PROJECT DESCRIPTION

MANUFACTURING SUPPLY CHAIN TRACEABILITY WITH BLOCKCHAIN RELATED TECHNOLOGY

Reference Implementation

Michael Pease
Keith Stouffer
Smart Connected Systems Division
Communications Technology Laboratory

Evan Wallace
Systems Integration Division
Engineering Laboratory

Harvey Reed
Steve Granata
The MITRE Corporation
McLean VA

FINAL

August 2023

blockchain_nccoe@nist.gov

This revision incorporates comments from the public.

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

This document describes a problem that is relevant to many industry sectors. NCCoE cybersecurity experts will address this challenge through collaboration with a Community of Interest, including vendors of cybersecurity solutions. The resulting reference design will detail an approach that can be incorporated across multiple sectors.

Abstract

Manufacturing supply chains are increasingly critical to maintaining the health, security, and the economic strength of the United States. As supply chains supporting Critical Infrastructure become more complex and the origins of products become harder to discern, efforts are emerging that improve traceability of goods by exchanging traceability data records using distributed ledger and other blockchain related technologies. Recent events and current economic conditions exposed the impact of disruptions in the security and continuity of the U.S. national manufacturing supply chain. This in turn, drew critical attention to the need to illuminate and secure the supply chain from numerous hazards and risks. Further, the U.S. manufacturing supply chain is susceptible to logistical disruptions, in addition to the effects of nefarious actors seeking fraudulent gain or attempting to sabotage or corrupt manufactured products. Improving the traceability of goods and materials that flow through the manufacturing supply chain may help mitigate these risks. This project will continue building on ongoing NCCoE efforts to demonstrate the role that blockchain related technologies may play to improve manufacturing supply chain traceability and integrity by exploring several use cases and the issues surrounding implementing supply chain traceability and will result in a freely available NIST Cybersecurity publication.

Keywords

blockchain; distributed ledger; distributed permissioned ledger; DLT; pedigree; provenance; supply chain traceability; traceability chain.

Disclaimer

Certain commercial entities, equipment, products, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST or NCCoE, nor is it intended to imply that the entities, equipment, products, or materials are necessarily the best available for the purpose.

Table of Contents

List of Figures

List of Tables

1 EXECUTIVE SUMMARY

Purpose

Manufacturing supply chains are increasingly critical to maintaining the health, security, and the economic strength of the United States. As supply chains supporting critical infrastructure become more complex and the origins of products become harder to discern, efforts are emerging that improve traceability of goods by exchanging traceability data records using ecosystems enabled by Distributed Ledger Technologies (DLTs) and other blockchain related technologies[1] that provide provenance and integrity.

This document describes a Minimum Viable Product (MVP) Reference Implementation (RI) of manufacturing supply chain ecosystems, to illustrate product traceability across microelectronic and OT (Operational Technology) supply chains to critical infrastructure operators. The MVP RI is a follow-on effort from NISTIR 8419 “Blockchain and Related Technologies to Support Manufacturing Supply Chain Traceability.” [1] In addition, the project seeks technical exchange and discussion with related groups (e.g., industry and standards groups [2] [3] [4]) to discover and refine relevant MVP use cases regarding data sharing of traceability information; data, pedigree and provenance integrity; and manufacturing supply chain wide traceability queries.

The choice of microelectronics and operational technology emphasizes the importance of manufacturing supply chain traceability, although the MVP RI should be understandable in other contexts and serve as an architectural approach for other supply chain domains and critical infrastructure sectors.

The choice of critical infrastructure as the consumer of the notional supply chain in this MVP effort emphasizes the importance of manufactured products, and constituent products and assemblies therein, used for purposes critical to civil society. This MVP approach may also be adapted to national security and other contexts. Ultimately, the stakeholders and participants in individual supply chains will determine which entities become members of their ecosystems. In addition to ecosystem membership, stakeholders and participants will mutually determine the composition of data elements and standards that define the data in their distributed ledger transactions. The ecosystem membership and associated data agreements are necessary to implement an MVP ecosystem for their industry sector or sub-sector.

This project has a goal to demonstrate traceability across manufacturing domain stakeholder blockchain related technologies enabled ecosystems [1] to determine authenticity of products for use in critical infrastructures. The project will continue building on NCCoE ongoing efforts to demonstrate the role that blockchain related technologies may play to improve manufacturing supply chain traceability. This project will result in a freely available NIST Cybersecurity Practice Guide. For the specific architecture used in this MVP, blockchain will be used as the as an

[1] The terms “Distributed Ledger Technologies (DLTs) and other Blockchain related technologies” are used interchangeable throughout the documents and refer to the family of technologies around permissioned ledgers, and confidential distributed ledgers that provide integrity, traceability, and identity information about items and who added them using byzantine fault tolerance consensus mechanisms.

example of blockchain related technologies; however, other implementations such as confidential DLT are also within the scope of possibilities for this work.

Scope

This project addresses key challenges in the manufacturing supply chain:

- Improve visibility, integrity, and permanence of manufacturing supply chain product pedigree. The initial claim of product authenticity by a manufacturer needs to survive the lifetime of the manufacturer through mergers, acquisitions, and dissolution.
- Improve visibility and integrity of provenance across tiers of manufacturers. The existing process of tracking provenance via bi-lateral exchange of traceability information between buyer and seller is: (a) complicated, and (b) non-permanent, where information may be lost or further obscured during mergers, acquisitions, and dissolutions.

This project describes and delivers a reference implementation of a potential manufacturing supply chain traceability mechanism that demonstrates:

- Manufacturers' ability to post traceability records to their respective industry ecosystem DLTs. Each traceability record written to the DLT links to the prior traceability record(s), going back to the original traceability record(s) (e.g., 'making' the product) where the traceability record links to the originating manufacturer.
- Establishing traceability record links and forms as an immutable traceability chain. Traceability records can link to multiple prior traceability records in the case of combining components in higher-order assemblies and products.
- Associating traceability records link to relevant context. In addition to linking to previous traceability records, traceability records point to relevant context such as the author (e.g., who wrote the record) and additional data in external repositories as needed.
- Establishing traceability record links to external data as required. In addition to the minimal data in the traceability record, traceability can link to external data as needed (with appropriate access controls) for larger data sets, images, audio, video, etc.

This project delivers an MVP RI that:

- Demonstrates manufacturers joining their respective blockchain-related technology-enabled ecosystems.
- Demonstrates manufacturers writing and linking traceability records.
- Demonstrates critical infrastructure operators reading the traceability chain to inform their assessment whether to employ the manufactured product.
- Uses microelectronics, operational technology, and critical infrastructure as example domains.
- Positions the MVP RI as a starting point for future research and refinement.

Assumptions/Challenges

The key project challenge is to explain and illustrate the traceability chain method with sufficient fidelity to indicate potential suitability for traceability of complex manufacturing supply chains, while avoiding detail which may be better suited for future refinement. The key assumption is that the MVP project, once complete, is a starting point for further research and refinement. Beyond the scope of the MVP, further topics can be explored, such as ecosystem governance, identity proofing, cyber-physical identification, as well as adaptation of the MVP by supply chains in other industry sectors and sub-sectors and/or further in-depth application of the MVP in manufacturing supply chains.

Background

Supply chain participants are motivated to increase traceability in complex manufacturing supply chains to mitigate risk of supply chain vulnerabilities [5]. Vulnerabilities can arise in any manufacturing supply chain and are exemplified by the OT (Operational Technology) domains. OT includes hardware, software, and managed services, where consequences of OT supply chain vulnerabilities can impact the daily operation of U.S. critical infrastructure [6]. Today, organizations lack the ability to readily distinguish between trustworthy and untrustworthy products. Having a repeatable, quick, and provable means to determine if a product is trustworthy is a critical foundation of cybersecurity supply chain risk management [7].

An ecosystem perspective of the manufacturing supply chain serves to define provable traceability for a subset (an ecosystem) of the manufacturing supply chain stakeholders (e.g., suppliers, critical infrastructure), and to share and store applicable product traceability data records (e.g., pedigree, provenance). Traceability requirements and their means of implementation will be unique for each ecosystem (e.g., microelectronics, operational technology, critical infrastructure).

Traceability data includes information about product provenance, pedigree, and other data as needed. Early industry ecosystem efforts indicate that the ecosystem perspective is useful and perhaps necessary to enable trusted and symmetric supply chain information sharing and migrate away from existing linear and bi-lateral information exchange. The existing status quo of bi-lateral information sharing is susceptible to incomplete coverages, differing implementations, corruption and alteration of data, and potential semantic gaps in data elements. A semantic gap may occur when a stakeholder multiple tiers away writes or conveys a traceability record that may not be fully understood or recognized downstream. Ecosystem-wide agreement on traceability information requirements, mitigates semantic gaps in understanding traceability data records within a manufacturing domain. This ecosystem perspective is layered atop, and does not replace, the existing and prevalent “per acquirer” perspective of supply chain management and security.

Across complex manufacturing supply chains, multiple ecosystems will arise and must themselves link traceability information across the ecosystems in order to establish trusted and symmetric traceability data, from commodities to final assemblies used in critical infrastructure, where products include hardware, software, and services [1]. The resulting traceability chain across industry ecosystems provides a path (links) to follow traceability records across ecosystems. The linking of traceability records can be performed with a small number of data fields. Further, traceability records can be specialized to meet the needs of various industry sectors as needed. The traceability links allow for multiple source components to be combined in an assembly, where the traceability record for the assembly can contain a list of constituent links back to the sourced components. This enables a tree structure of links, with a critical infrastructure acquirer ultimately receiving the root traceability record. The root traceability record can then be followed backwards, or upstream in the product supply chain, as necessary through ecosystems and across the chain of product traceability records.

2 Scenarios

Scenario Stakeholders and Ecosystems

The following ecosystems and manufacturing stakeholders are used in the MVP scenarios to illustrate the MVP traceability chain mechanism as shown in Figure 1: Manufacturers Participate in DLT:

- Three (3) distinct blockchain related technology (e.g., DLT) enabled ecosystems:
    1. Microelectronic Provider (MEP) manufacturing domain
    2. Operational Technology (OT) manufacturing domain
    3. Critical Infrastructure (CI) domain
- Three (3) distinct manufacturing stakeholders:
    1. MEP-001 – microelectronic manufacturer
    2. OT-001 – operational technology manufacturer
    3. CI-001 – critical infrastructure operator

The manufacturing stakeholders participate in an economic value chain, where value chain activities result in manufacture, making, and employing products. When products are made, included in assemblies, and ultimately used by the end operating environment, **traceability records** are written to the ecosystem DLT. This provides both permanence for the traceability chain, surviving company mergers, acquisitions, and dissolutions, and a simplification of navigating traceability chains. The manufacturing domain ecosystems evolve slower than the constituent manufacturing stakeholders, and once established persist over time, providing permanence to the traceability records.

The manufactured products used in the scenarios are assumed to be represented in data records, but not manifested physically or in software code. The relationships between the stakeholders and ecosystems used in the MVP are illustrated below.

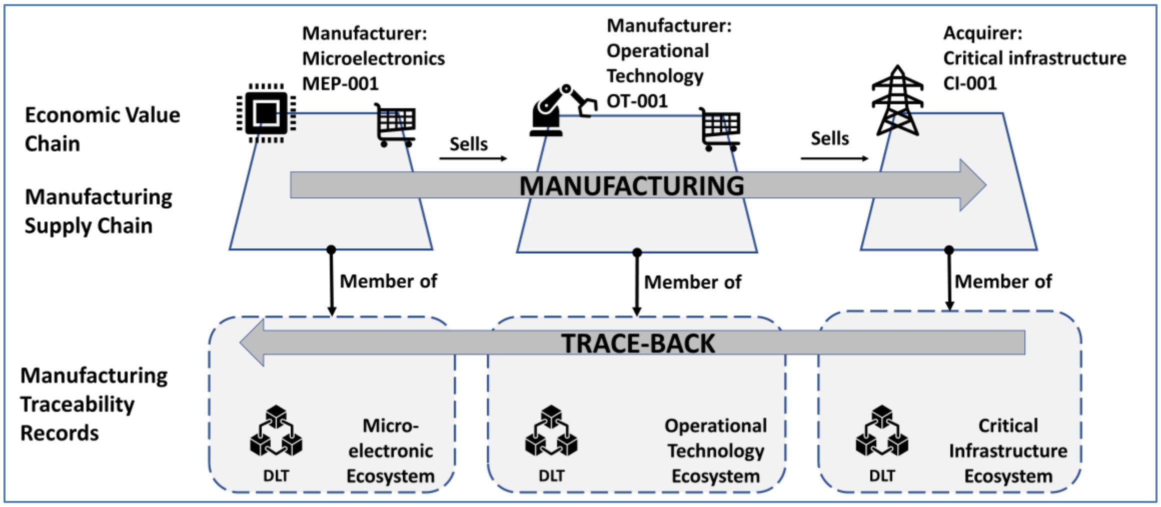
Economic Value Chain
Manufacturing Supply Chain
Manufacturing Traceability Records
Manufacturer: Microelectronics MEP-001
Manufacturer: Operational Technology OT-001
Acquirer: Critical infrastructure CI-001
Sells
Sells
MANUFACTURING
Member of
Member of
Member of
TRACE-BACK
DLT
Micro-electronic Ecosystem
DLT
Operational Technology Ecosystem
DLT
Critical Infrastructure Ecosystem

Figure 1: Manufacturers Participate in DLT Enabled Ecosystems to Record Traceability Records

Scenario 1: Supply chain manufactures operational technology assembly.

MVP Scenario 1 exercises the set of manufacturing domain ecosystems to produce and sell manufactured goods for procurement by critical infrastructure, recording traceability data to establish pedigree and provenance:

1. MEP-001 produces a chip and sells the chip to OT-001:
   a. Marks the chip with a unique ID
   b. MEP-001 creates a traceability record and writes it to the microelectronic traceability ecosystem. The traceability record has a URI pointer to internal private manufacturing data, the ID of the chip, and a digest of traceability manufacturing data including hashes as needed, and the identity of MEP-001 and purchaser OT-001.
   c. MEP-001 virtually delivers the chip to the purchaser, an operational technology manufacturer OT-001.
2. OT-001 records receipt of the virtual chip and applicable chip traceability data and writes a traceability record, in the operational technology ecosystem DLT, acknowledging receipt which contains the ID of MEP-001, OT-001, and the ID of the chip:
   a. OT-001 adds their software to the chip, where the software development steps are assumed to be traceable themselves, but (similar to the chip manufacturing above) doesn’t have to be demonstrated just referenced via URI.
   b. OT-001 adds the chip and software to an operational technology assembly and virtually delivers the operational technology assembly to critical infrastructure operator CI-001.
3. CI-001 records receipt of the operational technology assembly and writes a traceability record, in the critical infrastructure ecosystem DLT, acknowledging receipt which contains the ID of OT-001, and the ID of the operational technology assembly.
   a. CI-001 starts a process to verify authenticity of the operational technology assembly.
4. Include additional chip and software deliveries which are invalid.
   a. Emulate fraudulent parts to test whether authenticity queries (see Scenario 2) can detect the fraudulent manufactured goods.

Scenario #1 (sub parts 1-3) is notionally illustrated below.

Diagram depicting three distinct manufacturing ecosystems: MEP-001 for microelectronic manufacturers, OT-001 for operational technology manufacturers, and CI-001 for critical infrastructure acquirers. The diagram highlights the different activities occurring within each ecosystem and the creation of traceability records in ecosystem-specific distributed ledger technology implementations.

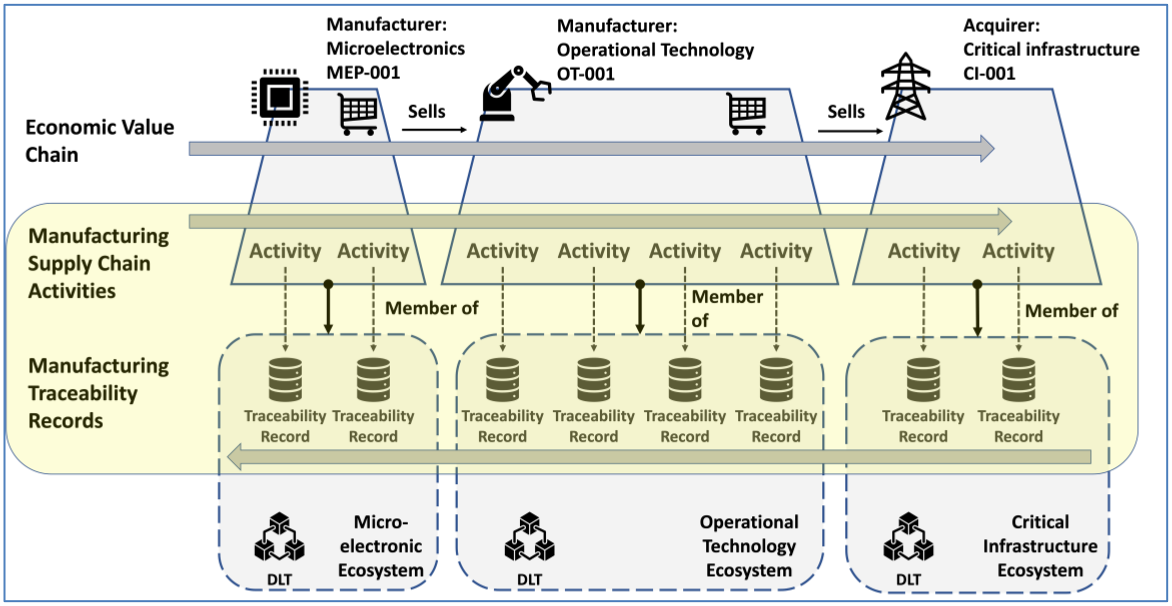

Figure 2: Traceability Chain Mirrors the Manufacturing Supply Chain in Reverse

Scenario 2: CI-001 uses the traceability chain to query the operational technology ecosystem and validate the authenticity of the operational technology assembly.

MVP Scenario 2 exercises the query facility of each ecosystem to determine if a received manufactured good is authentic, by querying traceability records written to ecosystems during manufacturing, for example:

1. OT-001 queries the microelectronic ecosystem using the chip ID as a primary query parameter.
2. CI-001 queries the operational technology ecosystem using the operational technology assembly ID as a primary query parameter.

Note: The scenario can include generated faults (counterfeit data records) to simulate general supply chain issues and identify how supply chain trackability can assist with detection. Generated faults may include:

- Swapping the genuine manufactured good (altering product ID), at point of sale, with a counterfeit part.
- Generate faults for chips, and the operational technology assembly which represent counterfeiting between manufacturer and acquirer.
- Generate faults for software which represent subversion of the software development process internal to OT-001.

Scenario 3: After installation, CI-001 performs statistical quality check to re-verify authenticity of the operational technology assembly.

MVP Scenario 3 also exercises traceability query facilities of each ecosystem as in Scenario 2. However, with a difference that the goods being verified are parts that are already in use in the critical infrastructure. This scenario demonstrates how the traceability ecosystems can continue to protect critical infrastructure after manufactured goods are in use. The MVP scenario will include generated faults to simulate a malicious actor swapping a valid manufactured good for a counterfeit and potentially malicious manufactured good.

All Scenarios: Traceability Chain

A traceability chain is a chain of linked traceability records. A traceability record is a transaction recorded to an ecosystem DLT, which is tamper evident and difficult to destroy. The manufacturing traceability records are of the sub-types: make, assemble, transport, receive, employ. The data fields in the sub-types are developed further in Section 3 below. The traceability record sub-types link to each other, providing an immutable traceability chain.

The diagram below illustrates the traceability record sub-types, and how they can be linked to form a traceability chain.

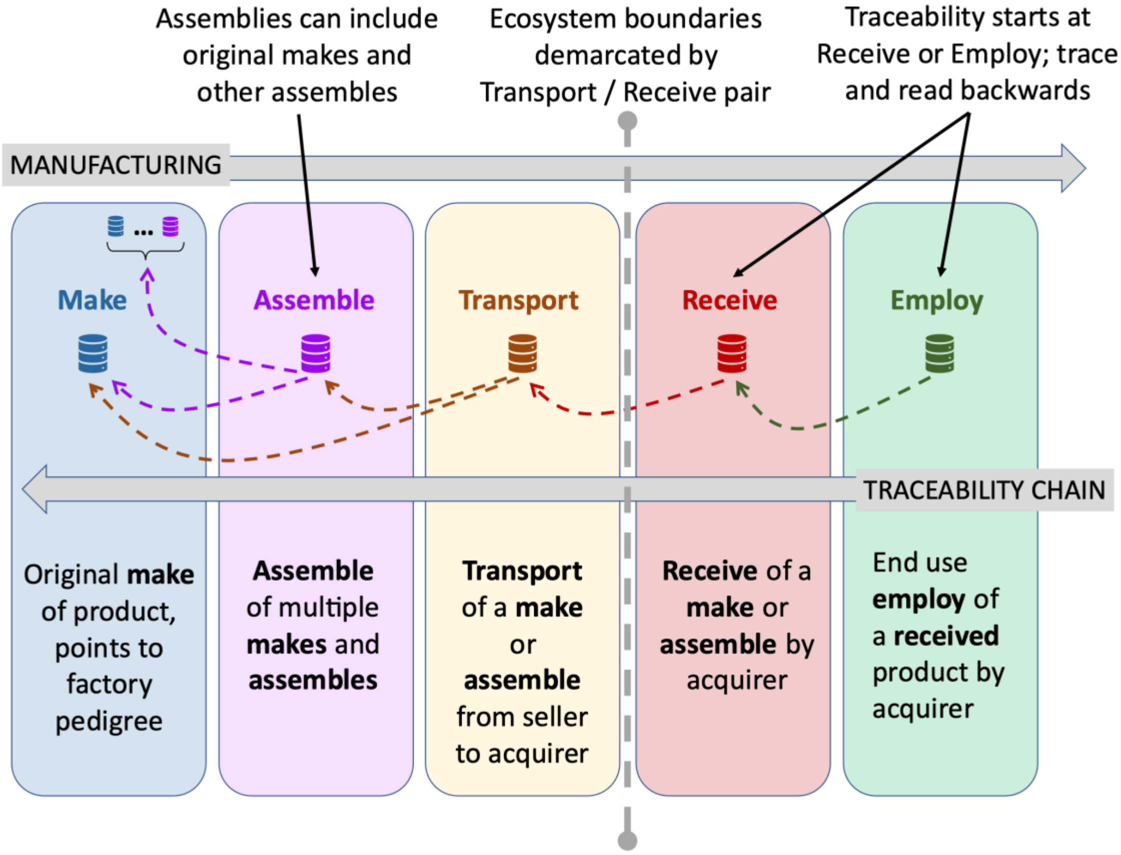
Assemblies can include original makes and other assembles
Ecosystem boundaries demarcated by Transport / Receive pair
Traceability starts at Receive or Employ; trace and read backwards
MANUFACTURING
Make
Assemble
Transport
Receive
Employ
TRACEABILITY CHAIN
Original make of product, points to factory pedigree
Assemble of multiple makes and assembles
Transport of a make or assemble from seller to acquirer
Receive of a make or assemble by acquirer
End use employ of a received product by acquirer

Figure 3: Traceability Records Form a Traceability Chain

The three scenarios above describe manufacturing actors making chips, software, and assembling them into operational technology, then selling the resulting assembly to a critical infrastructure entity.

Scenario #1 Revisited: Illustrated with Traceability Data Types

The primary purpose of the MVP is to illustrate traceability records linked in traceability chains, across the chip, operational technology, and critical infrastructure ecosystems, performing activities as outlined in the above scenarios.

Figure 4: Traceability Chain Mirrors the Manufacturing Supply Chain in Reverse is an illustrated lifecycle of Scenario #1 which creates and uses a manufacturing supply chain traceability chain across ecosystems. The lifecycle steps are denoted by circular numbered markers 1-8:

1. Chip manufacturer MEP-001 makes a chip and writes a **make-chip traceability record** with a statement of authentic product pedigree (summation of factory internal process, provenance, certification, testing, etc.) and links to the factory.
2. Chip manufacturer MEP-001 transports (ships, uploads, etc.) the chip to a buyer operational technology manufacturer OT-001, in a different ecosystem, and writes a **transport traceability record** which links to the **make-chip traceability** record, which in turn links to the factory.
3. OT-001 receives (loading dock, downloads, etc.) the chip, and writes a **receive traceability record** which links to the prior **transport traceability record**.
4. OT-001 makes software for the chip for use in an OT assembly and writes a **make-software traceability record** with a statement of authentic product pedigree (summation of software development internal process, software bill of materials (SBOM), etc.).
5. OT-001 makes an OT assembly with the chip, software, (could also include sensors, actuators, etc.), and writes an **assemble traceability record**, which includes the OT assembly pedigree, and links to the **chip receive traceability record** and the **make software traceability records**.
6. OT-001 transmits (ships, uploads, etc.) the OT assembly to a critical infrastructure CI-001 buyer, in a different ecosystem, and writes a **transport traceability record** which links to the **assembly traceability record**.
7. CI-001 receives OT assembly and writes a **receive traceability record** which links to the prior **transport traceability record**. The security officer for CI-001 uses the **receives traceability record** to trace-back through the traceability chain backward for pedigree and provenance information which informs the decision as to whether the OT assembly should be employed in the infrastructure.
8. The critical infrastructure acquirer CI-001 decides whether to employ the OT assembly and writes an **employ traceability record** that links back to the **receive traceability record**. The **employ traceability record** includes a link to the acquirer's decision documentation whether to employ the product, as well as documentation of where the product is employed, if the decision is to employ the product. Thus, this **employ traceability record** explains both the rationale of the employment decision and the capacity in which the employed product will be used. This **employ traceability record** enables periodic future inspection to determine whether the product may have been substituted inappropriately, thereby serving as a means to discover security risk vectors described in Scenario #3.

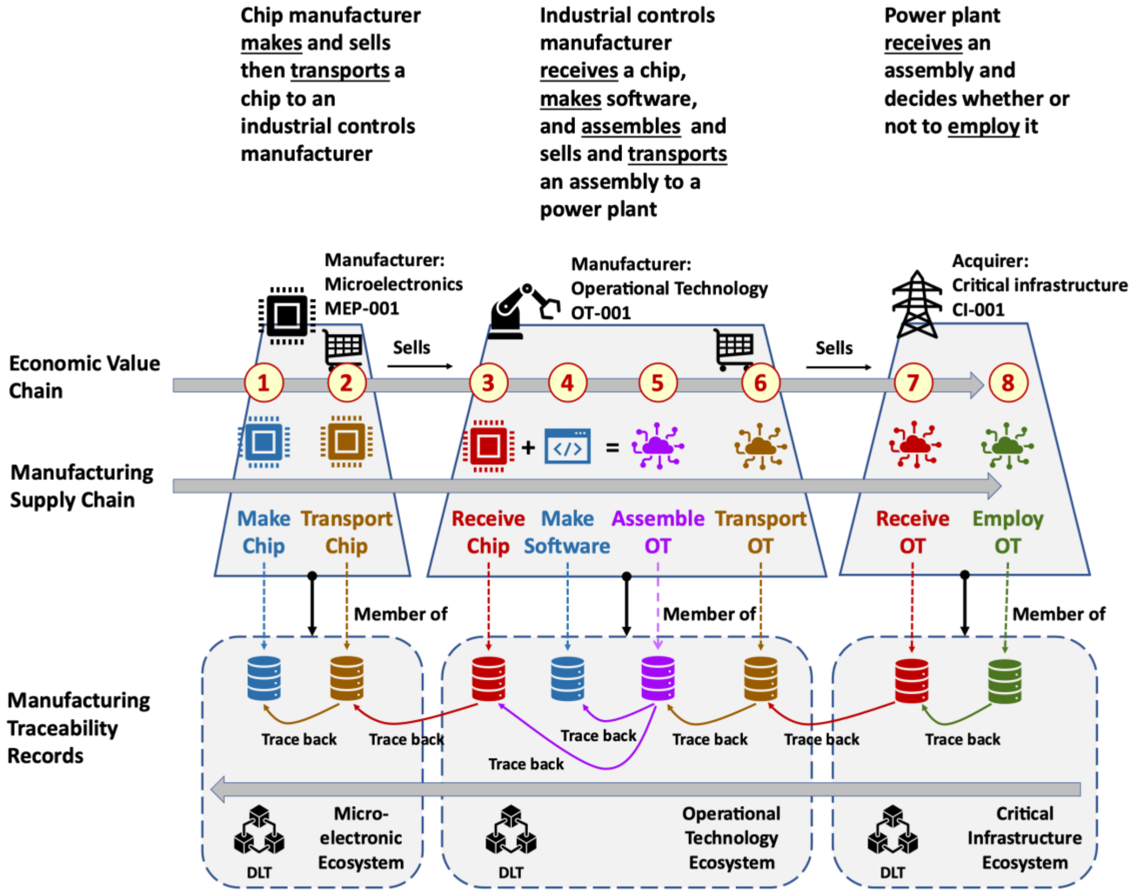
Chip manufacturer makes and sells then transports a chip to an industrial controls manufacturer
Industrial controls manufacturer receives a chip, makes software, and assembles and sells and transports an assembly to a power plant
Power plant receives an assembly and decides whether or not to employ it
Manufacturer: Microelectronics MEP-001
Manufacturer: Operational Technology OT-001
Acquirer: Critical infrastructure CI-001
Economic Value Chain
Sells
Sells
1
2
3
4
5
6
7
8
+
=
Manufacturing Supply Chain
Make Chip
Transport Chip
Receive Chip
Make Software
Assemble OT
Transport OT
Receive OT
Employ OT
Member of
Member of
Member of
Manufacturing Traceability Records
Trace back
Trace back
Trace back
Trace back
Trace back
Trace back
Trace back
DLT
Micro-electronic Ecosystem
DLT
Operational Technology Ecosystem
DLT
Critical Infrastructure Ecosystem

Figure 4: Traceability Chain Mirrors the Manufacturing Supply Chain in Reverse

Each Make traceability record is written to the DLT in the applicable manufacturing ecosystem at the time of the applicable manufacturing step, by the pertinent actors. Each subsequent traceability record is similarly written to an ecosystem DLT and points back to the preceding applicable traceability record. The subsequent traceability record could be in a different ecosystem and must contain sufficient metadata to support traceability chain query. This forms an immutable manufacturing traceability chain. This chain can later be ‘crawled’ backward through applicable ecosystem DLTs to read the whole traceability chain for full pedigree and provenance information, as described in Scenario #2. The manufacturing traceability chain provides provable manufacturer claims of authentic product pedigree, and provable provenance as the product moves through the supply chain.

In Figure 4 above, note that the direction of supply chain actions is depicted left-to-right as a timeline. The supply chain actions are recorded in traceability records, also written in timeline order. As new traceability records are written, the previous applicable traceability records are referenced with a hash-link. Thus, when the traceability chain is used to trace back, for example at the point in timeline of the Employ traceability record, the records are read from right-to-left.

Fully expanded, the shape of the manufacturing supply chain is a tree, and the shape of the corresponding manufacturing traceability chain is the same tree in reverse. **The primary objective of the MVP is to construct the traceability chain (linked traceability records) described above.**

Note: While this MVP will not require smart contracts, the MVP does not preclude the addition of smart contracts to illustrate additional financial and other transactional activities in the context of specific manufacturing traceability record ecosystem DLT transactions.

3 High-Level Architecture

Overview

The high-level architecture below develops the structure of the MVP components, expressed in a server/host architecture context. The high-level architecture description then continues to develop the data structure of traceability records and the resulting traceability chain, by stepping through the lifecycle of using traceability records to create a traceability chain.

Components and Server Architecture

Figure 5: Component and Server Architecture depicts the MVP components. The architecture separates the ecosystem hosts to emphasize that ecosystems (and DLT instances within) operate, evolve, and innovate independently. The single MVP identity provider provides the ecosystems with a consistent identity scheme. The scenarios are driven by, and results recorded in, a Scenario Dashboard as a separate component.

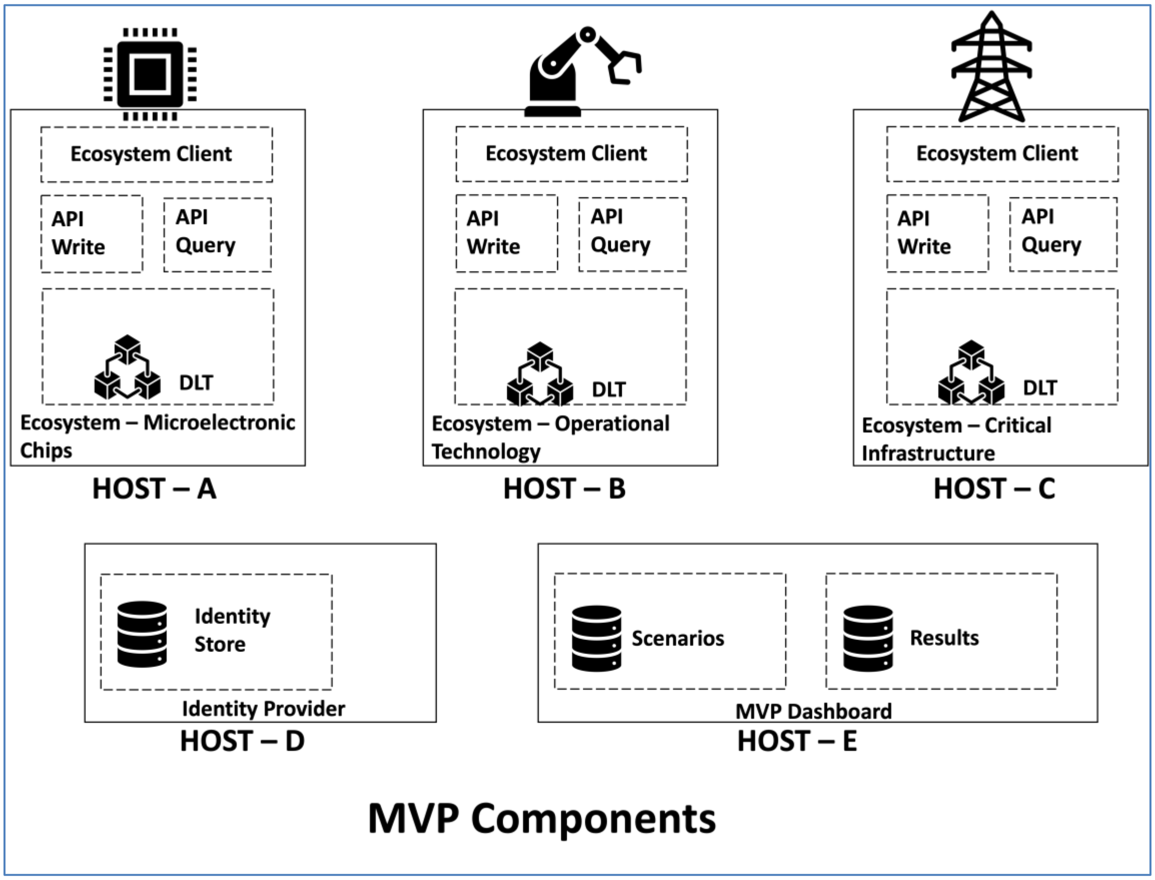

Figure 5: Component and Server Architecture

Identity (role-based)

The MVP assumes one identity provider and a single flat identity space across ecosystems; however, the transaction records contain information to support configurations where ecosystems may use independent identity providers. Identifiers can be simple labels, although in production, identities may be based on Credentials Community Group Decentralized Identifiers (W3C DID) emerging standards. Further, in production each ecosystem governance will independently generate their own identities. Identities for the MVP are role-based, and related to activities of Make, Assemble, Transport, Receive, and Employ.

Ecosystems (DLT and API component)

When a traceability chain is crawled, each link to the preceding traceability record can be followed, even to a different ecosystem, to the preceding traceability record. This link includes a hash of the preceding traceability record. For the MVP, the hash of the preceding traceability record can serve as simple authorization to access the preceding traceability record. The hash linking of traceability records is conceptually similar to linking blocks in a DLT, except a traceability chain is an inverted tree not a linear chain and spans multiple ecosystems. Thus, the traceability chain is a higher order data construct retaining the property of tamper evident data. Retaining the tamper evident property requires that each DLT also have the data integrity property of being tamper evident.

The use of multiple DLTs reflects the nature of manufacturing supply chains which include many affinity groups of similar manufacturers, and importantly can agree on the data fields of traceability records which are pertinent to the types of goods they manufacture. Using independent groupings of manufacturers which share affinity with each other gives rise to enabling ecosystems each of which use a DLT. These ecosystems expose an access-controlled application program interface (API) to the enclosed DLT within. Thus, the ecosystems and enclosed DLTs retain confidentiality, and the supporting DLTs provide data integrity. Further, the affinity groups can independently and incrementally update their ecosystems and DLTs, enabling rapid adaptation and encouraging adoption.

Note that critical infrastructures may adopt traceability ecosystems at a slower rate than the relevant manufacturing supply chains. Alternately, in the early phases of adoption, the critical infrastructure operating environments can store the traceability records (e.g., receive, employ) in their enterprise asset management and vulnerability analysis systems. If ecosystems are adopted by critical infrastructure operating environments, the traceability records can be stored there.

DLT

Each ecosystem will have an independent instance of DLTs. The DLT selected can be the same or differing types across the ecosystems.

Traceability Chain Lifecycle

The sequence of diagrams below illustrates the notional lifecycle of manufacturing traceability records written to industry ecosystem DLTs, and the resultant persistent and immutable traceability chain. The notional lifecycle informs the explication of traceability data types. The diagrams are accompanied by a high-level description of data associated by traceability records. Following the diagrams is a table of traceability records with a summary of applicable data fields.

NOTE: The number of ecosystems and where products are made below, is different from the MVP scenarios above. This difference highlights the flexibility of the traceability chain approach which is intended to accommodate an arbitrary number of stakeholders in an arbitrary number of ecosystems. Nonetheless, the data field requirements for each of the Make, Assemble, Transport, Receive, and Employ traceability records are the same in any situation.

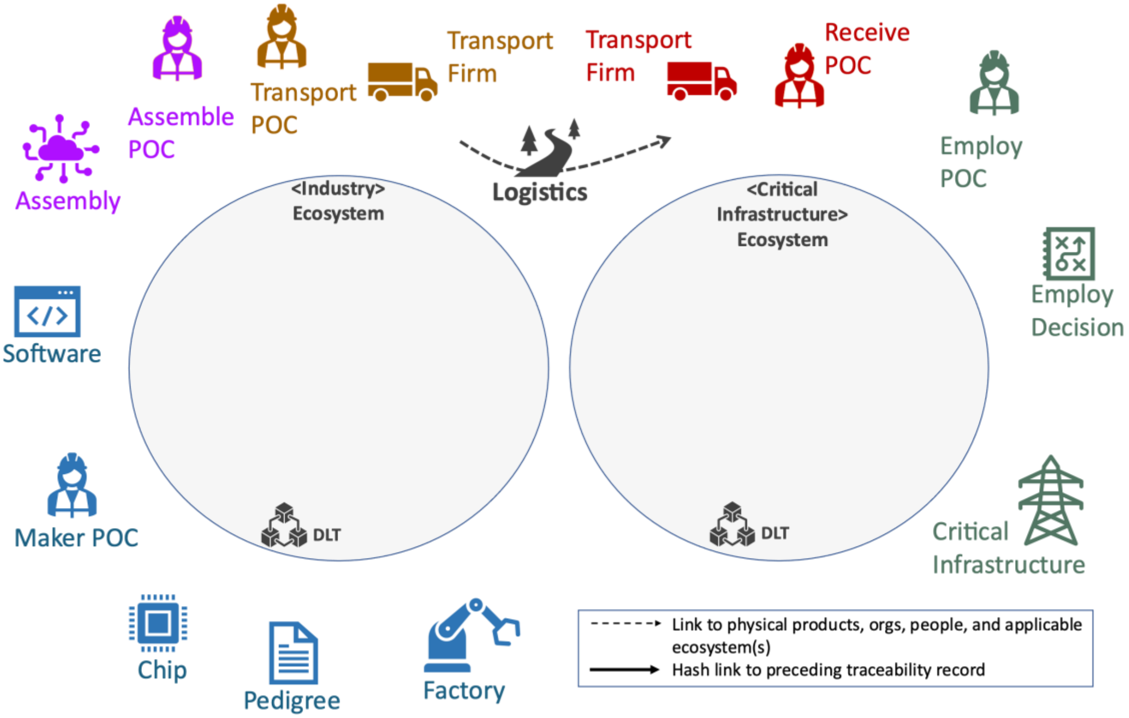

Figure 6: Traceability Chain Lifecycle - Actors

The actors include people and organizations (e.g., factories, critical infrastructure, transport firms), the ecosystems which group actors and enable actors to write transactions (e.g., traceability records), and the object of traceability (e.g., chip). The people actors are grouped into Make, Assemble, Transport, Receive, and Employ, responsible for those respective activities and are the author of the respective traceability records.

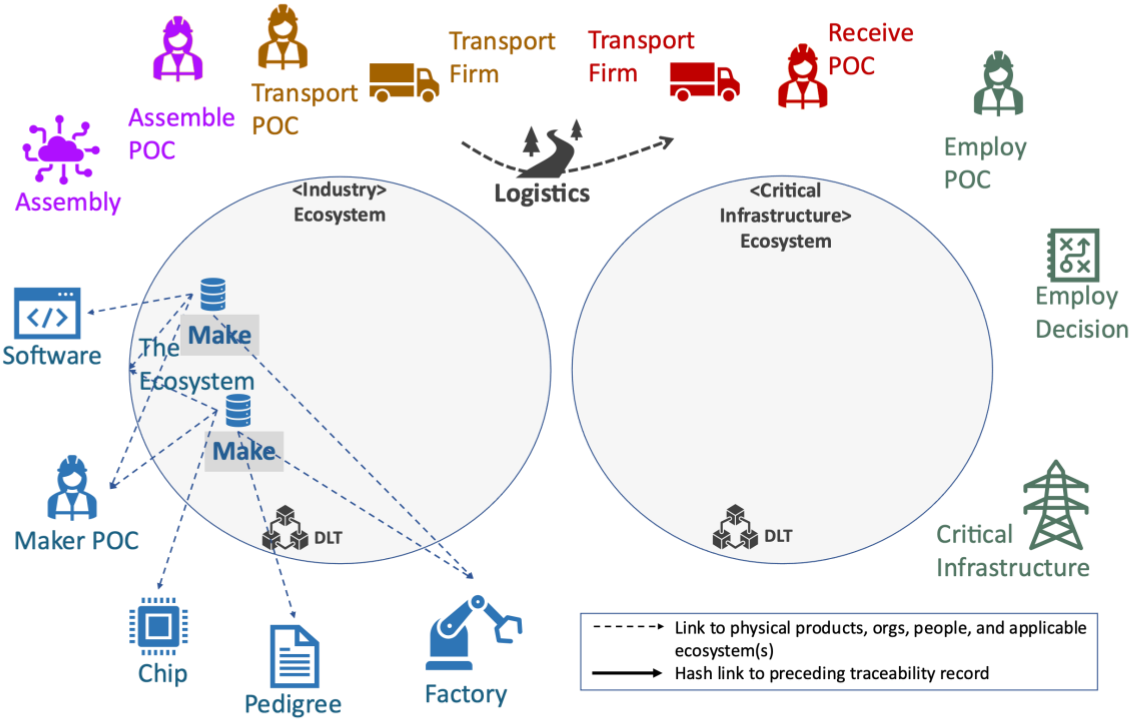

Figure 7: Notional Traceability Chain Lifecycle - Make

The Make POC writes a Make traceability record to the <industry> ecosystem. The make traceability record includes the Maker POC ID, the Product ID (e.g., chip), link to the Pedigree summary, and link to the Factory (if needed and agreed can query for more detailed pedigree). Another make traceability record is similarly written for software.

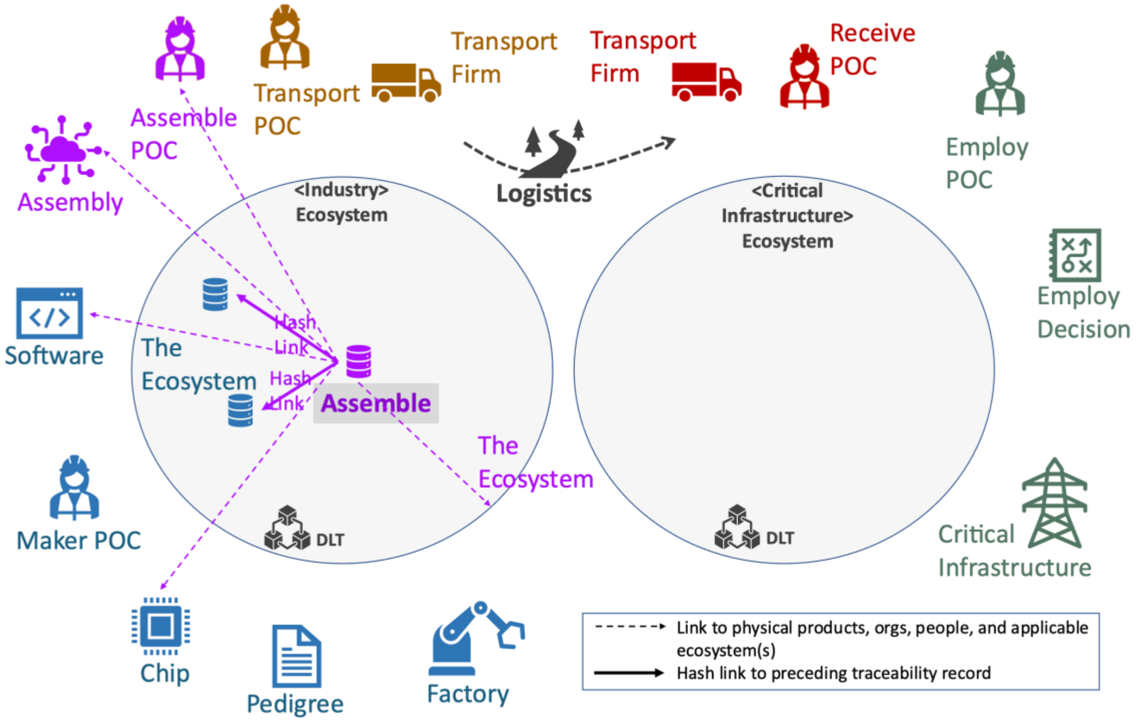

Figure 8: Notional Traceability Chain Lifecycle – Assemble

The Assemble POC writes an assemble traceability record to the <industry> ecosystem. The assemble traceability record includes a list (in this case two) of included products.

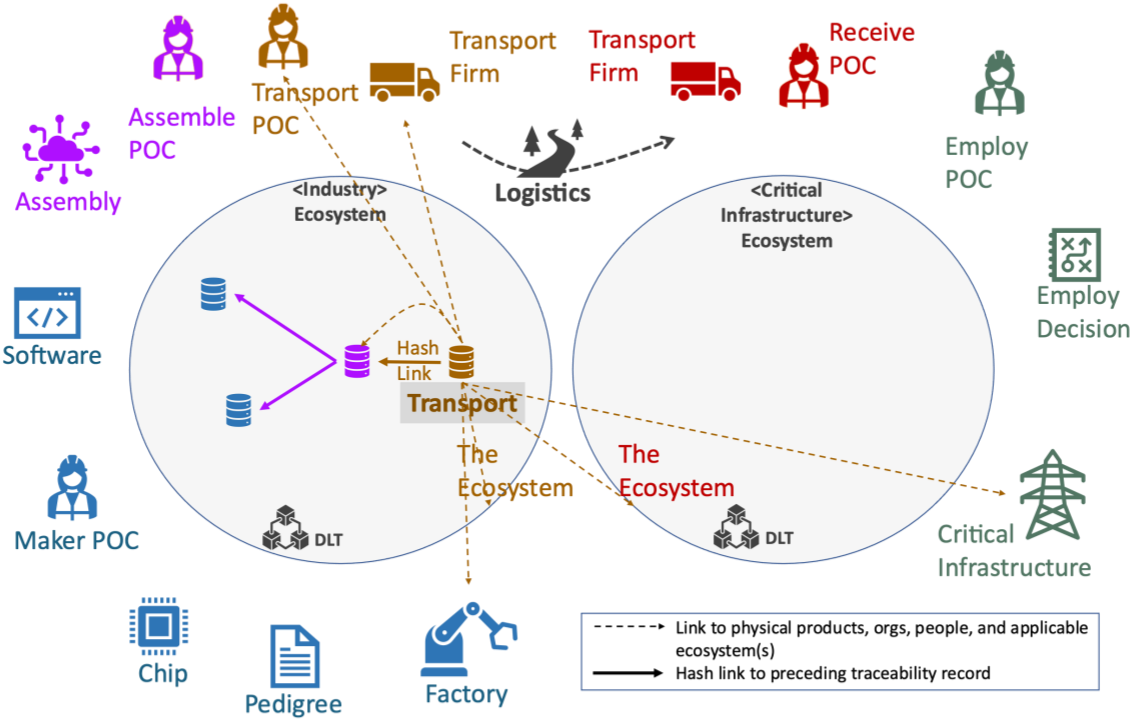

Figure 9: Notional Traceability Chain Lifecycle – Transport

The Transport POC writes a Transport traceability record to the <industry> ecosystem. The Transport traceability record includes the Transport POC ID, the Product ID (e.g., chip), the Factory ID, the original Make traceability record, the destination ecosystem, the destination org ID (e.g., critical infrastructure).

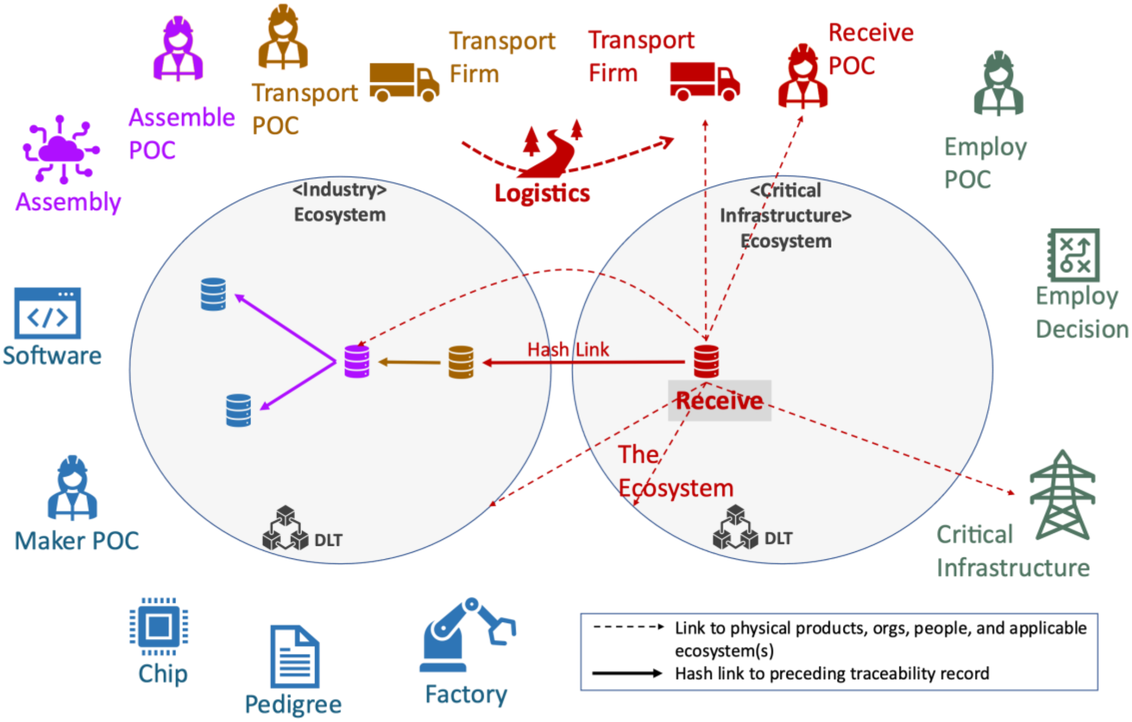

Figure 10: Notional Traceability Chain Lifecycle – Receive

The Receive POC writes a Receive traceability record to the <critical infrastructure> ecosystem. The Receive traceability record includes the Receive POC ID, the Product ID (e.g., chip), the Transport traceability record, the destination ecosystem, the destination org ID (e.g., critical infrastructure).

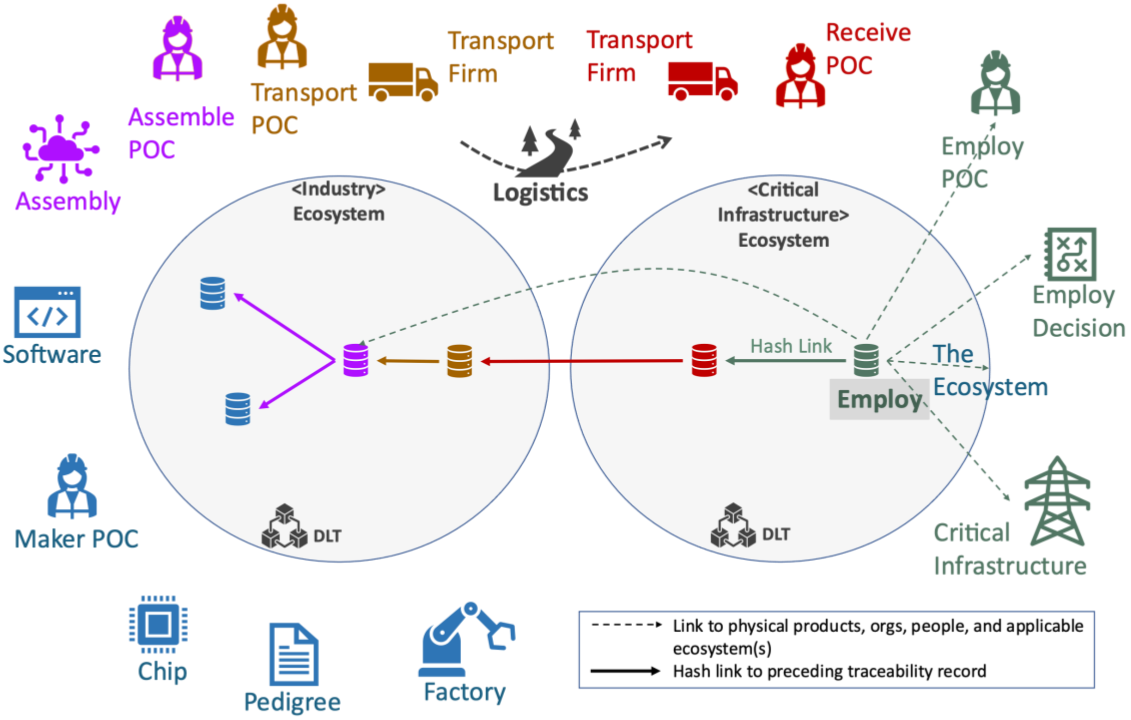

Figure 11: Notional Traceability Lifecycle – Employ

The Employ POC writes a Employ traceability record to the <critical infrastructure> ecosystem. The Employ traceability record includes the Employ POC ID, the Product ID (e.g., chip), the Receive traceability record, the destination ecosystem, the destination org ID (e.g., critical infrastructure).

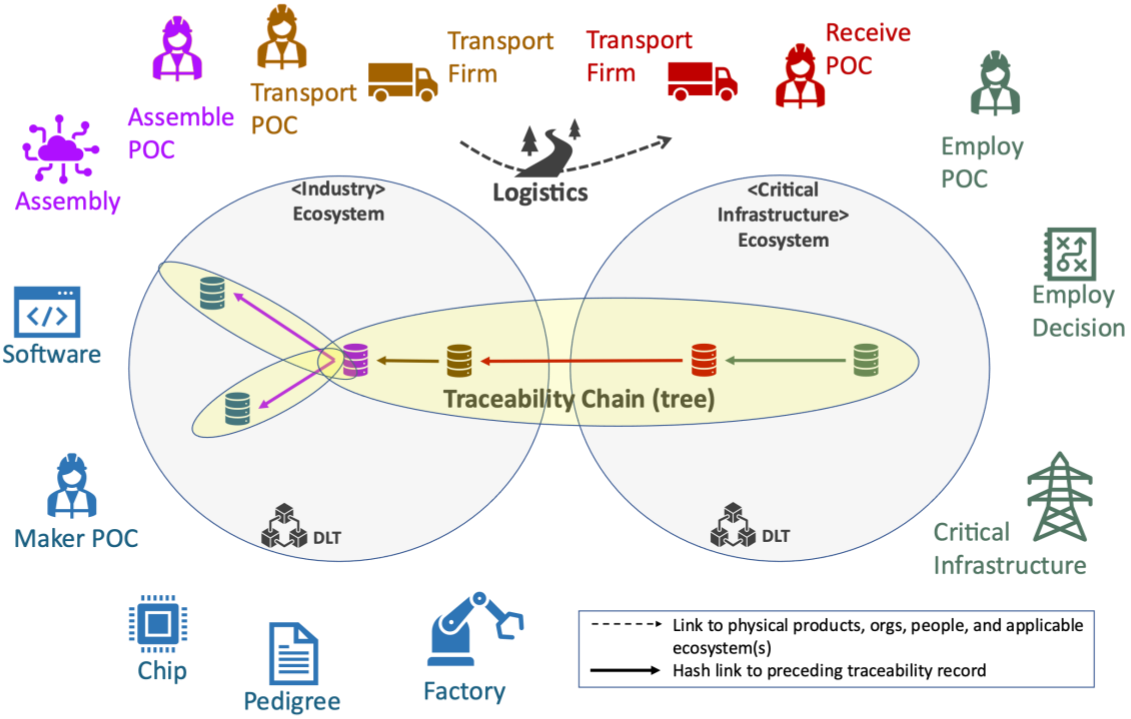

Figure 12: Notional Traceability Chain – Full Chain

The resulting traceability chain is depicted as a singular object, composed of constituent traceability records, which can be read starting at the final receive (or employ) traceability record, and tracing back to the original make records.

Traceability Record Data Types

Traceability records are written as DLT transactions, of which the data types for the DLT transaction data payload are specialized and sub-typed according to use. The traceability DLT transactions are written by an ecosystem authorized actor as the activity or transaction is confirmed to occur and will include any required back linked (hash link) to the preceding traceability record(s). While having the actors directly support each action might be desirable, it may not always be possible. The MVP does not stipulate or assume any requirement on actors who perform actions and those that record transactions to the ecosystem. We recognize the possibility that an organization may limit the number of authorized personnel capable of submitting transactions to the ecosystem. In these situations, the actor submitting the transaction could be different from the actor(s) performing the actions. For example, an organization may have a department that coordinates transportation logistics for products. Once the transportation logistics are coordinated, the relevant information is submitted to the actor authorized to submit the transaction information to the ecosystem as confirmation the activity has been completed. We also recognize that automation is another factor to consider when reviewing these transaction flows. Some organization may implement automation and allow authorized devices to act as the actors for supporting some transaction submissions. For this initial MVP, we are focusing on the allocation of the transaction types to the appropriate ecosystems establishing a foundation that will allow exploration of these different actor scenarios as future research.

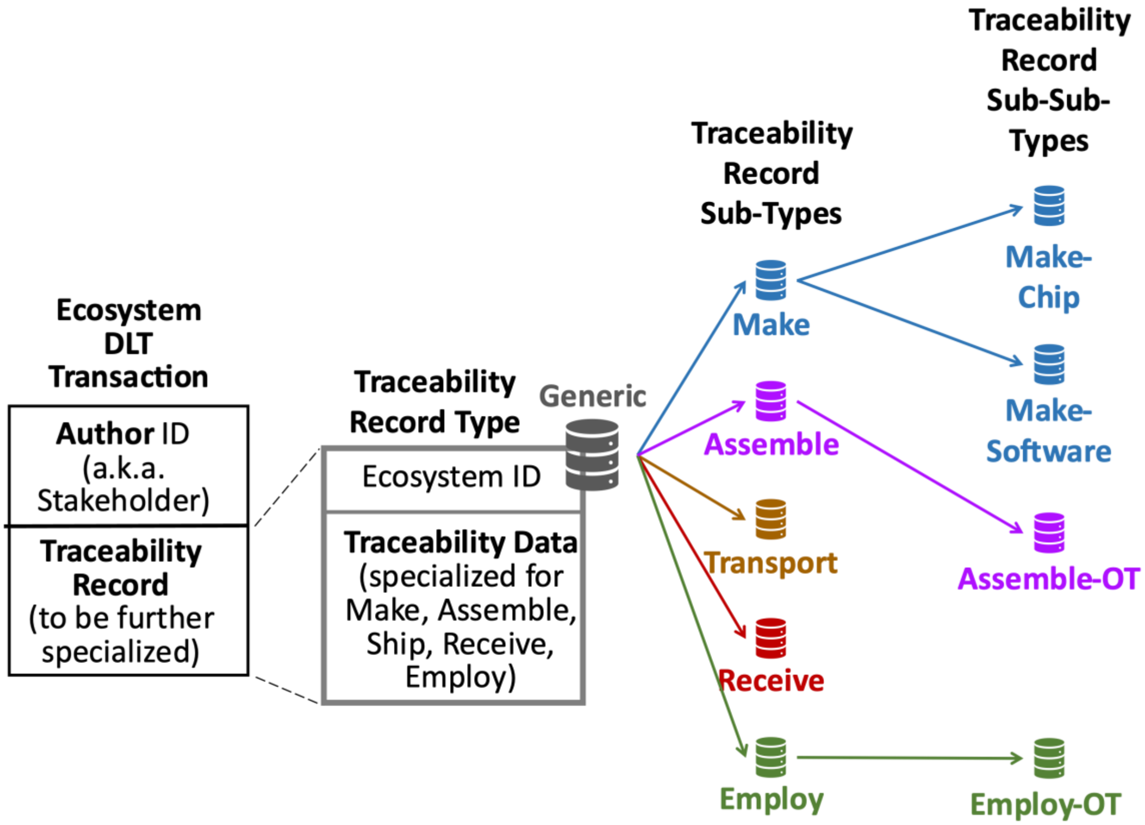
Ecosystem DLT Transaction
Author ID (a.k.a. Stakeholder)
Traceability Record (to be further specialized)
Traceability Record Type
Generic
Ecosystem ID
Traceability Data (specialized for Make, Assemble, Ship, Receive, Employ)
Traceability Record Sub-Types
Make
Assemble
Transport
Receive
Employ
Traceability Record Sub-Sub-Types
Make-Chip
Make-Software
Assemble-OT
Employ-OT

Figure 13: Traceability Data Types

Note that the actor submitting the transaction is also called 'author.' This generic traceability record type is specialized to sub-types based on the activity category (Make, Assembly, Ship, Receive, Employ). Make and Employ sub-types can be further specialized again to sub-sub-types for the specific industry type (e.g., make-chip, make-software). All concrete traceability records for Make and Employ are instances of a sub-sub-type (e.g., Make-Chip, Make-Software) and are extensible to support applicable industry standards. The fields shown for the MVP, represent a subset of metadata believed to be needed to support cross-ecosystem traceability and may either be already part of data standards or may extend data standards for the ecosystems. Additional information such as links to external datasets, SBOM, or other regulatory or mandated information could be incorporated into the sub-type definitions to extend the transactions for supporting different industries.

Another key aspect for Make transactions is the incorporation of an immutable identifier for the item. Establishing and embedding immutable identifiers is beyond the scope for this initial MVP; however, the Make transactions must include these identifiers to support traceability. For Make-Software events, we are implementing multiple values that may be required to uniquely identify each instance of the software. For the MVP, we are utilizing a hash combined with a unique serial number to provide identification and traceability as an initial starting point. We recognize that as industry standards evolve, the metadata for these data types may require adjustment to ensure they properly capture the unique identities for hardware and software defined by industry.

An important note for the MVP is that the Transport and Receive traceability records serve as generic provenance links and are not specialized to relevant industry for this MVP project. This structure of traceability types, sub-types, and sub-sub-types are initial considerations for standards development. The generic sub-types (Make, Assemble, Transport, Receive, Employ) are described in the table below. The Make records below can be further sub-sub-typed as Make-Chip and Make-Software to illustrate the flexibility and extensibility of the Traceability Chain approach.

The MVP scenarios include writing Make-Chip and Make-Software sub-sub-typed traceability records. These records are derived from Make sub-type traceability records. For the MVP, the additional fields in the sub-sub-type

Table 1: Traceability Record Sub-Type Data Fields

Traceability Record Types	Data Fields	Notes
Top level (generic)	• DLT user address • Traceability Record (see below sub-types)	The DLT user address is a public key, derived from the user private key; the user is the relevant stakeholder and an individual (not organization). Decentralized identity standards orgs are working the complex issues regarding organizational identity.
Make Sub-type	• Ecosystem ID (origination) • Factory ID (organization) • Maker POC	Factory is in (origination) ecosystem
Make-Chip Sub-sub-type	• (Extend from Make Sub-type) • Chip Product ID • Chip Pedigree Statement	
Make-Software Sub-sub-type	• (Extend from Make Sub-type) • Software Product ID • Software Pedigree Statement	
Assemble Sub-type	• Ecosystem ID (origination) • Assembly ID • Assemble POC • For each product included in the assembly ○ Ecosystem ID ○ Hash-link to Make traceability record ○ Product ID in Make traceability record	Assemble can refer to assemble / make records in the same ecosystem, and/or receive records from prior ecosystems Assemble traceability records are the branching nodes in the traceability chain/tree
Transport Sub-type	• Ecosystem ID (origination) • Factory ID (origination) • Transport POC • Transport Firm • Ecosystem ID (destination) • Consuming ID (destination organization)	Transport record is in origination ecosystem

Traceability Record Types	Data Fields	Notes
	• Hash-link to Assemble or Make traceability record • Product ID (assemble or simple make)	
Receive	• Ecosystem ID (origination) • Ecosystem ID (destination) • Transport Firm • Receive POC • Hash link to transport record • Product ID (assemble or simple make) • Consuming ID (destination organization)	Receive record is in destination ecosystem
Employ	• Ecosystem ID (final use in critical infrastructure, or equivalent) • Critical Infrastructure (or equivalent) ID • Employ POC • Hash link to receive record • Product ID (assemble or simple make) • Link to employ decision	The employ decision is the document which summarizes the decision to use the product, and where in the critical infrastructure (or equivalent) the product is used.

Cybersecurity Factors

The MVP is primarily concerned with the security and integrity of the overall traceability chain. The security and integrity of the traceability chain is predicated on the security and integrity of each traceability record written to the data stores of each respective ecosystem. The MVP focuses on two of the three pillars of cybersecurity: confidentiality (permissioned based access requirements), and integrity (transaction immutability). The MVP seeks data stores such as DLT (including blockchain) which provide tamper-evident properties supporting data integrity. Further, each ecosystem implements and protects its own instance of DLT, providing confidentiality.

Architectural Notes

MVP Project

The MVP project includes many technical aspects of supply chain, data, and identity technology. We anticipate that multiple industry contributors will be required to implement the MVP. This project also assumes notional agreement around simplified traceability data types, which in a real industry sector adoption would be subject to negotiation and agreement, the same as any shared data standard, and likely would vary in particular supply chains and/or different industry sectors and sub-sectors.

Ecosystems

The MVP will implement specific manufacturing and critical infrastructure domains: (a) microelectronic chip manufacturers, (b) operational technology manufacturers, and (c) critical infrastructure. While the concepts are illustrated in the MVP using DLT (including blockchain) and a specific set of suppliers and infrastructure, the concepts can be applied to other manufacturing supply chain domains and critical infrastructures. Ultimately, the stakeholders and participants in individual supply chains will determine which entities become members of their particular ecosystem. Further, stakeholders can participate in multiple ecosystems, for example a large company with multiple divisions each participating in ecosystems of in disparate manufacturing domains. In addition to membership, ecosystem stakeholders and participants mutually will determine the composition of standards and data elements that constitute their distributed ledger, as necessary to accomplish an MVP ecosystem for their industry sector or sub-sector.

Ecosystem Stakeholders and Identity

MVP manufacturers and critical infrastructure operators are stakeholders of their respective manufacturing ecosystems. Each stakeholder has an identity which is unique across the MVP. For example, a critical infrastructure operator who has previously accessed a traceability record, can understand the identity of a microelectronic or operational technology ecosystem, and the manufacturer stakeholder, who wrote the traceability record. Accessing a traceability record within an ecosystem is performed by providing the hash link to the traceability record to the query facility of the respective ecosystem, as simplified data access management for this MVP. For example, a power plant operator will accept the shipment of an OT assembly, and in parallel accept the corresponding transport traceability record for the OT assembly, writing a receive traceability record to acknowledge. This receive traceability record contains links to the preceding ecosystem and transport traceability record, which can be used to follow the traceability chain in reverse. This constraint simplifies the data access management aspect of the MVP implementation.

Identity Technology and Standards

Identity standards are currently being developed with important progress in the W3C suite of decentralized Identity specifications. There are open questions about what the manufacturing supply chain traceability ecosystem identity standards should be in the future. This MVP is intended to be a foundational starting point for refinement of future manufacturing supply chain traceability ecosystem identity standards. The section High Level Architecture above discusses a simple role-based identity scheme for use in this MVP project, intended to be supplanted by identity standards, both individual and organizational, as they become available.

Ecosystem Operations

The MVP illustrates select aspects of writing and reading manufacturing supply chain traceability records. A full implementation will include additional features, governance, and operational models that will leverage the specific blockchain related technologies being used. NIST IR 8419 [1] describes industry case studies which include an example where the ecosystem is operated by a consortium (e.g., Mediledger, pharma industry) where the consortium uses a third-party company to build and operate the ecosystem DLT and related code. Other operating models are possible, and beyond the scope of the MVP.

Distributed Ledger Technology

Each MVP ecosystem (manufacturing and critical infrastructure) will include an instance of permissioned DLT independent from the other ecosystem DLTs (no sharing of DLT or blockchain implementation across ecosystems). Beyond that, there is no requirement to employ a specific type of DLT or blockchain other than to use a type of permissioned blockchain technology or DLT with similar tamper evident properties. Recommendation to keep the MVP simplified is to use the same type of DLT or blockchain technology in each instance of ecosystem DLT. Note that smart contracts are optional for the MVP.

Transaction Data

The traceability record data in the MVP ecosystem DLTs will be notional and representative of industry domain traceability data; however, it will not be based on specific standards (see “Data Standards” below) in order to facilitate rapid implementation. The new concept in the MVP is the mechanism to create and read a traceability chain (tree) across manufacturing ecosystems.

The MVP DLT transaction data (traceability records) is intended to be minimal in size and complexity. The transaction data can include notional pointers to manufacturer's private manufacturing data to indicate that a critical infrastructure operator could, if mutually agreed, use the traceability data to access internal manufacturer process data. Access to the private manufacturing data is controlled by the manufacturer, is expected to be negotiated with purchasers (other suppliers and critical infrastructure operators) and is not written to the ecosystem DLT. This notional pointer can be used in scenarios below to illustrate anticipated real world forensic activities to verify authenticity in certain traceability use cases.

Data Standards

This MVP is intended to be a foundational starting point for refinement of future manufacturing supply chain traceability ecosystem data standards. Subsequent refinements to the MVP could incorporate future traceability record standards, specific to each industry. The section High-Level Architecture above discusses a set of notional traceability record data types for use in this project.

Integration

This MVP includes integration as well as technology. This MVP is a starting point for researching and demonstrating cross manufacturing supply chain exchange of traceability information. Future research could explore data and identity standards, and different modes of organizing and governing ecosystems.

Supply Chain Actions beyond End Customer Employ

This MVP is focused on the manufacturing supply chain actions (Make, Assemble, Transport, and Receive) prior to and including end customer disposition (Employ) and recorded in the respective traceability records. As shown below, there are opportunities for future research in exploring what may be linked to the Employ record after initial deployment, as a result of sustainment chain activities (Repair, Update, etc.).

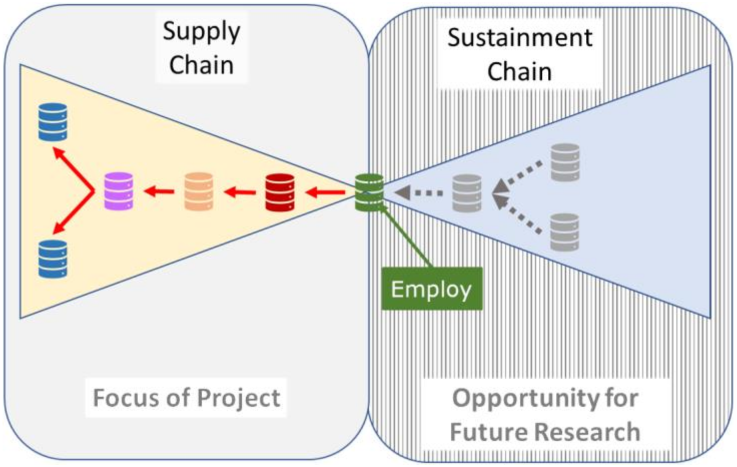

Figure 14: Sustainment Chain Opportunity for Future Research

Further, there are opportunities for future research in connecting other data to the Employ records, such as enterprise IT integration, operational data, and more. These additional data contexts can provide useful insights during forensic investigations which use the Traceability Chain for historical data.

Component List

All components below are intended to be implemented in software and data (not physical components).

- MEP Ecosystem
  - Instance of DLT or other blockchain related technology (can be the same technology across ecosystems)
  - Instance of query facility (can be the same technology across ecosystems)
  - Stakeholders (e.g., MEP-001), each with MVP-wide unique identity
  - Chips, each with unique identity, synthetic factory pedigree data
- OT Ecosystem
  - Instance of DLT or other blockchain related technology (can be the same technology across ecosystems)
  - Instance of query facility (can be the same technology across ecosystems)
  - Stakeholders (e.g., OT-001), each with MVP-wide unique identity
  - Software, each with unique identity, synthetic factory pedigree data
  - Assemblies (chip + software + [optional: sensors, mechanical device]), each with unique identity, synthetic factory pedigree data
- CI Ecosystem
  - Instance of DLT or other blockchain related technology (can be the same technology across ecosystems)
  - Instance of query facility (can be the same technology across ecosystems)
  - Stakeholders (e.g., CI-001), each with MVP-wide unique identity
  - Critical infrastructure, each with unique identity, synthetic pedigree data
- MVP Dashboard with functions:
  - Initialize (clear data)
  - Scenario 1, execute scenario, display activity, save results
  - Scenario 2, execute scenario, display activity, save results
  - Scenario 3, execute scenario, display activity, save results

MVP Requirements

1. Create ecosystems and actors per the Component List above and in concordance with the high-level architecture.
2. Create data types per **Error! Reference source not found.** above.
3. Execute scenarios per the Scenario section above and capture results.

4 Relevant Standards and Guidance

The Traceability Chain MVP does not implement specific supply chain data standards since the MVP is intended to be illustrative not proscriptive. Nonetheless, some standards were explored to inform project team discussions.

List of standards used for this project:

Table 2: Standards and Guidance

Standards Body	Nomenclature	Name
Global Semiconductor Alliance	WP-19	Using a Virtual Identifier Thread for Root of Trust and Reliability

Standards Body	Nomenclature	Name
SEMI	SEMI T20	SEMI T20 - Specification for Authentication of Semiconductors and Related Products
IETF	IETF SCITT (chartered)	Supply Chain Integrity, Transparency, and Trust

5 Security Control Map

This table maps the characteristics of the commercial products that the NCCoE will apply to this cybersecurity challenge to the applicable standards and best practices described in the Framework for Improving Critical Infrastructure Cybersecurity, and to other NIST activities. This exercise is meant to demonstrate the real-world applicability of standards and best practices but does not imply that products with these characteristics will meet an industry's requirements for regulatory approval or accreditation.

Table 3: Security Control Map

Cybersecurity Framework v1.1			SP 800-53 R5
Function	Category	Subcategory	
Identify (ID)	Supply Chain Risk Management (ID.SC)	ID.SC-3: Contracts with suppliers and third-party partners are used to implement appropriate measures designed to meet the objectives of an organization's cybersecurity program and Cyber Supply Chain Risk Management Plan.	SA-9, SA-11, SA-12, PM-9 SR-6
		ID.SC-4: Suppliers and third-party partners are routinely assessed using audits, test results, or other forms of evaluations to confirm they are meeting their contractual obligations.	AU-2, AU-6, AU-12, AU-16, PS-7. SA-9, SA-12 SR-6
	Asset Management (ID.AM)	ID.AM-1: Physical devices and systems within the organization are inventoried	CM-8, PM-5
		ID.AM-2: Software platforms and applications within the organization are inventoried	CM-8, PM-5
Protect (PR)	Identity Management, Authentication, and Access Control (PR.AC)	PR.AC-6: Identities are proofed and bound to credentials and asserted in interactions	AC-1, AC-2, AC-3, AC-16, AC-19, AC-24, IA-1, IA-2, IA-4, IA-5, IA-8, PE-2, PS-3
	Data Security (PR.DS)	PR.DS-1: Data-at-rest is protected	MP-8, SC-12, SC-28

Cybersecurity Framework v1.1			
Function	Category	Subcategory	SP 800-53 R5
		PR.DS-6: Integrity checking mechanisms are used to verify software, firmware, and information integrity	SC-16, SI-7
		PR.DS-8: Integrity checking mechanisms are used to verify hardware integrity	CM-2
Detect (DE)	Detection Processes (DE.DP)	DE.DP-2: Detection activities comply with all applicable requirements	AC-25, CA-2, CA-7, SA-18, SI-4, PM-14
NA	NA	NA	SR-4
NA	NA	NA	SR-7
NA	NA	NA	SR-11

Appendix A References

[1] K. Stouffer, M. Pease, J. Lubell, E. Wallace, H. Reed, V. Martin, S. Granata, A. Noh and C. Freeberg, "Blockchain and Related Technologies to Support Manufacturing Supply Chain Traceability: Needs and Industry Perspectives," National Institute of Standards and Technology, Gaithersburg, MD, 2022. Available: https://doi.org/10.6028/NIST.IR.8419

[2] "Supply Chain Integrity, Transparency, and Trust (scitt)," Internet Engineering Task Force (IETF), [Online]. Available: https://datatracker.ietf.org/wg/scitt/about/. [Accessed 3 April 2023].

[3] "Decentralized Identifiers (DIDs) v1.0," World Wide Web Consortium (W3C) , [Online]. Available: https://www.w3.org/TR/did-core/. [Accessed 3 April 2023].

[4] "Trusted IoT Ecosystem Security (TIES)," Global Semiconductor Alliance (GSA), [Online]. Available: https://www.gsaglobal.org/iot/ties/. [Accessed 3 April 2023].

[5] "Supply Chain Traceability," MIT Sustainable Supply Chain Lab, [Online]. Available: https://sustainable.mit.edu/supply-chain-traceability/. [Accessed 8 March 2023].

[6] "Information and communications Technology Supply Chain Risk Management," DHS CISA, [Online]. Available: https://www.cisa.gov/information-and-communications-technology-supply-chain-risk-management. [Accessed 8 March 2023].

[7] "Supply Chain Assurance," National Institute of Standards (NIST) National Cybersecurity Center of Excellence (NCCoE), [Online]. Available: https://www.nccoe.nist.gov/supply-chain-assurance. [Accessed 8 March 2023].

Appendix B Acronyms and Abbreviations

API	Application Program Interface
CI	Critical Infrastructure
DID	Decentralized Identifier
DLT	Distributed Ledger Technologies
MVP	Minimum Viable Product
NCCoE	National Cybersecurity Center of Excellence
NIST	National Institute of Standards and Technology
OT	Operational Technology
POC	Point of Contact
RI	Reference Implementation
SBOM	Software Bill of Materials
W3C	World Wide Web Consortium