only as long as necessary and use it strictly for the traceability purposes for which it was collected.
- **Transparency:** Ensure individuals whose data may be collected during traceability events are informed about how their data is used and shared.

D.3. Balancing High-Assurance Identity and Privacy Risks

Cryptographic object identifiers (e.g., Tracked Entity IDs) are foundational for verifiable traceability. However, when deployed into operational environments (e.g., during warranty claims or system maintenance), these identifiers may introduce privacy risks if correlated with end-user activity or location data.

To address the tension between integrity assurance and privacy protection, ecosystem operators should use privacy risk management tools, such as Privacy Impact Assessments (PIAs) or the NIST Privacy Risk Assessment Methodology (PRAM)[3], to evaluate traceability risks in downstream use cases. Ecosystems are encouraged to incorporate privacy-enhancing technologies (e.g., zero-knowledge proofs, selective disclosure)[4] into their architectural decisions to balance traceability utility with privacy obligations.

D.4. Threat Modeling and Ecosystem Risk Posture

Ecosystems vary widely in terms of industry context and operational complexity. Therefore, each ecosystem should perform ongoing threat modeling and risk analysis to identify potential attack vectors, ranging from record tampering and unauthorized access to metadata inference, and define the appropriate level of security controls.

Ecosystems should align their threat modeling practices with established frameworks, such as NIST SP 800-30 for risk assessments and the NIST Cybersecurity Framework, to structure their risk responses. By integrating zero-trust principles and minimizing trust assumptions across ecosystem boundaries, organizations can evolve their security postures to meet operational needs.

D.5. Operational Resilience

In addition to access controls and privacy measures, ecosystems should plan for operational resilience by implementing:

[3] The NIST Privacy Risk Assessment Methodology (PRAM) helps organizations analyze, assess, and prioritize privacy risks to determine how to respond and select appropriate solution. The PRAM can be found at https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/resources.
[4] To learn more about differential privacy, see NIST Special Publication 800-226, Guidelines for Evaluating Differential Privacy Guarantees. Available at https://doi.org/10.6028/NIST.SP.800-226.

- **Audit and Monitoring:** Maintaining secure logs of traceability record access to enable the early detection of anomalous behavior.
- **Data Encryption:** Ensuring records and related metadata are encrypted both in transit and at rest using industry-accepted standards.
- **Incident Response:** Establishing plans for responding to cybersecurity incidents, including preserving forensic data and restoring trust in affected traceability chains.

Appendix E. Potential Future Directions for the Meta-Framework

The Meta-Framework is designed to establish a principle-based approach to defining structural data patterns for a core set of supply chain events. These foundational event templates (*Make, Assemble, Store, Ship, Receive,* and *Employ*) capture the necessary pedigree information for the product, the responsible organization, timestamps, and, as appropriate, links to preceding record(s). Each manufacturer and industry sector can conceptualize and refine the template to fit their specific operational needs.

Because supply chains and product lifecycles are constantly evolving, this framework is conceptually extensible. The structural data patterns enable the introduction of new supply chain lifecycle event types as needed. However, how these new events should be modeled requires further study to ensure they can be integrated while maintaining the semantic interoperability and predecessor linkage required to support efficient traceback during SCRM analysis. This appendix outlines preliminary areas for future research and community collaboration and contains illustrative examples of additional event templates that could support broader approaches to lifecycle traceability.

E.1. Expanding Traceability to Lifecycle Chain Traceability Records

After the supply chain receives a product, the acquirer conducts an SCRM analysis to determine whether it is valid, unaltered, and suitable for the intended use. If accepted, the acquirer generates an Employ record, which is cryptographically linked to the predecessor Receive event. The Employ record serves as the anchor for the product's active lifecycle and may indicate the results of the initial SCRM analysis.

Because consistently maintained records in traceability chains significantly reduce the time required to perform SCRM analysis, acquirers are better positioned to conduct periodic checks and continuous monitoring for new risks post-deployment. As the acquirer performs ongoing lifecycle management, a key area for future research is determining which new post-Employ lifecycle event records may be valuable to template and use in traceability chains.

Based on preliminary brainstorming and initial feedback from the community of interest, Table 7 outlines several candidate events for this expanded lifecycle visibility.

Table 7. Candidate Post-Employ Lifecycle Chain Events

Candidate Event	Description
Return	Could capture when a product is returned by an end customer. Recording returns as traceability events contain verifiable proof that a product has been removed from service or is no longer in the customer’s possession.
Recall	Could trace a product back to the customer during a manufacturer-initiated recall. If a customer is also a manufacturer or integrator, research is needed to determine how they could sequentially pass the recall notification to their downstream customers, enabling efficient recall cascading.

Candidate Event	Description
Refurbish	Could capture maintenance actions, such as major software upgrades or hardware replacements. Research must determine how a Refurbish record ensures all significant product modifications are documented and properly linked to both the original asset and any newly integrated components.
Transfer	Could trace the handoff of custody, ownership, or operational control (Chain of Title) between organizations or environments. This would support scenarios involving leasing, resale, or cross-border movement.
Dispose	Could capture end-of-life actions, such as disconnection from IT/Operational Technology (OT) environments, physical destruction, or secure e-waste processing, confirming the product is no longer available for redeployment.

To illustrate the types of structural questions this future research must address, consider the initial brainstorming surrounding the potential Refurbish event listed above. If an acquirer learns that a hardware component within a deployed product must be upgraded, a Refurbish record could theoretically be generated. As illustrated conceptually in Fig. 9, additional research is needed to define how this event could feature multiple linkages: pointing back to (a) the original product's Employ record (or another post-Employ record) and (b) the Receive record of the newly acquired replacement component.

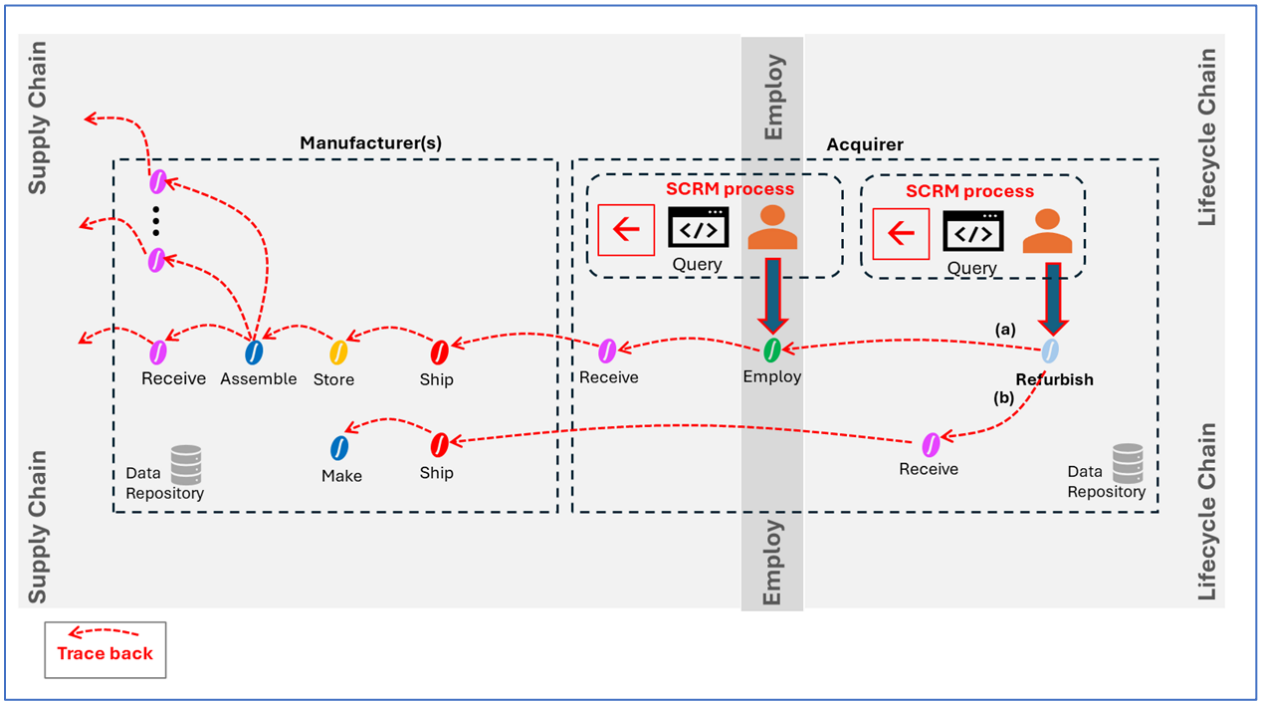

Fig. 9: Additional Lifecycle Opportunities for Future Research (Refurbish)

Defining the precise structural mechanics for Refurbish, along with the other candidate events in Table 7, will require similar brainstorming and extensive coordination with industry stakeholders to ensure the resulting data templates accurately reflect real-world lifecycle management practices.

Finally, while full data lifecycle for traceability records or embedded product data is not directly governed within the Meta-Framework, organizations implementing these patterns should establish data governance policies that address record retention, archiving, and disposal. In cases where decommissioned products contain embedded data (e.g., logs, user information, cryptographic credentials), appropriate system-level procedures for secure deletion or sanitization should be applied independent of the traceability layer.

E.2. Additional Supply Chain Traceability Records

In addition to post-*Employ* lifecycle events, future research should explore new inter-organizational supply chain events to capture additional data on manufacturing and logistics. Table 8 describes conceptual candidates for these expanded supply chain event templates. However, these initial concepts will likely evolve, and some may need to be broken down into multiple distinct event types.

Table 8. Candidate New Supply Chain Events

Candidate Event	Description
Precursor	**Could trace raw materials or base elements (e.g., silica used in semiconductor manufacturing) through the earliest stages of production, providing a more comprehensive view of the deep supply chain.**
Process / Convert	**Could capture events for continuous flow, batch, and other transformative manufacturing processes. Research is needed to determine how ecosystems can structurally distinguish between continuous material production, batch processing, and discrete component assembly.**
Split	**Could capture scenarios where a single product, material lot, or shipment is divided into multiple distinct entities while maintaining traceability to the original source (e.g., cutting a silicon wafer into individual chips). Investigating this event type could help ensure that the relationship between parent materials and derived components is securely documented.**
Modify	**Could capture material changes to an existing component prior to final assembly (e.g., component-level firmware flashing, reworking a defective part, or applying specialized heat treatments).**
Transport	**Could enhance visibility into the logistics phases between the Ship and Receive events. Further research is required to determine whether "Transport" is best represented as a single event or as a complex subset of logistical event types (e.g., intermodal transfers, cross-docking, customs processing). Additionally, future iterations of existing Ship and Receive event structures may need to be revisited to fully integrate with these new logistical linkages.**

Appendix F. Meta-Framework Reference Implementation

The Meta-Framework Reference Implementation (RI) is a minimum viable product (MVP) developed jointly by NIST and MITRE based on the examples and information provided in the second public draft of NIST IR 8536 (https://csrc.nist.gov/pubs/ir/8536/2pd). It is designed to investigate and demonstrate the concepts presented in this publication, with a specific focus on supply chain traceability across independently governed ecosystems. The current codebase is an early-stage research artifact intended strictly for laboratory experimentation, prototyping, and exploration, rather than a production-grade software solution.

F.1. RI MVP Project Goals

The primary goals of the RI are to:

- Demonstrate multi-ecosystem traceability.
- Validate the data model, traceability links, and data integrity established by the Meta-Framework.
- Support continued laboratory experimentation regarding cybersecurity, interoperability, and governance through modular, configuration-driven components.

F.2. RI MVP High-Level Architecture

The RI establishes three example ecosystems representing distinct roles within a manufacturing supply chain: Ecosystem-A (a Microelectronics manufacturer (ME-001)), Ecosystem-B (an Operational Technology Integrator (OT-001)), and Ecosystem-C (a Critical Infrastructure Acquirer (CI-001)). Because the Meta-Framework is technology-neutral, the RI defaults to using simple text files for its trusted data repositories, with NoSQL available as a configuration option. It also contains a lightweight Python web server dashboard for viewing ecosystem records and navigating traceability links.

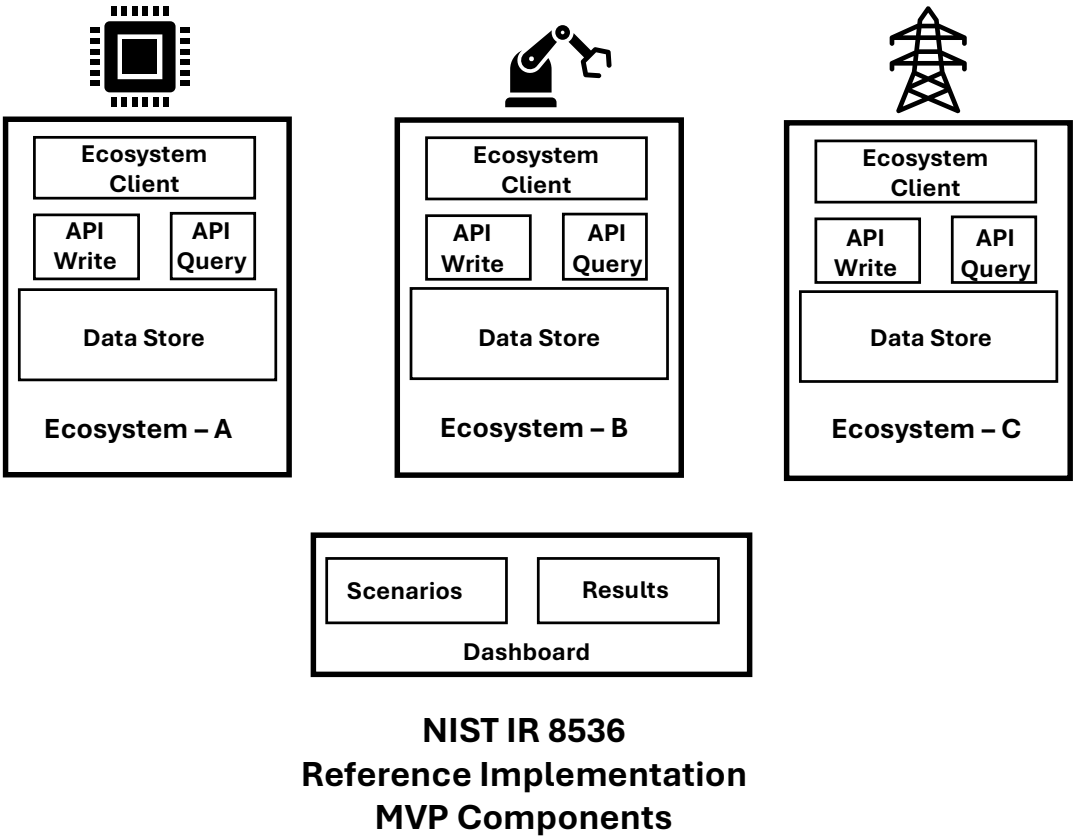

Fig. 10. NIST IR 8536 Reference Implementation Components

F.3. Supply Chain Event Mapping

To demonstrate the creation of a continuous, cryptographically linked provenance chain, the implementation contains scripts that generate and submit specific supply chain events across the three ecosystems. The events are mapped as follows:

- Microelectronics Manufacturer (ME-001): Executes the "Make" and "Ship" event records.
- OT Integrator (OT-001): Executes the "Receive" and "Assemble" event records.
- Critical Infrastructure Acquirer (CI-001): Executes the "Employ" event record.

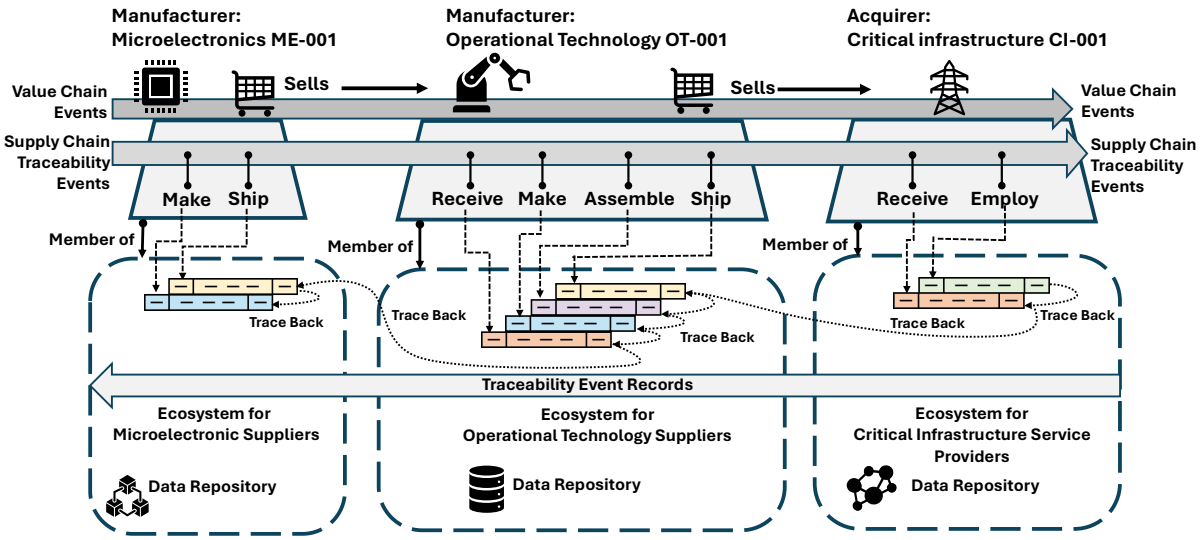

Fig. 11. Simplified Example of Ecosystems to Support Lab Experimentations

F.4. Accessing the Repository

The software is Python-based and has been tested on Windows, Mac, and Linux. The complete source code, installation instructions, technical documentation, and guidance for community contribution are publicly available.

Researchers and implementers can access the repository at:

https://github.com/usnistgov/ir8536-reference-implementation