US 20230237290A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2023/0237290 A1**

YANNUZZI et al. (43) **Pub. Date: Jul. 27, 2023**

(54) **RELABELING-RESISTANT ITEM IDENTIFIERS USED ACROSS DOMAINS**

(71) Applicant: **Cisco Technology, Inc.**, San Jose, CA (US)

(72) Inventors: **Marcelo YANNUZZI**, Vufflens-La-Ville (CH); **John MONAGHAN**, West Dunbartonshire (GB); **Joel Abraham OBSTFELD**, Bushey (GB)

(21) Appl. No.: **17/583,444**

(22) Filed: **Jan. 25, 2022**

Publication Classification

(51) **Int. Cl.**

G06K 7/10	(2006.01)
G06Q 10/08	(2006.01)
G06K 19/06	(2006.01)

(52) **U.S. Cl.**

CPC ***G06K 7/10861*** (2013.01); ***G06Q 10/087*** (2013.01); ***G06K 19/06037*** (2013.01)

(57) **ABSTRACT**

According to one or more embodiments of the disclosure, a device makes a determination that an asset identification system has been deployed at an organization. The device automatically deploys, based on the determination, a connector for the asset identification system to capture labeling information indicative of transport of one or more physical assets via the organization. The device receives, from the connector, a binding request tag notification indicative of transport of a particular physical asset of the one or more physical assets via the organization. The device associates the binding request tag notification with one or more identifiers for the particular physical asset used by one or more other organizations.

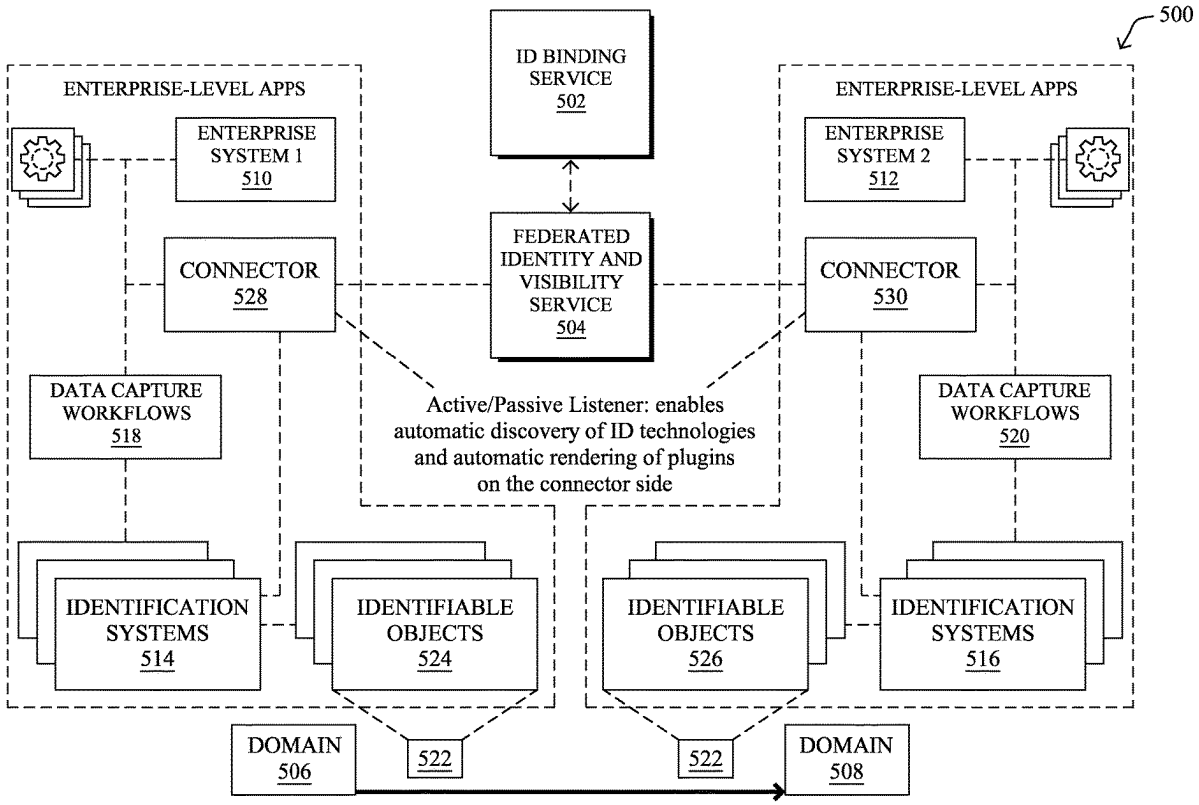

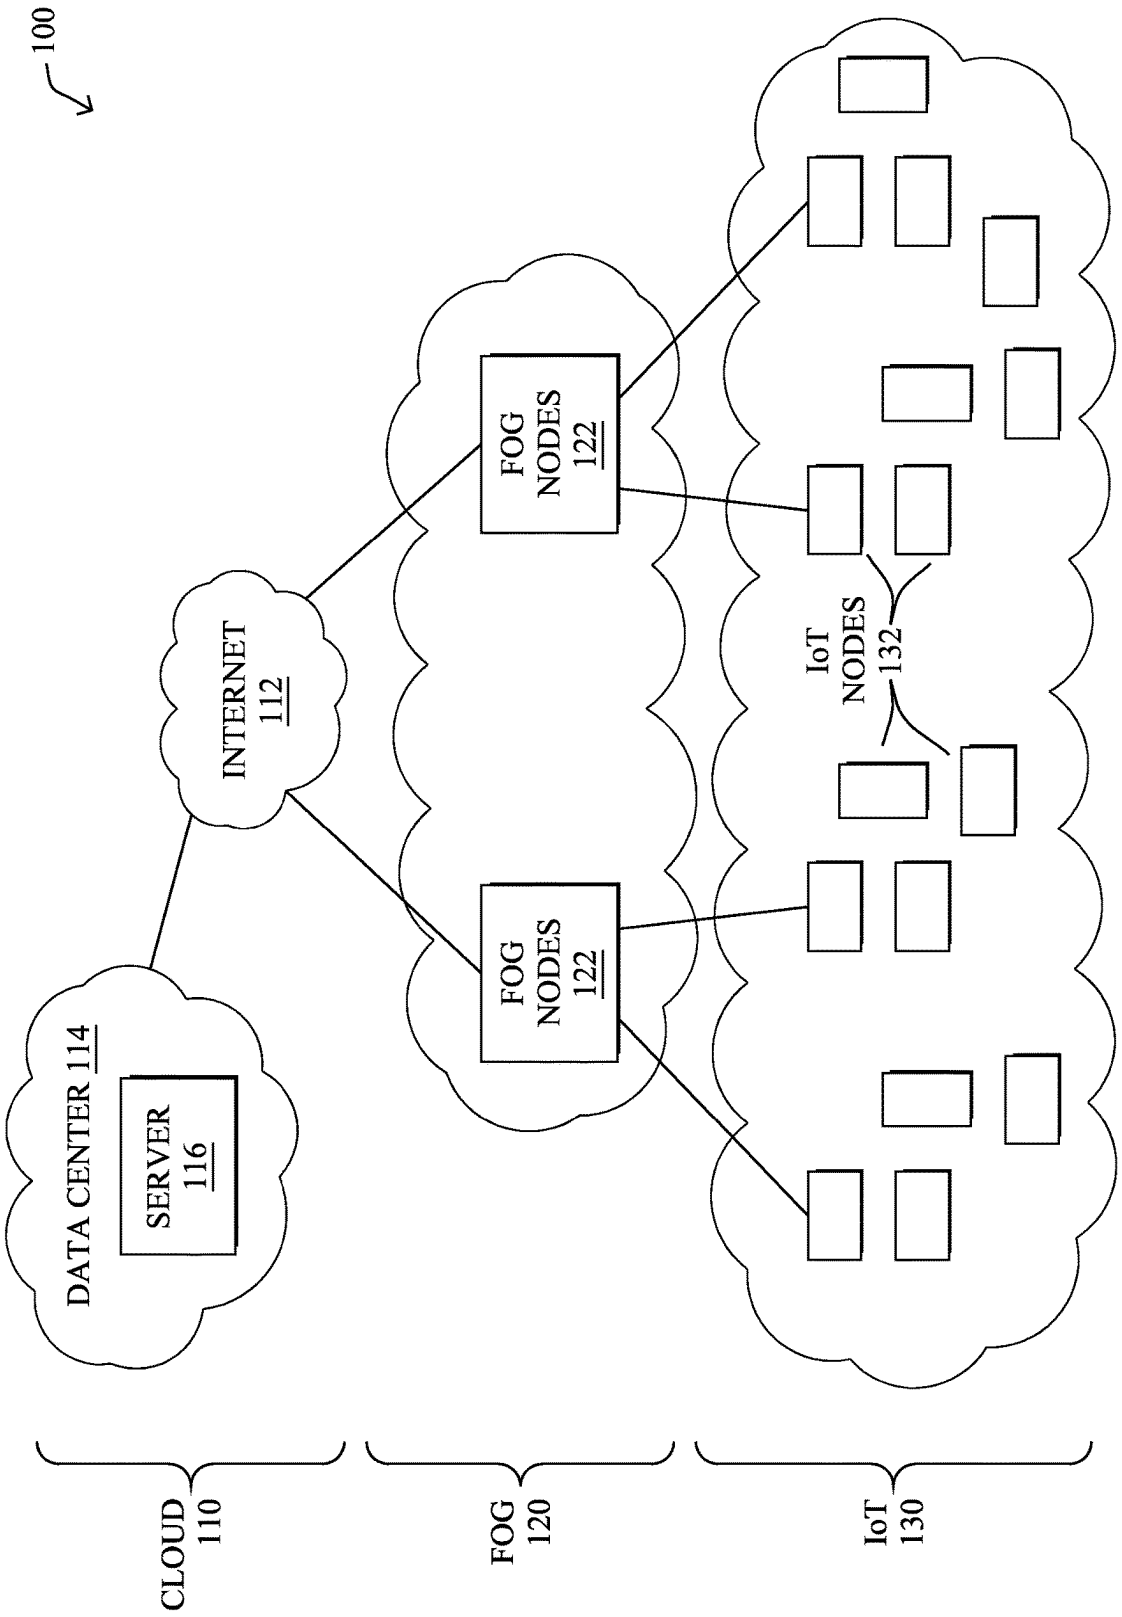
100
DATA CENTER 114
SERVER 116
INTERNET 112
FOG NODES 122
FOG NODES 122
IoT NODES 132
CLOUD 110
FOG 120
IoT 130

FIG. 1

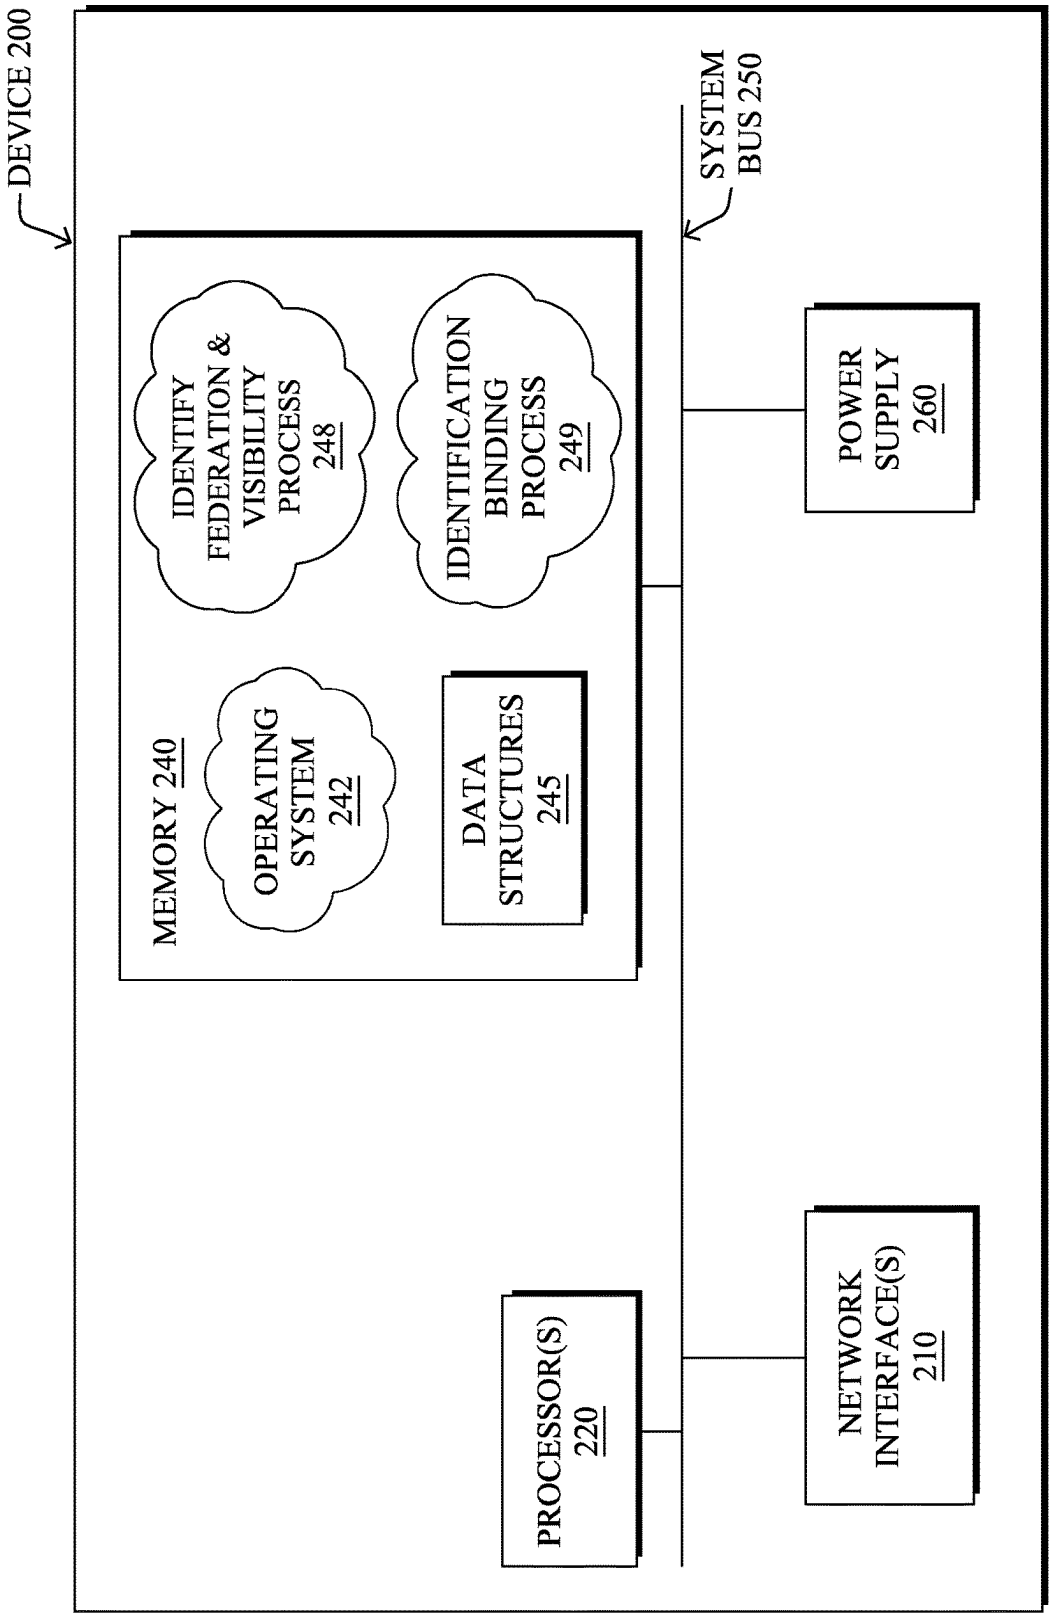
DEVICE 200
SYSTEM BUS 250
MEMORY 240
OPERATING SYSTEM 242
DATA STRUCTURES 245
IDENTIFY FEDERATION & VISIBILITY PROCESS 248
IDENTIFICATION BINDING PROCESS 249
POWER SUPPLY 260
PROCESSOR(S) 220
NETWORK INTERFACE(S) 210

FIG. 2

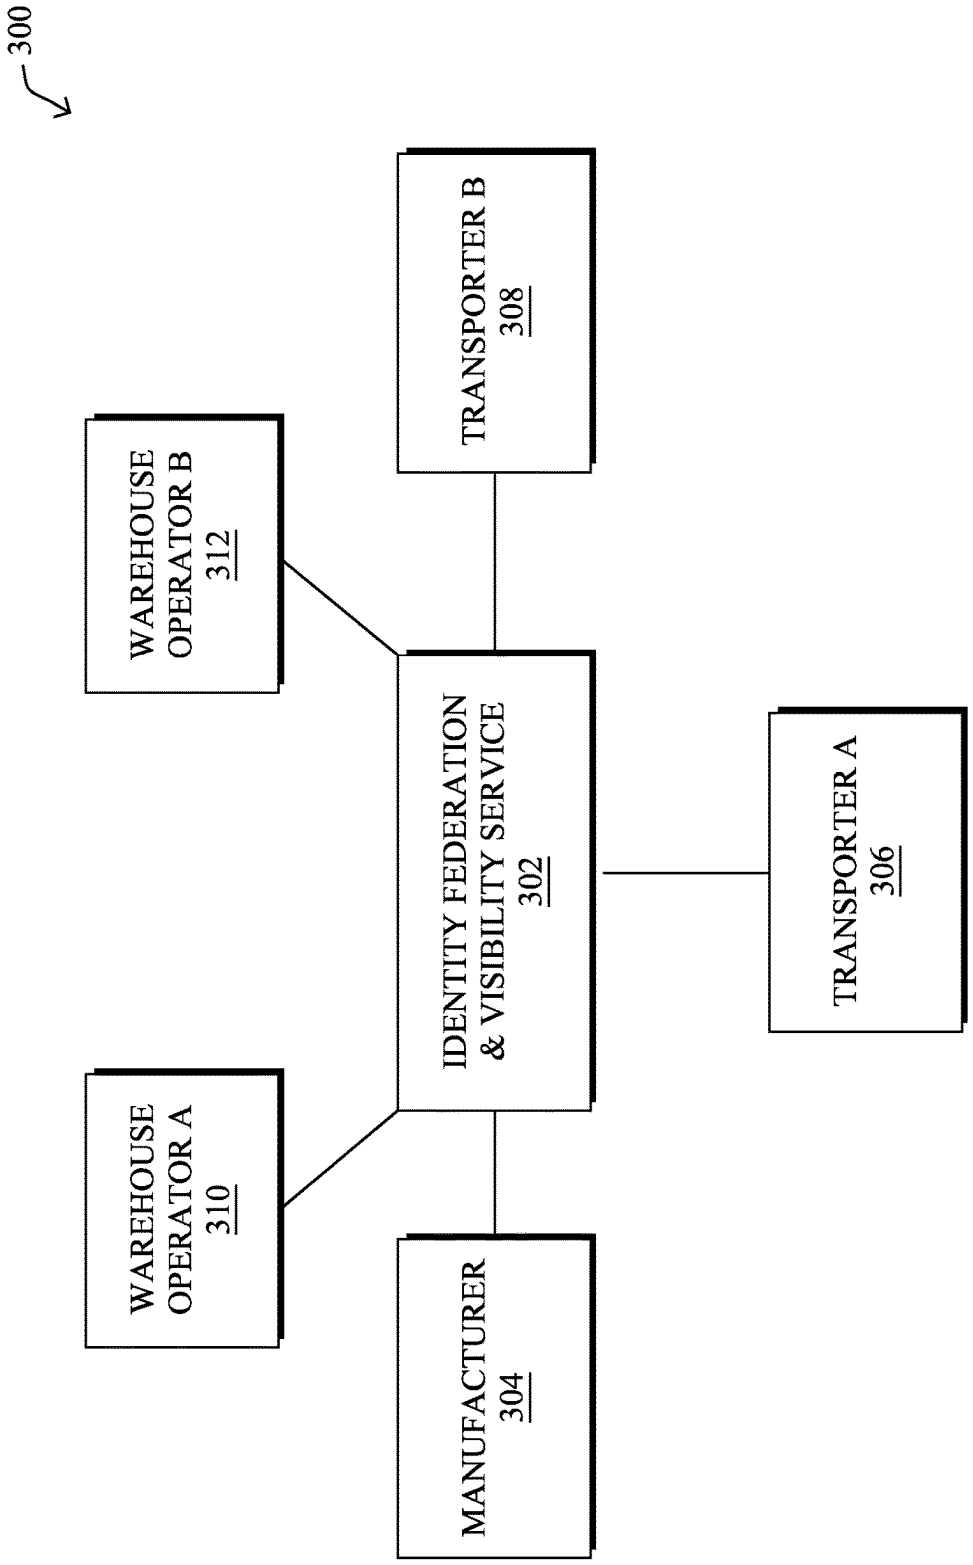
300
TRANSPORTER B
308
WAREHOUSE
OPERATOR B
312
IDENTITY FEDERATION
& VISIBILITY SERVICE
302
TRANSPORTER A
306
WAREHOUSE
OPERATOR A
310
MANUFACTURER
304

FIG. 3

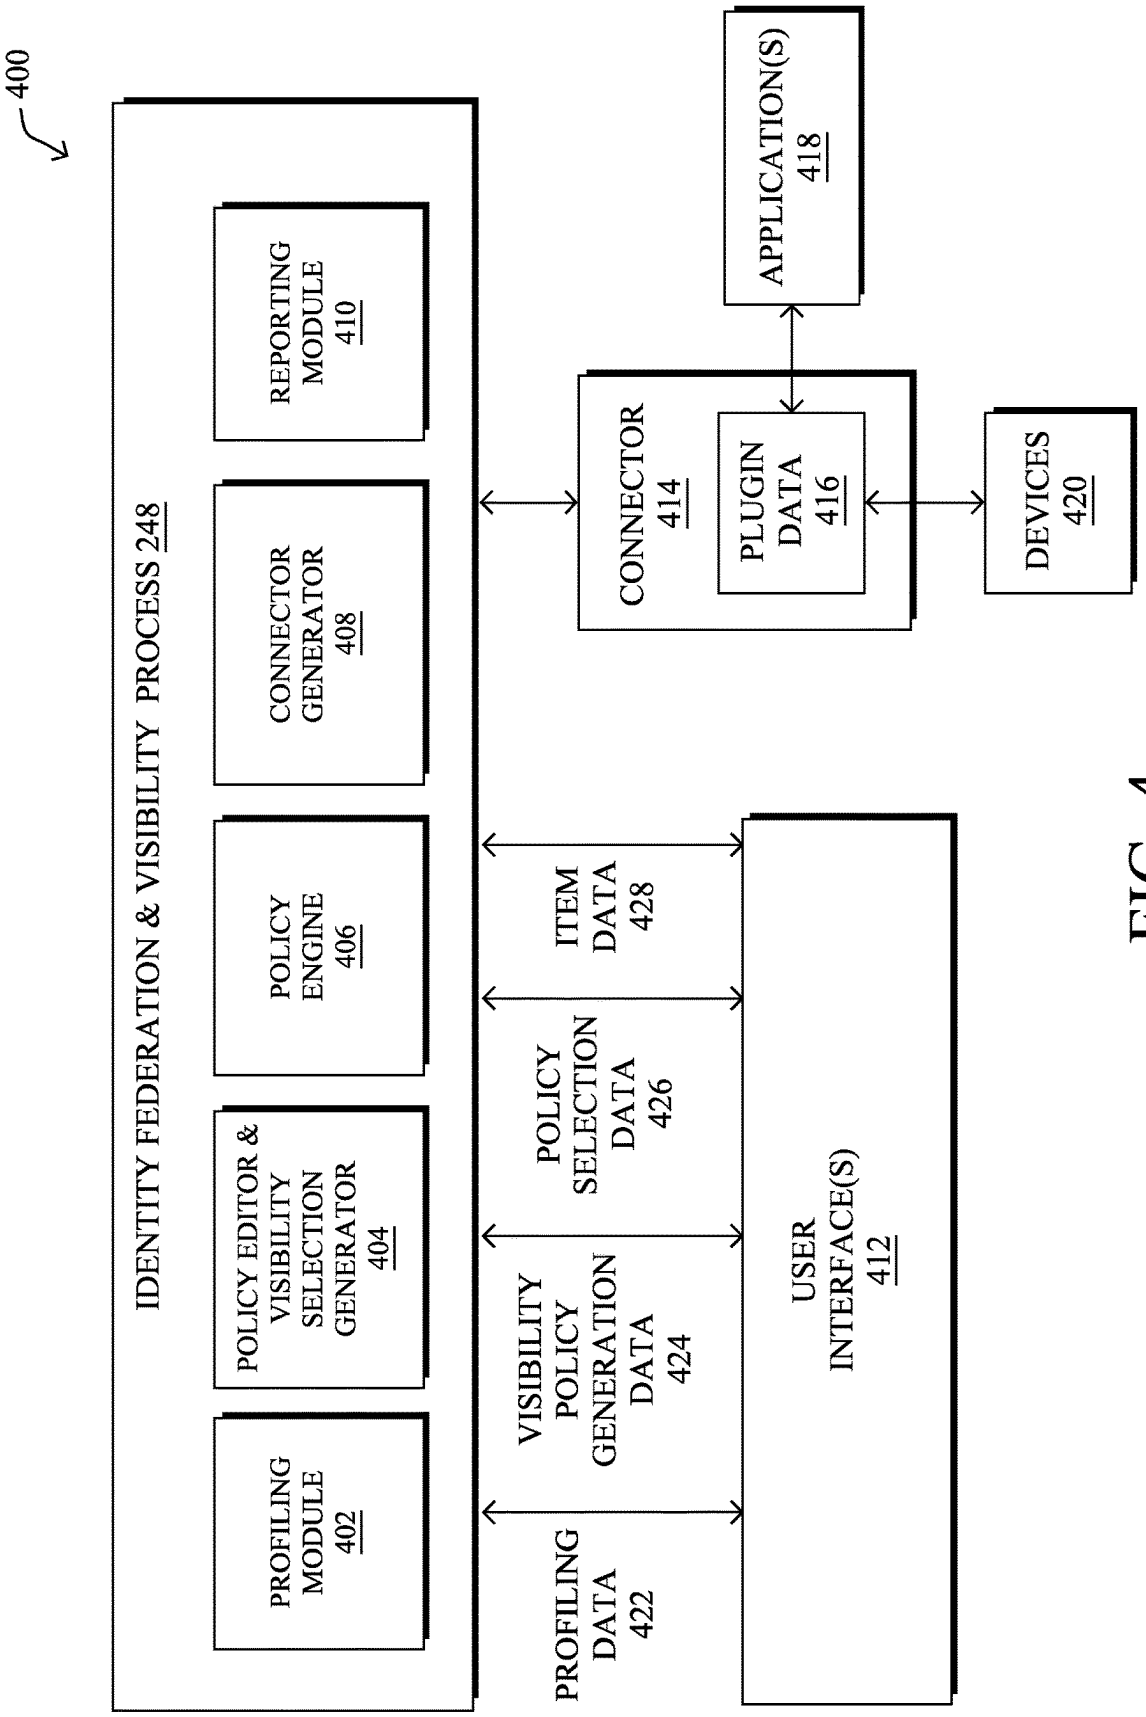
400
IDENTITY FEDERATION & VISIBILITY PROCESS 248
PROFILING MODULE 402
POLICY EDITOR & VISIBILITY SELECTION GENERATOR 404
POLICY ENGINE 406
CONNECTOR GENERATOR 408
REPORTING MODULE 410
PROFILING DATA 422
VISIBILITY POLICY GENERATION DATA 424
POLICY SELECTION DATA 426
ITEM DATA 428
USER INTERFACE(S) 412
CONNECTOR 414
PLUGIN DATA 416
APPLICATION(S) 418
DEVICES 420

FIG. 4

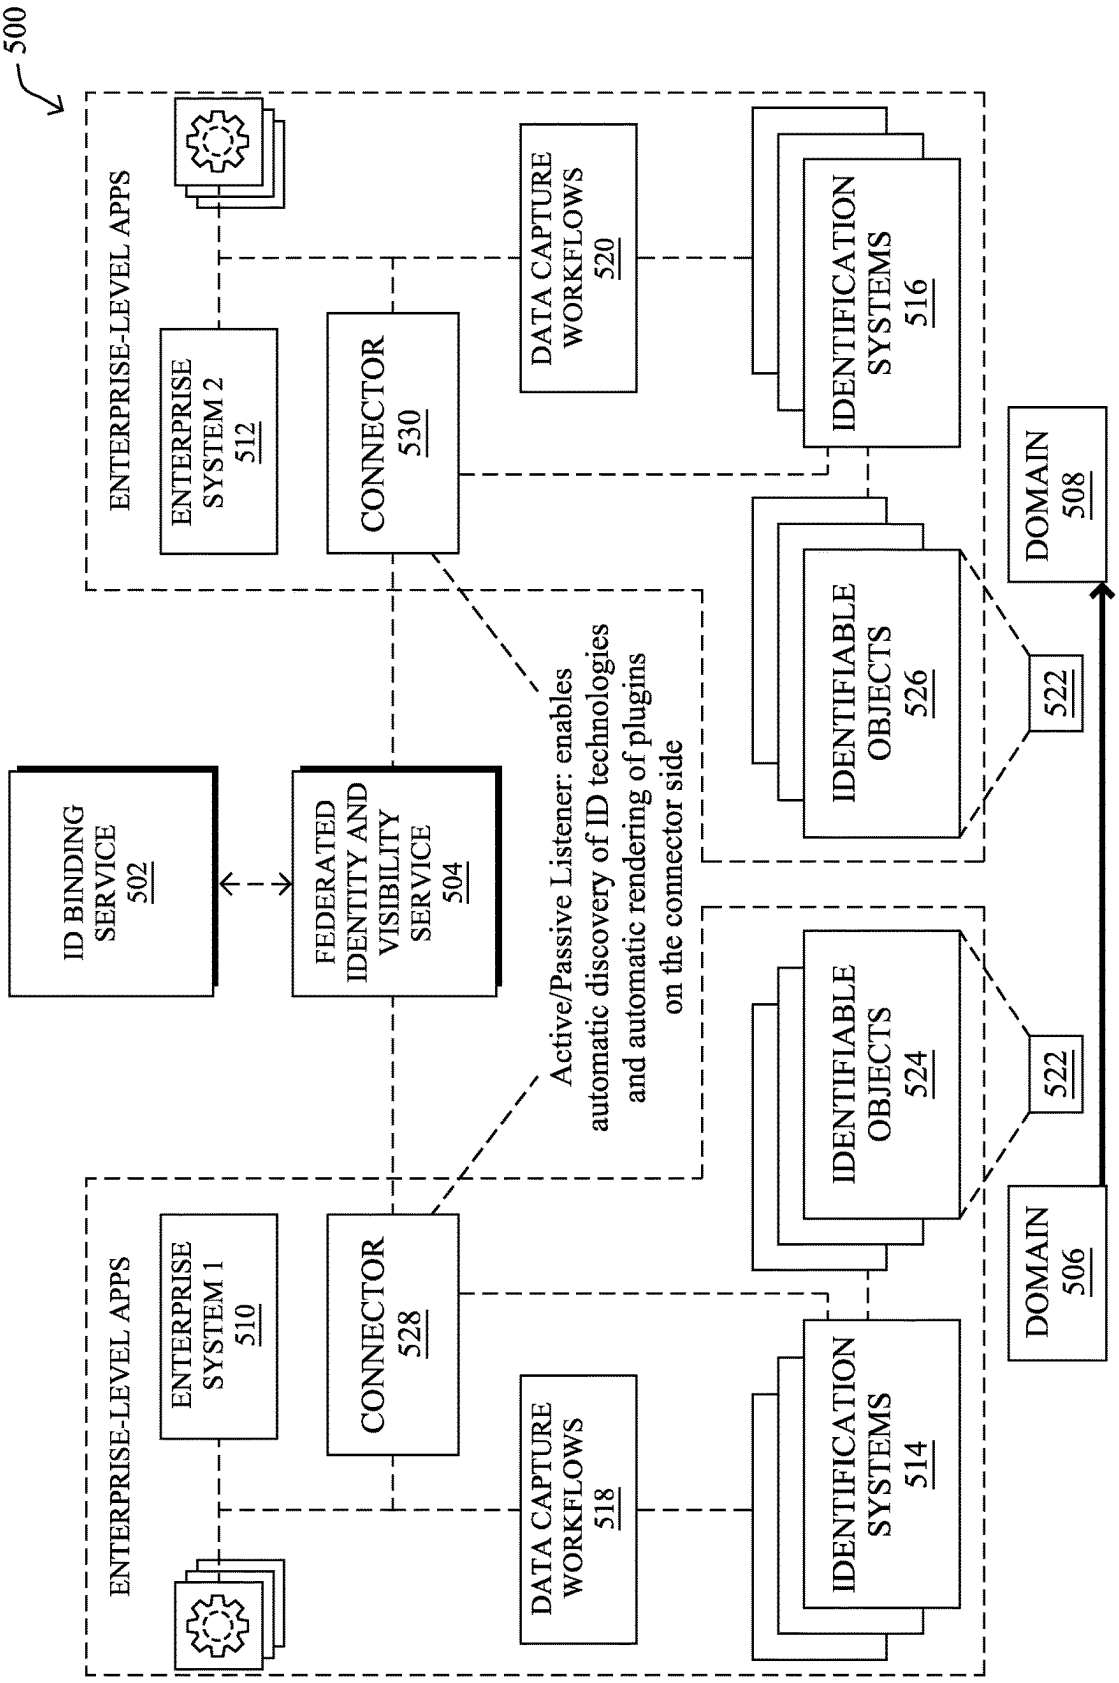
500
ENTERPRISE-LEVEL APPS
ENTERPRISE SYSTEM 2 512
CONNECTOR 530
DATA CAPTURE WORKFLOWS 520
IDENTIFICATION SYSTEMS 516
IDENTIFIABLE OBJECTS 526
DOMAIN 508
522
ID BINDING SERVICE 502
FEDERATED IDENTITY AND VISIBILITY SERVICE 504
Active/Passive Listener: enables automatic discovery of ID technologies and automatic rendering of plugins on the connector side
ENTERPRISE-LEVEL APPS
ENTERPRISE SYSTEM 1 510
CONNECTOR 528
DATA CAPTURE WORKFLOWS 518
IDENTIFICATION SYSTEMS 514
IDENTIFIABLE OBJECTS 524
522
DOMAIN 506

FIG. 5

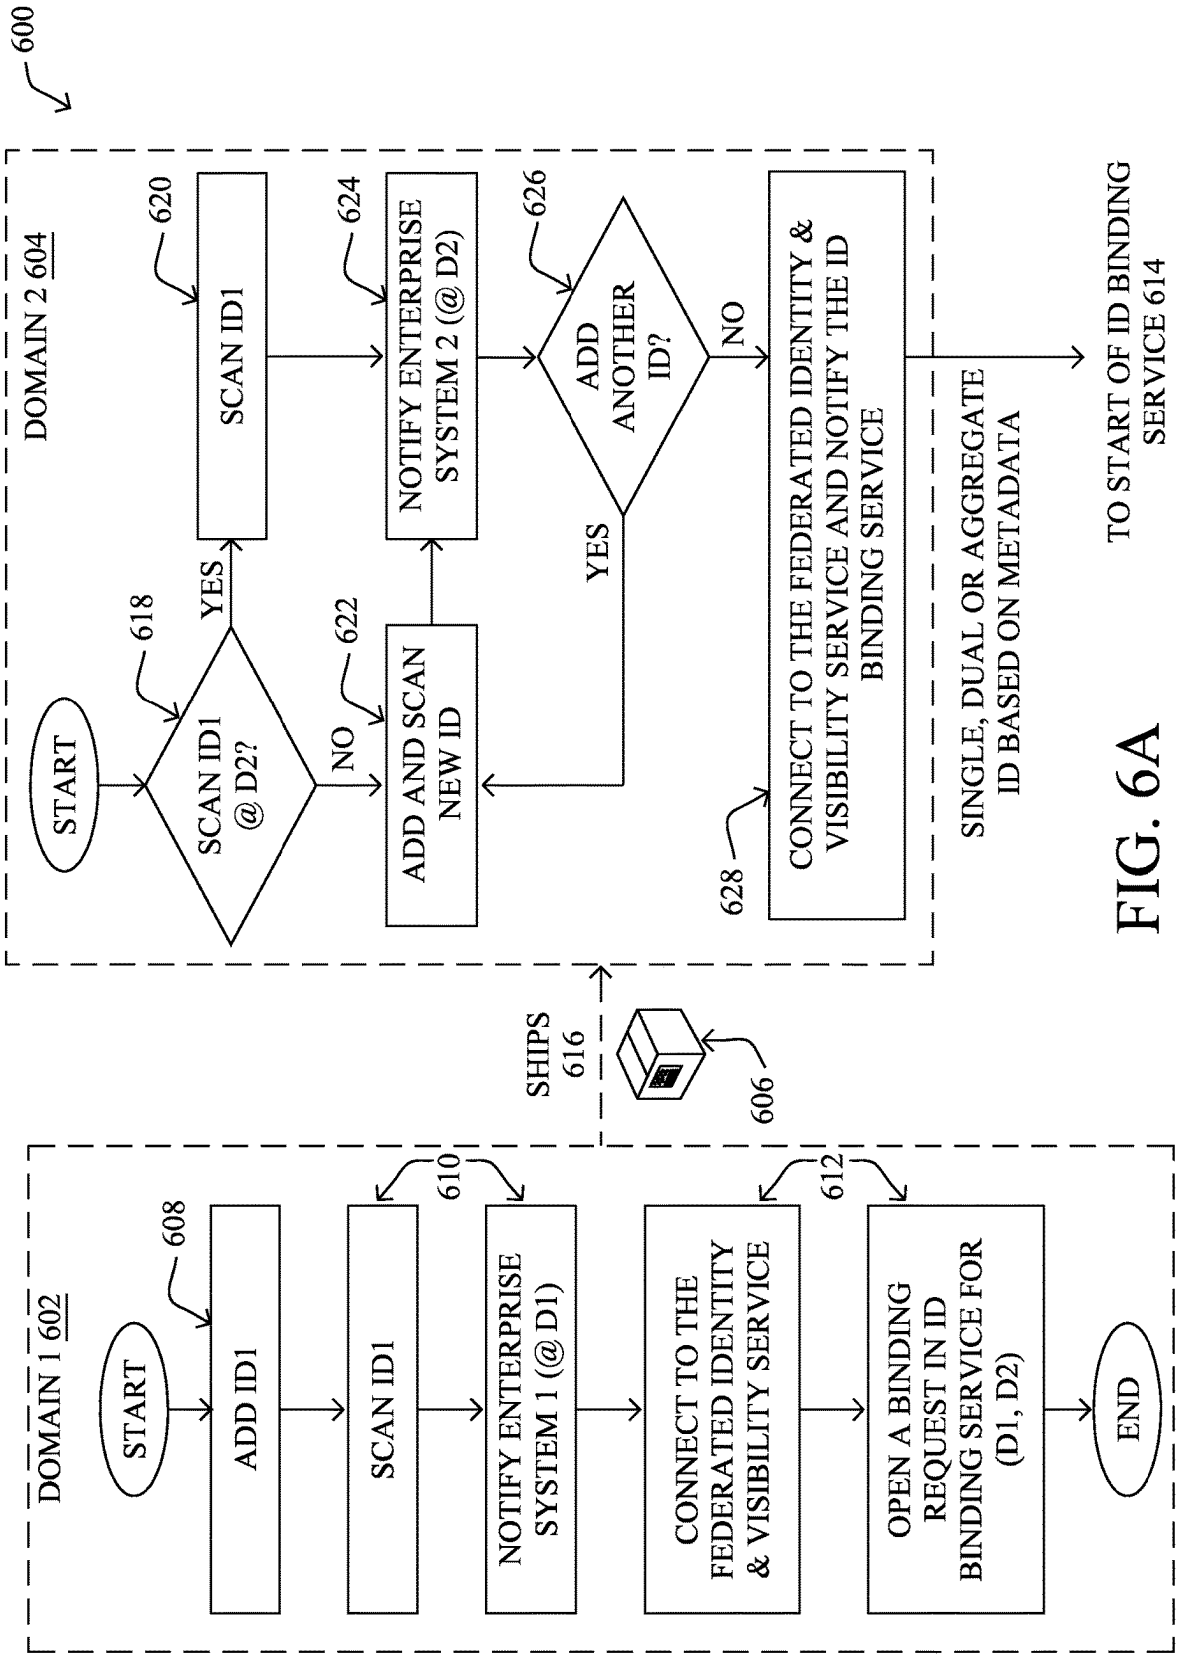
600
DOMAIN 1 602
START
608
ADD ID1
SCAN ID1
610
NOTIFY ENTERPRISE SYSTEM 1 (@ D1)
CONNECT TO THE FEDERATED IDENTITY & VISIBILITY SERVICE
612
OPEN A BINDING REQUEST IN ID BINDING SERVICE FOR (D1, D2)
END
SHIPS 616
606
DOMAIN 2 604
START
618
SCAN ID1 @ D2?
YES
620
SCAN ID1
NO
622
ADD AND SCAN NEW ID
624
NOTIFY ENTERPRISE SYSTEM 2 (@ D2)
626
ADD ANOTHER ID?
YES
NO
628
CONNECT TO THE FEDERATED IDENTITY & VISIBILITY SERVICE AND NOTIFY THE ID BINDING SERVICE
SINGLE, DUAL OR AGGREGATE ID BASED ON METADATA
TO START OF ID BINDING SERVICE 614

FIG. 6A

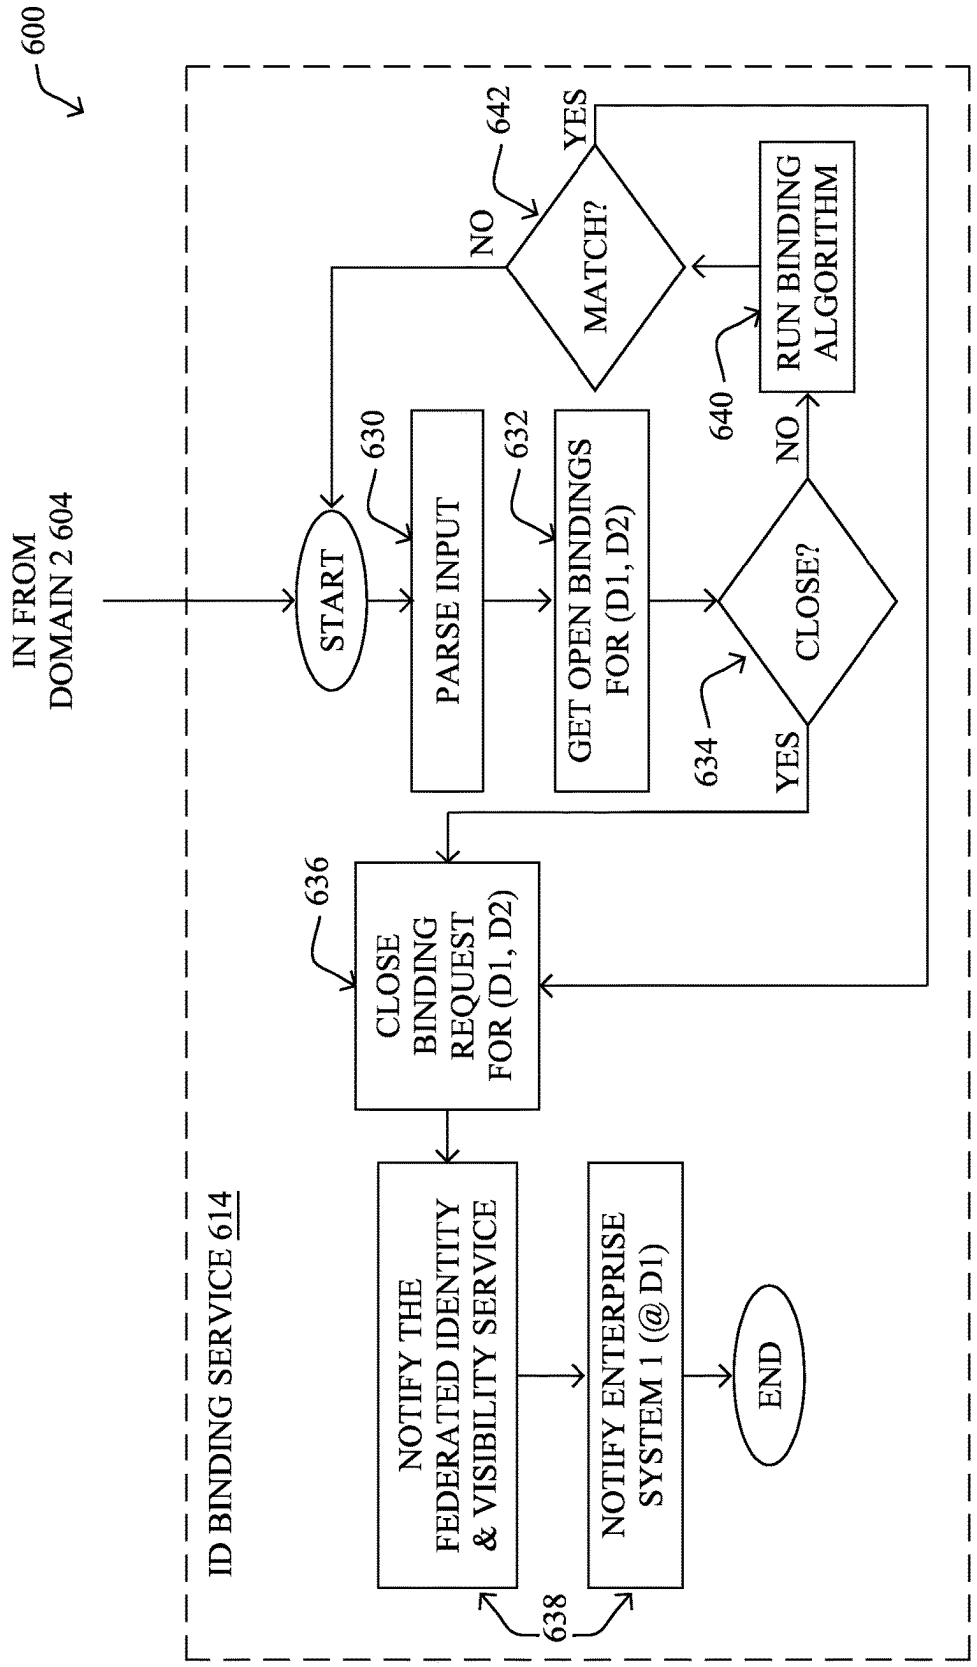
600
ID BINDING SERVICE 614
IN FROM DOMAIN 2 604
START
630
PARSE INPUT
632
GET OPEN BINDINGS FOR (D1, D2)
634
CLOSE?
NO
YES
640
RUN BINDING ALGORITHM
642
MATCH?
NO
YES
636
CLOSE BINDING REQUEST FOR (D1, D2)
638
NOTIFY THE FEDERATED IDENTITY & VISIBILITY SERVICE
NOTIFY ENTERPRISE SYSTEM 1 (@ D1)
END

FIG. 6B

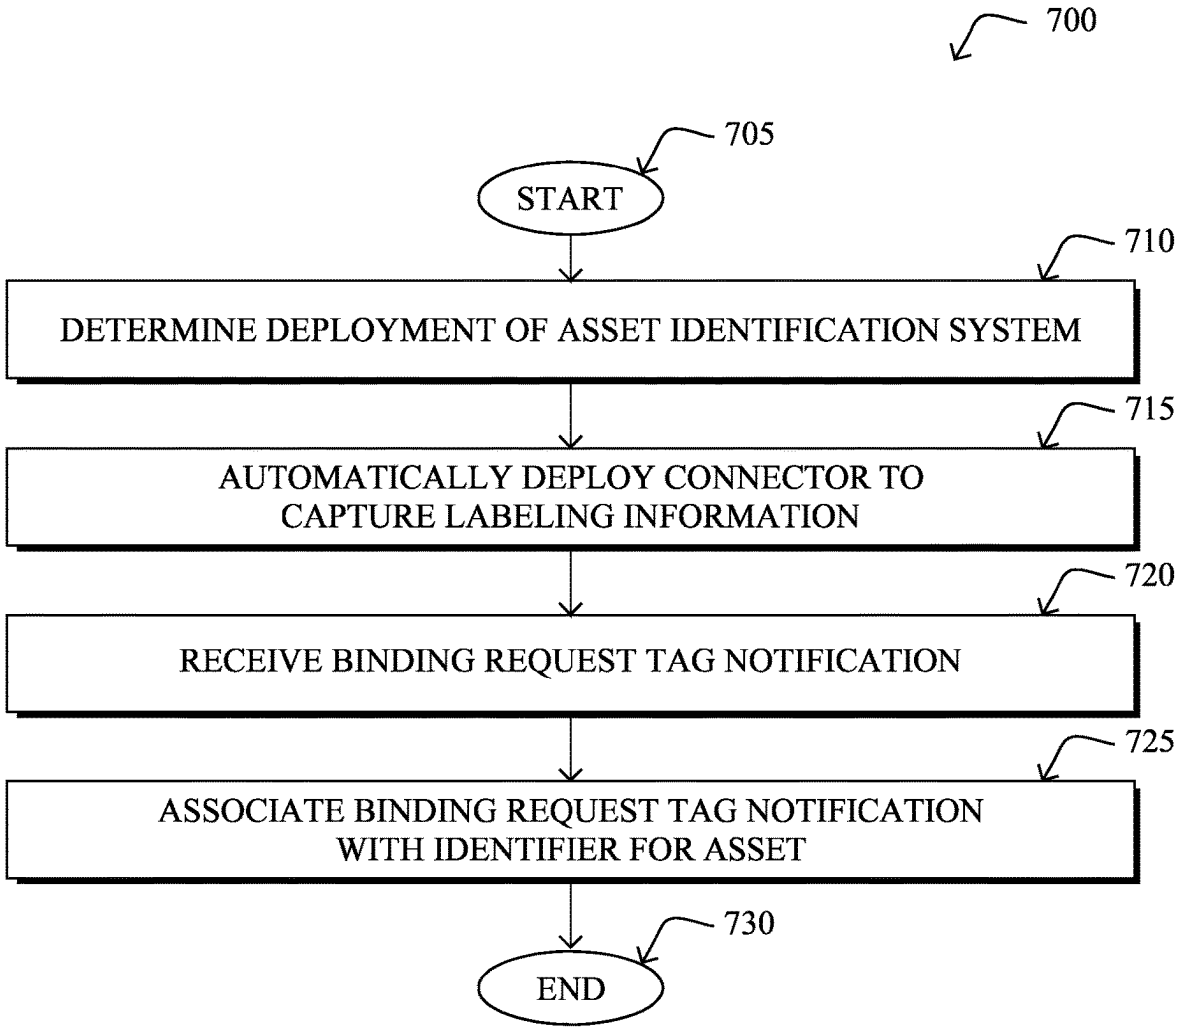
700
705
START
710
DETERMINE DEPLOYMENT OF ASSET IDENTIFICATION SYSTEM
715
AUTOMATICALLY DEPLOY CONNECTOR TO CAPTURE LABELING INFORMATION
720
RECEIVE BINDING REQUEST TAG NOTIFICATION
725
ASSOCIATE BINDING REQUEST TAG NOTIFICATION WITH IDENTIFIER FOR ASSET
730
END

FIG. 7

RELABELING-RESISTANT ITEM IDENTIFIERS USED ACROSS DOMAINS

TECHNICAL FIELD

[0001] The present disclosure relates generally to computer networks, and, more particularly, to relabeling-resistant item identifiers used across domains.

BACKGROUND

[0002] In a supply chain, a particular set of assets may change multiple hands as it moves from a source to a destination, for example, from a shipper via a carrier to a warehouse. Each of these stops may be understood as a domain of the supply chain. Each organization that controls the items, however brief, oftentimes uses different technologies and/or different identifier systems for tracking the particular set of assets. For example, a shipper may identify its shipment of items using barcodes, but the barcodes may be meaningless to other organizations that it received the goods from (e.g., a manufacturer) or organizations that it sends the goods to (e.g., e.g., a third-party warehouse). These other organizations are oftentimes unable to recognize the barcodes and, therefore, are unable to ascertain any contextual information that the shipping organization may have associated to individual barcodes.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The embodiments herein may be better understood by referring to the following description in conjunction with the accompanying drawings in which like reference numerals indicate identically or functionally similar elements, of which:

[0004] FIG. **1** illustrates an example network;

[0005] FIG. **2** illustrates an example network device/node;

[0006] FIG. **3** illustrates an example identity federation and visibility service;

[0007] FIG. **4** illustrates an example architecture for an identity federation and visibility service;

[0008] FIG. **5** illustrates an example architecture for an identification binding system for an identity federation and visibility service;

[0009] FIGS. **6**A-**6**B illustrate an example flow diagram for an identification binding system for an identity federation and visibility service; and

[0010] FIG. **7** illustrates an example simplified procedure for relabeling-resistant item identifiers used across domains.

DESCRIPTION OF EXAMPLE EMBODIMENTS

Overview

[0011] According to one or more embodiments of the disclosure, a device makes a determination that an asset identification system has been deployed at an organization. The device automatically deploys, based on the determination, a connector for the asset identification system to capture labeling information indicative of transport of one or more physical assets via the organization. The device receives, from the connector, a binding request tag notification indicative of transport of a particular physical asset of the one or more physical assets via the organization. The device associates the binding request tag notification with one or more identifiers for the particular physical asset used by one or more other organizations.

Description

[0012] A computer network is a geographically distributed collection of nodes interconnected by communication links and segments for transporting data between end nodes, such as personal computers and workstations, or other devices, such as sensors, etc. Many types of networks are available, ranging from local area networks (LANs) to wide area networks (WANs). LANs typically connect the nodes over dedicated private communications links located in the same general physical location, such as a building or campus. WANs, on the other hand, typically connect geographically dispersed nodes over long-distance communications links, such as common carrier telephone lines, optical lightpaths, synchronous optical networks (SONET), synchronous digital hierarchy (SDH) links, or Powerline Communications, and others. Other types of networks, such as field area networks (FANs), neighborhood area networks (NANs), personal area networks (PANs), etc. may also make up the components of any given computer network.

[0013] In various embodiments, computer networks may include an Internet of Things network. Loosely, the term "Internet of Things" or "IoT" (or "Internet of Everything" or "IoE") refers to uniquely identifiable objects (things) and their virtual representations in a network-based architecture. In particular, the IoT involves the ability to connect more than just computers and communications devices, but rather the ability to connect "objects" in general, such as lights, appliances, vehicles, heating, ventilating, and air-conditioning (HVAC), windows and window shades and blinds, doors, locks, etc. The "Internet of Things" thus generally refers to the interconnection of objects (e.g., smart objects), such as sensors and actuators, over a computer network (e.g., via IP), which may be the public Internet or a private network.

[0014] Often, IoT networks operate within a shared-media mesh networks, such as wireless or Powerline Communication networks, etc., and are often on what is referred to as Low-Power and Lossy Networks (LLNs), which are a class of network in which both the routers and their interconnect are constrained. That is, LLN devices/routers typically operate with constraints, e.g., processing power, memory, and/or energy (battery), and their interconnects are characterized by, illustratively, high loss rates, low data rates, and/or instability. IoT networks are comprised of anything from a few dozen to thousands or even millions of devices, and support point-to-point traffic (between devices inside the network), point-to-multipoint traffic (from a central control point such as a root node to a subset of devices inside the network), and multipoint-to-point traffic (from devices inside the network towards a central control point).

[0015] Fog computing (also known as edge computing, near edge computing, far edge computing, etc.) is a distributed approach of cloud implementation that acts as an intermediate layer from local networks (e.g., IoT networks) to the cloud (e.g., centralized and/or shared resources, as will be understood by those skilled in the art). That is, generally, fog computing entails using devices at the network edge to provide application services, including computation, networking, and storage, to the local nodes in the network, in contrast to cloud-based approaches that rely on remote data centers/cloud environments for the services. To this end, a fog node is a functional node that is deployed close to fog endpoints to provide computing, storage, and networking resources and services. Multiple fog nodes orga-

nized or configured together form a fog system, to implement a particular solution. Fog nodes and fog systems can have the same or complementary capabilities, in various implementations. That is, each individual fog node does not have to implement the entire spectrum of capabilities. Instead, the fog capabilities may be distributed across multiple fog nodes and systems, which may collaborate to help each other to provide the desired services. In other words, a fog system can include any number of virtualized services and/or data stores that are spread across the distributed fog nodes. This may include a master-slave configuration, publish-subscribe configuration, or peer-to-peer configuration.

[0016] Low power and Lossy Networks (LLNs), e.g., certain sensor networks, may be used in a myriad of applications such as for "Smart Grid" and "Smart Cities." A number of challenges in LLNs have been presented, such as:

[0017] 1) Links are generally lossy, such that a Packet Delivery Rate/Ratio (PDR) can dramatically vary due to various sources of interferences, e.g., considerably affecting the bit error rate (BER);

[0018] 2) Links are generally low bandwidth, such that control plane traffic must generally be bounded and negligible compared to the low rate data traffic;

[0019] 3) There are a number of use cases that require specifying a set of link and node metrics, some of them being dynamic, thus requiring specific smoothing functions to avoid routing instability, considerably draining bandwidth and energy;

[0020] 4) Constraint-routing may be required by some applications, e.g., to establish routing paths that will avoid non-encrypted links, nodes running low on energy, etc.;

[0021] 5) Scale of the networks may become very large, e.g., on the order of several thousands to millions of nodes; and

[0022] 6) Nodes may be constrained with a low memory, a reduced processing capability, a low power supply (e.g., battery).

[0023] In other words, LLNs are a class of network in which both the routers and their interconnect are constrained: LLN routers typically operate with constraints, e.g., processing power, memory, and/or energy (battery), and their interconnects are characterized by, illustratively, high loss rates, low data rates, and/or instability. LLNs are comprised of anything from a few dozen and up to thousands or even millions of LLN routers, and support point-to-point traffic (between devices inside the LLN), point-to-multipoint traffic (from a central control point to a subset of devices inside the LLN) and multipoint-to-point traffic (from devices inside the LLN towards a central control point).

[0024] An example implementation of LLNs is an "Internet of Things" network. Loosely, the term "Internet of Things" or "IoT" may be used by those in the art to refer to uniquely identifiable objects (things) and their virtual representations in a network-based architecture. In particular, the next frontier in the evolution of the Internet is the ability to connect more than just computers and communications devices, but rather the ability to connect "objects" in general, such as lights, appliances, vehicles, HVAC (heating, ventilating, and air-conditioning), windows and window shades and blinds, doors, locks, etc. The "Internet of Things" thus generally refers to the interconnection of objects (e.g., smart objects), such as sensors and actuators, over a computer network (e.g., IP), which may be the Public Internet or a private network. Such devices have been used in the industry for decades, usually in the form of non-IP or proprietary protocols that are connected to IP networks by way of protocol translation gateways. With the emergence of a myriad of applications, such as the smart grid advanced metering infrastructure (AMI), smart cities, and building and industrial automation, and cars (e.g., that can interconnect millions of objects for sensing things like power quality, tire pressure, and temperature and that can actuate engines and lights), it has been of the utmost importance to extend the IP protocol suite for these networks.

[0025] FIG. **1** is a schematic block diagram of an example simplified computer network **100** illustratively comprising nodes/devices at various levels of the network, interconnected by various methods of communication. For instance, the links may be wired links or shared media (e.g., wireless links, powerline communication links, etc.)

[0026] where certain nodes, such as, e.g., routers, sensors, computers, etc., may be in communication with other devices, e.g., based on connectivity, distance, signal strength, current operational status, location, etc.

[0027] Specifically, as shown in the example network **100**, three illustrative layers are shown, namely cloud layer **110**, fog layer **120**, and IoT device layer **130**. Illustratively, the cloud layer **110** may comprise general connectivity via the Internet **112**, and may contain one or more datacenters **114** with one or more centralized servers **116** or other devices, as will be appreciated by those skilled in the art. Within the fog layer **120**, various fog nodes/devices **122** (e.g., with fog modules, described below) may execute various fog computing resources on network edge devices, as opposed to datacenter/cloud-based servers or on the endpoint nodes **132** themselves of the IoT device layer **130**. For example, fog nodes/devices **122** may include edge routers and/or other networking devices that provide connectivity between cloud layer **110** and IoT device layer **130**. Data packets (e.g., traffic and/or messages sent between the devices/nodes) may be exchanged among the nodes/devices of the computer network **100** using predefined network communication protocols such as certain known wired protocols, wireless protocols, powerline communication protocols, or other shared-media protocols where appropriate. In this context, a protocol consists of a set of rules defining how the nodes interact with each other.

[0028] Those skilled in the art will understand that any number of nodes, devices, links, etc. may be used in the computer network, and that the view shown herein is for simplicity. Also, those skilled in the art will further understand that while the network is shown in a certain orientation, the network **100** is merely an example illustration that is not meant to limit the disclosure.

[0029] Data packets (e.g., traffic and/or messages) may be exchanged among the nodes/devices of the computer network **100** using predefined network communication protocols such as certain known wired protocols, wireless protocols (e.g., IEEE Std. 802.15.4, Wi-Fi, Bluetooth®, DECT-Ultra Low Energy, LoRa, etc.), powerline communication protocols, or other shared-media protocols where appropriate. In this context, a protocol consists of a set of rules defining how the nodes interact with each other.

[0030] FIG. **2** is a schematic block diagram of an example node/device **200** (e.g., an apparatus) that may be used with one or more embodiments described herein. As shown, device **200** may comprise one or more communication interfaces **210** (e.g., wired, wireless, etc.), at least one

processor **220**, and a memory **240** interconnected by a system bus **250**, as well as a power supply **260** (e.g., battery, plug-in, etc.).

[0031] Communication interface(s) **210** may be coupled to processor **220** and include the mechanical, electrical, and signaling circuitry for communicating data over a communication link. To this end, communication interface(s) **210** may be configured to transmit and/or receive data using a variety of different communication protocols, such as TCP/IP, UDP, etc. Note that the device **200** may have multiple different types of communication interface(s) **210**, e.g., wireless and wired/physical connections, and that the view herein is merely for illustration.

[0032] The memory **240** comprises a plurality of storage locations that are addressable by the processor(s) **220** and the communication interface(s) **210** for storing software programs and data structures associated with the embodiments described herein. The processor **220** may comprise necessary elements or logic adapted to execute the software programs and manipulate the data structures **245**. An operating system **242**, portions of which are typically resident in memory **240** and executed by the processor(s), functionally organizes the node by, inter alia, invoking network operations in support of software processors and/or services executing on the device. These software processors and/or services may comprise an identity federation and visibility process **248** and an identification binding process **249**.

[0033] It will be apparent to those skilled in the art that other processor and memory types, including various computer-readable media, may be used to store and execute program instructions pertaining to the techniques described herein. Also, while the description illustrates various processes, it is expressly contemplated that various processes may be embodied as modules configured to operate in accordance with the techniques herein (e.g., according to the functionality of a similar process). Further, while processes may be shown and/or described separately, those skilled in the art will appreciate that processes may be routines or modules within other processes.

[0034] Modern supply chains typically span multiple organizations, such as the shipper of an item, any number of carriers, and the target destination of the item. As the item travels towards its destination, its digital representation may undergo a number of transformations. For instance, the identity of the item under the responsibility of a first carrier is likely to be different than the identity of the same item under the responsibility of a second carrier.

[0035] Even within a single organization, the identification of a particular item may vary due to the different technologies that shippers may employ. For instance, an item may be tagged using a barcode, a Quick Response (QR) code, radio frequency identification (RFID) tag, Bluetooth Low Energy (BLE) tag, cellular tag, or the like. It is also quite common for an item to be re-tagged when passing from one carrier to another (e.g., the new carrier puts a new barcode on the item being shipped). In doing so, this effectively changes the digital identity of the item. As a result, supply chains that span multiple organizations often lack end-to-end transparency.

[0036] As noted above, visibility represents one of the main areas of focus in supply chains, as they undergo a digital transformation. In general, visibility refers to the ability to answer questions such as the following:

[0037] What is this item?

[0038] Where is a particular item?

[0039] What is the state of a particular item?

[0040] Etc.

[0041] Without visibility across the supply chain, it is quite difficult to make informed decisions. Despite this, modern supply chains tend to be fractured in terms of visibility, due to the myriad of organizations that may be involved.

[0042] A major challenge in automating end-to-end supply chains is that there is no one-size-fits-all identification technology that has received universal support by the industry. Today, for instance, some products may be tagged with radio frequency identification (RFID) transponders, others may use barcodes, Bluetooth Low Energy (BLE) or alternative technologies in the future, such as ultra-wideband (MB) or cellular (e.g., 5G) active tags. These technologies have different identity spaces and formats, and many tagged products have different identity providers. For example, the identity of a product while under the responsibility of a first carrier is very likely to be different than the identity of the same product while under the responsibility of a second carrier. In fact, it is quite common for a product to be "re-tagged" when passing from one carrier to another, effectively changing its digital identity.

[0043] Even if all the supply chains worldwide would agree upon a single Universally Unique Identifier (UUID) and a common traceability technology, still dealing with different identity providers is unavoidable by the mere construction of the supply chain. For instance, many warehouses receive, store, and ship products manufactured by different entities, and therefore, those products will have different identity providers. In fact, supply chains are inherently federated ecosystems, but their members lack the ability to amalgamate the different identities and identity providers involved in this type of federation.

[0044] Accordingly, there is a real-world demand to verify, and attest to, the identity of items that are transported via a (contactless) supply chain across different organizations. One potential approach to achieving this would be to build complex integrations between the systems of two organizations (e.g., a shipper and a carrier). However, this typically requires both organizations to use the same software or for one organization to 'adapt' to the system of another organization. Adaptation of a common system is typically not practical, when across an entire supply chain, as there may be many different organizations that have responsibility for an item.

Item Identity Federation and Visibility as a Service

[0045] The techniques herein introduce an item identity federation and visibility service that allows common policy lists to be rendered, automatically, for exchanging data between organizations based on their profiles (e.g., a manufacturer and a warehouse, a transporter and a warehouse, etc.). In some aspects, the techniques herein also allow users to select which information they wish to request from other organizations (e.g., a. 'visibility intent'), as well as the information they would like to publish to other organizations (e.g., a 'visibility offering'), In further aspects, the techniques herein allow for the automatic matching of visibility policies among organizations: 1.) based on the identification

technologies used (e.g., RFID, barcodes, BLE, etc.) and 2.) based on the specified visibility intents and the visibility offerings.

[0046] Operationally, the federated identity method introduced herein allow for information about an item to be shared across different organizations, systems, and/or technologies. In such a model, instead of a single identity system being used, requests for verification and/or validation of an identity may be passed to a specified identity provider. Note that prior approaches to this have mainly focused on the authentication and authorization of users or other data consumers. However, the requirements in terms of item identification in a contactless supply chain differ significantly from what these types of approaches offer. For instance, inventory visibility typically entails challenges identification and notification across domains, rather than those relating to authentication or guest network access. Indeed, an RFID interrogator or barcode scanner will not need remote authentication of a passive RFID transponder or barcode, since such a functional requirement does not (currently) exist. Moreover, many of the devices used to tag (and identify) inventory do not need guest/Internet access (e.g., RFID BLE, barcodes, UWB, etc.).

[0047] In addition, simplicity and ease of use is paramount in supply chains. Indeed, the trend in this sector is to increase the level of automation of the large majority of the processes. In so doing, organizations are choosing to lease a significant part of the equipment and infrastructure needed to run the business, rather than purchase such systems themselves. This applies to a wide variety of equipment, ranging from advanced systems such as Automated Mobile Robots (AMRs) and automated forklifts to industrial RFID interrogators, meaning that, there is a rapidly increasing number of third-party identities and devices that need to interact with the local communications infrastructures and systems of an organization.

[0048] FIG. **3** illustrates an example identity federation and visibility service, according to various embodiments. As shown, system **300** may include an identity federation and visibility service **302** that allows any organization (e.g., a member of a supply chain) to easily initiate an identity federation. For instance, as shown, consider a typical supply chain that involves the following organizations: a manufacturer **304** that ships an item to a destination. During transit, any number of transporters, such as transporters **306-308** (e.g., a first transporter, a second transporter, etc.), may receive, transport, and hand off the item to another organization. In addition, there may be any number of other organizations involved in the supply chain, such as warehouse operators **310-312** (e.g., a first warehouse operator, a second warehouse operator, etc.), that store the shipped item while in transit and/or on delivery.

[0049] By way of example, consider the case in which manufacturer **304** is to send an item for delivery to a third-party warehouse operator **312** (e.g., a retailer). To do so, manufacturer **304** may pass the item to transporter **306**. In turn, transporter **306** may ship the item and, at some point, deposit the item at a warehouse operated by warehouse operator **310**. From there, transporter **308** may pick up the item from the first warehouse and convey it through its own channels until finally delivering the item to its final destination, a warehouse operated by warehouse operator **312**.

[0050] At the core of system **300** is identity federation and visibility service **302**, which may be provided by one or more devices, such as device **200**, through the execution of identity federation and visibility process **248**. In some instances, identity federation and visibility service **302** may be provided by an organization that differs from those of manufacturer **304**, transporters **306-308**, warehouse operators **310-312**. In other instances, identity federation and visibility service **302** may be provided by any of these organizations.

[0051] After creating an identity federation via identity federation and visibility service **302**, the creator can invite other organizations to join the federation and start using it. More specifically, the techniques herein allow even non-technically savvy users to create and/or join an identity federation via identity federation and visibility service **302** in a rapid manner. Once established, the identity federation provided by identity federation and visibility service **302** allows the authorized participant organizations to share and view information about the item throughout its traversal of the supply chain.

[0052] For instance, assume that each of the organizations **304-312** have registered with identity federation and visibility service **302** and participate in the identity federation of the item shipped by manufacturer **304**. Each of these organizations may independently specify their intents in terms of what data they wish to receive regarding the item, as well as in terms of what data they are willing to share. Assume, for instance, that manufacturer **304** wishes to know the state of its item in terms of where it is, when it arrived at its current location, when it leaves, the level of inventory in its final destination, etc. Through the use of identity federation and visibility service **302**, the other organizations **306-312** may provide updated information to identity federation and visibility service **302** regarding the item, automatically, allowing manufacturer **304** to access this information via identity federation and visibility service **302**.

[0053] FIG. **4** illustrates an example architecture **400** for an identity federation and visibility service, such as identity federation and visibility service **302** in FIG. **3**. At the core of architecture **400** is identity federation and visibility process **248** which may comprise any or all of the following components: a profiling module **402**, a policy editor & visibility selection generator **404**, a policy engine **406**, a connector generator **408**, and/or a reporting module **410**. As would be appreciated, the functionalities of these components may be combined or omitted, as desired. In addition, these components may be implemented on a singular device or in a distributed manner, in which case the combination of executing devices can be viewed as their own singular device for purposes of executing identity federation and visibility process **248**.

[0054] During execution, identity federation and visibility process **248** may communicate with any number of user interface(s) **412** operated by any number of people across any number of organizations. For instance, user interface(s) **412** may comprise desktop computers, laptop computers, smart phones, tablet devices, wearable electronic devices, or the like.

[0055] According to various embodiments, a domain expert for an organization may use user interface(s) **412** to specify profiling data **422** to identity federation and visibility process **248**. In turn, profiling module **402** may use profiling data **422** to define a set of selectable visibility policies for the organization or a domain of the organization. In other words, profiling module **402** may generate and expose a set of

pre-defined visibility offerings and intents that could be applied and/or modified later on by a particular organization. For instance, profiling data **422** may define different organization types, such as, but not limited to, manufacturers, warehouse operators, transporters, etc., in the case of a supply chain. In other embodiments, identity federation and visibility process **248** may be used to set up a federation in other use cases, as well, such as healthcare, insurance, or the like. As would be appreciated, the domain expert(s) may be unaffiliated with the specific organization participating in an identity federation and may simply provide a base set of defaults for process **248** for the usage domain (e.g., supply chain, healthcare, etc.).

[0056] The domain expert(s) may also provide visibility policy generation data **424** that allows policy editor & visibility selection generator **404** to automatically generate predefined policy templates. In general, visibility policy generation data **424** indicates the types of data that an organization with a particular profile may be able to share and/or the types of data that it may wish to be able to access. For instance, a policy template for a particular type of organization, as indicated by its profile, may specify a predefined set of preferences for its visibility intent and visibility offerings. In addition, the choices available for selection in the template may also depend on the organization type of the participant and/or those with whom the organization is to interact as part of the supply chain. For instance, a Manufacturer profile may be associated to a Manufacturer policy template, while a Warehouse might be associated to a Warehouse policy template. Based on these two policy templates and considering the (Manufacturer, Warehouse) pair, policy editor & visibility selection generator **404** may auto-render a predefined set of choices that each member of the pair can select to express its visibility intent and offerings. Likewise, policy editor & visibility selection generator **404** may auto-render a predefined set of choices for a (Transporter, Warehouse) pair or a (Manufacturer, Transporter) pair. The list of auto-rendered policies may vary depending on the members profiles. For example, the visibility choices auto-rendered by policy editor & visibility selection generator **404** for (Manufacturer_1, Warehouse_1) and (Manufacturer_2, Warehouse_2) would be the same, while the one for (Transporter_1, Warehouse_1) might differ.

[0057] Thus, as a result of the processing by profiling module **402** and policy editor & selection generator **404**, identity federation and visibility process **248** now has a set of default visibility policies that indicate the set of 'intents' of the organization or domain in terms of which data that it is able to share (e.g., its 'visibility offerings'), as well as any information that it can access from another organization (e.g., its 'visibility intent').

[0058] Once the above initializations have occurred, specific organizations and their users may be onboarded by process **248**. Such a registration may, for instance, associate a particular organization with a profile type and, accordingly, a set of default policy templates for them. For instance, a particular user at Warehouse A may provide details to process **248** regarding their organization (e.g., location information, contact information, etc.). In addition, this information may also indicate the specifics of the capabilities of the systems of the organization, such as its software systems, identification technologies in use, and/or any other information that may be captured regarding the participating organization. For instance, this information may indicate, for a particular organization or domain, which identification technologies are in use by that organization, such as RFID tags, barcodes, Wi-Fi, cellular, BLE tags, UWB, and the like.

[0059] In various embodiments, a user associated with a particular organization registered with process **248** may specify policy selection data **426**, which is used by connector generator **408** to construct and deploy the appropriate connector(s) **414**, to facilitate the corresponding sharing of data across organizations. Typically, policy selection data **426** comprises a selection of the visibility intent and/or visibility offering of that user. For instance, policy selection data **426** may take the form of checkbox input that allows a user to select from among the template policy options generated previously by policy editor & visibility selection generator **404** and based on the specific profiles of the organizations involved.

[0060] Regarding visibility selection, the identity federation and visibility service may provide display data to a user interface associated with a 3rd party warehouse (e.g., Warehouse W), thereby allowing a user of that organization to specify policy selection data **426** for a Manufacturer M (e.g., manufacturer **304**). For instance, display data may take the form of a user interface (e.g., graphical user interface-based) checklist via which the user is able to specify the policy target(s), send policy data regarding what types of information should be sent, as well as the corresponding receive policy data. Similarly, other display data may be sent to a user interface associated with Manufacturer M, also a participant in the identity federation, that allows that organization to specify its own policy target(s), send policy data, and receive policy data, with respect to Warehouse W.

[0061] As would be appreciated, the options available to a user may vary depending on the organization type, the peer's organization type, identification technology, etc., and selectable pairs may be auto-rendered based on profile pairs (e.g., a manufacturer-warehouse pair, a warehouse-carrier pair, etc.). For instance, a warehouse operator may specify any or all of the following as part of a send policy: notify arrival, notify departure, in stock query, quantity in stock query and specify any or all of the following as part of a receive policy: estimated time of departure (ETD) updates, asset description, or asset state. In contrast, a manufacturer may be able to specify any or all of the following as part of a send policy: ETD updates, asset description, or asset state. In addition, the manufacturer may specify any or all of the following as part of a receive policy: notify arrival, notify departure, in stock query, quantity in stock query. Another example selection may indicate the location of a particular item (e.g., in terms of coordinates), which may be of interest with respect to a transporter.

[0062] According to various embodiments, policy engine **406** may match visibility intents and visibility offerings specified via policy selection data **426** across any number of organizations or domains. For instance, one organization may wish to receive a notification when another organization receives any of its shipped goods. If the second organization specifies a visibility offering that matches the visibility intent of the first organization, policy engine **406** may implement the corresponding data sharing policy between the two organizations. In other words, the role of policy engine **406** may be to: 1.) hook, map, and manage the send and receive policies produced by policy editor &

visibility selection generator **404**, 2.) perform matching among those send and receive policies selected via policy selection data **426**, and/or 3.) distribute the policies to the identity federation core and to the applicable connectors, such as connector **414**, as detailed below.

[0063] As would be appreciated, the set of available options for an organization, as well as its selected visibility offerings and intents, can change over time. Indeed, even though a default set of selectable options may be configured for an organization, these options can be modified over time, as needed, in some embodiments. For instance, say an organization outfits its warehouse with BLE scanners. In such a case, a person associated with that organization may specify this new offering as an option to identity federation and visibility process **248**. In turn, the new options will be auto-rendered and made available to the users of the federation. Similarly, a user may adjust their policy selection data **426** over time, such as when additional information is desired, certain information is no longer of interest, etc. Indeed, even though process **248** may present users with default options to set visibility policies, these options could be modified over time. For instance, in some embodiments, this could be instrumented either via a new configuration by an authorized user or programmatically (e.g., through **424**), including the introduction of new visibility policies and data models.

[0064] Identity federation and visibility process **248** may also maintain any number of identity federations across open, semi-private, or private consortia of the various organizations. In addition, members may be part of multiple identity federations maintained by identity federation and visibility process **248**, simultaneously.

[0065] In some embodiments, an identity federation may be implemented as an unmanaged service, through the execution of identity federation and visibility process **248**, with no requirement for an identity federation provider to operate the federation. This allows the federation to begin functioning as soon as a first organization establishes it and invites a second organization to participate in the federation and start exchanging data.

[0066] In the case in which a user opts to start a new identity federation, the user may specify this to process **248**, such as the name of the new identity federation, invitees to the federation, and the like. In one embodiment, identity federation and visibility process **248** may also suggest invitees to the user via user interface **412** and/or allow the user to pick invitees from a predefined list. This may be based, for instance, on the prior selections of the user and/or organization for other identity federations, the most common invitees (e.g., particular transporters, etc.), a template defined by the user, or the like.

[0067] As would be appreciated, not all of the invitees specified by a user may be registered with identity federation and visibility process **248**. In such cases, the user setting up a new federation may also specify information such as any or all of the following:

[0068] 1. The name of the invitee (e.g., 'Warehouse W1').

[0069] 2. An email address or other contact information to which identity federation and visibility process **248** may send an invitation for the identity federation.

[0070] 3. A member role that indicates the allowed activities for the invitee within the federation (e.g., "member," "co-owner," "member with the right to invite others," etc.).

[0071] 4. Membership duration information (e.g., one day, one week, permanent, etc.).

[0072] 5. Other information regarding the invitee (e.g., the location of the invitee, notes, etc.).

[0073] Note that, in some instances, the creator of an identity federation via identity federation and visibility process **248** may also delegate the ability to invite others to one or more invitees. This is an important capability, as many logistics and transportation companies rely on several layers of subcontracting in order to deliver an outcome. By extending invitations to subcontractors so that they can participate in the identity federation and exchange information about an item, visibility of the item can be greatly improved. Once the process is complete, identity federation and visibility process **248** may send invitations to the selected invitees (e.g., by sending links to identity federation and visibility process **248** by email, text message, etc.).

[0074] In general, an invitation to join an identity federation may identify the federation to the invitee and may include security token information that identifies the invitee to identity federation and visibility process **248**. Once registered with identity federation and visibility process **248**, or logged into an existing account, the invitee may enroll with the created identity federation. During this step, each party may select whether its profile should be public or not (e.g., whether the company/entity name should be publicly available and listed). An organization may also maintain different accounts and/or user roles, such as when the organization participates in different identity federations. For instance, a member may create internal tenant accounts, such as a 'viewer' account, an 'admin' account, etc.

[0075] According to various embodiments, connector generator **408** of identity federation and visibility process **248** may be configured to generate any number of 'connectors' for the participants in a federation, based on matches between their visibility offerings and visibility intents. For instance, connector generator **408** may generate connector **414** that encompasses the software necessary to interface with the item identification technologies used by a particular organization/participant in an identity federation.

[0076] As would be appreciated, item information may depend heavily on the internal systems of the source organization. For instance, different organizations may maintain item information in various enterprise-level applications or systems, such as an Enterprise Resource Planning (ERP) system, a Warehouse Management System (WMS), a Warehouse Execution System (WES), a Transport Management System (TMS), or the like. To this end, connector generator **408** may be configured to select plugin data **416** for inclusion in a connector (e.g., connector **414**) that facilitates the sharing of information from that resource. For instance, assume that an organization uses one or more application(s) **418** to maintain item information, such as a TMS. In such a case, plugin data **416** may include the information needed to interface with the TMS of the organization (e.g., credential information, etc.), to obtain or update information about an item in transport by the organization. Similarly, a connector **414** may be configured to share data from particular device (s) **420**, such as an RFID reader, etc., via the federation.

[0077] Typically, the communications enabled by plugin data **416** may be limited to accessing only the information needed to support the visibility intents and visibility offerings associated with the various participants, as determined and enforced by policy engine **406**. For example, assume that Warehouse W1 receives a pallet of RFID-tagged items from Manufacturer M. As soon as an RFID reader of Warehouse W1 scan the items, the corresponding information may be captured into the application(s) **418** of Warehouse W1 and captured by way of plugin data **416**, which may comprise one or more plugins for those application(s) **418**. In reading the information encoded on the RFID tag, the system is able to determine the identity provider as being Manufacturer M. This then determines that the information should be 'routed' according to the policy associated with Manufacturer M. Note that this data flow may occur because Manufacturer M specified a visibility intent of "notify arrival" that matched a visibility offering selected by Warehouse W1.

[0078] In further embodiments, connector **414** may include plugin data **416** configured to interface with the device(s) **420** of the organization, directly. For example, connector **414** may receive data in parallel from any number of deployed RFID scanners in addition to, or in lieu of, interfacing with a WMS or other application **418** of that organization. For instance, connector **414** may obtain scan information directly from a scanner, but may obtain inventory count information from a WMS or other application **418**.

[0079] According to various embodiments, connector generator **408** may construct a connector **414** based on the profile of the organization, its identification technologies, and/or its corresponding plugins. In general, connector **414** functions as an entity that allows an organization to connect to identity federation and visibility process **248** and exchange data flows. For security reasons, connector **414** may also include a certificate signed by connector generator **408**, to certify its identity to the identity federation and visibility process **248**. Likewise, the identity federation and visibility process **248** may also handle a certificate to certify its identity to the organization's connector **414**.

[0080] Assume, for instance, that a participant in an identity federation has specified the following: Profile="Warehouse"; Technologies="(RFID, Barcodes)"; Plugin="Oracle's WMS version X"; CERT="W_auto-signed". In this example, connector **414** may take the form of a software element (e.g., an image), which, once instantiated, embeds the processes providing the following functionality:

[0081] A list of policies that the participant can select to define its visibility intent, based on its profile, which may be modifiable on a per-participant basis.

[0082] A list of policies that the participant can select to define its visibility offering, based on its profile, which may also be modifiable on a per-participant basis.

[0083] The ability to parse and forward Electronic Product Codes (EPCs) from application(s) **418** and/or device **420**, which is the standardized identity format used both by RFID tags and Barcodes.

[0084] Plugin data **416** to parse and forward a set of messages to/from Oracle's WMS (e.g., an application **418**).

[0085] Code to create secure communications with identity federation and visibility process **248** that are authenticated using the certificate information provided by connector generator **408**.

[0086] In some embodiments, a new connector **414** may be instantiated for each identity federation in which an organization participates. In another embodiment, a single connector **414** may include the ability to slice and isolate item data from different identity federations, as in the case in which the organization is participating in multiple identity federations. In further embodiments, connector **414** may also be configured to store and cache policies (e.g., by implementing a localized form of policy engine **406**).

[0087] Once connector generator **408** has generated a connector **414**, identity federation and visibility process **248** may deploy it to its target organization either on a push basis or on a pull basis (e.g., in response to the organization requesting a download of connector **414**). As noted, the specific configuration of connector **414** may depend on the identification technologies and/or plugins specified by the target organization. For instance, if RFID and barcodes were chosen, connector **414** may be configured as yet another data consumer in the existing data capture workflows of the organization, as well as interface with internal application(s) **418**. In a further enhancement, connector **414** may also interface directly with the endpoint devices and/or networking equipment responsible for capturing and reporting the item data. For instance, connector **414** may be configured to interface with BLE beacons, UWB anchors, Layer **2** switches, wireless access points, wireless access point controllers, or the like. To support this, connector generator **408** may include the corresponding plugin data **416** into connector **414**, such as the plugin data needed to interface with application programming interfaces (APIs) of a network switch, barcode reader, etc.

[0088] Once connector **414** has been deployed, it may automatically and securely connect to identity federation and visibility process **248** (e.g., the service) and/or directly with other connectors of other participating organizations in the federation. This may be achieved by establishing a mutual transport layer security (mTLS) tunnel, for instance. Thus, as shown previously in FIG. **3**, identity federation and visibility service **302** may communicate with organizations **304-312** via secure tunnels that are established with corresponding connectors **414** deployed to those organizations. In turn, the connectors may report the required item data to service **302**, as needed.

[0089] Policy engine **406** is also responsible for enforcing the visibility policies generated by policy editor & visibility selection generator **404**. As would be appreciated, the visibility policies between organizations may differ. Indeed, the policies for exchanging data between a transporter and a warehouse operator may differ from the policies for exchanging data between a manufacturer and a warehouse operator. Thus, a warehouse operator may share certain information with a carrier that the shipper (e.g., the manufacturer) may not be allowed to view.

[0090] Policy enforcement by policy engine **406** can be achieved by placing restrictions on reporting module **410**, which is responsible for reporting any item data collected by a connector **414** to any of the other participants of the federation that match that data. In turn, reporting module **410** may provide that collected item data as item data **428** to the authorized organizations via user interface(s) **412**. Note that the functionalities of reporting module **410** may also be implemented as part of connector **414**, thereby allowing

connector **414** to report item data **428** to one or more other connectors, directly, for presentation to a user interface **412**. In various embodiments, item data **428** may be sent via the corresponding connector(s) **414** deployed to the destination organization(s) authorized to receive the item data. This allows, in some instances, item data **428** to be fed into the system(s) of that organization via the plugin data of the connector. For instance, in some cases, item data **428** may automatically populate the ERP system of the organization receiving item data **428**.

[0091] By way of illustration of the operation of the full system, consider again the case shown in FIG. **3**. Assume that manufacturer **304** ships an item with a passive RFID tag to warehouse operator **310**. An RFID reader in the warehouse may read the data in the RFID tag, providing local visibility to the systems of warehouse operator **310** (e.g., a WMS application, etc.). The item data read by the RFID interrogator thus reaches the connector deployed to warehouse operator **310**, allowing the connector to either find a matching entry or filter the data. Various embodiments could be used to enable such functionality at the connector level. For example, connectors might cache the ID of the owners (or identity providers) along with the corresponding visibility policy. Another option could be to push policy and updates to the connectors. The connectors might also be stateless and might not be endowed with any caching mechanism, in which case, the messages will always hit the federation.

[0092] In more complex examples, visibility policies may allow for different levels of aggregation, such as by aggregating groups of item identifiers. For instance, a 'notify arrival' policy may be applied to an entire pallet of items, rather than for each of the individual items on the pallet.

[0093] In some embodiments, the identity federation may be used for exchanging both identity-related data, as well as context and state-related data, subject to the visibility policies described above. In other words, identity federation and visibility service **302** may serve as both the control plane and data plane, to enable the desired levels of visibility among parties in the supply chain. In other embodiments, the identity federation may only carry control plane data and may provide trusted pointers (e.g., endpoint addresses), so that the parties can exchange context and state-related data directly among themselves. In a further enhancement, the identity federation may support a hybrid model, providing both control and data plane functions for certain types of identities, profile members or visibility functions, while providing control and data plane separation for others.

[0094] In further enhancements, e identity federation may support any or all of the following:

- **[0095]** more flexible profiles and templates, such as customizable templates;
- **[0096]** programmable/extensible policies;
- **[0097]** programmable/extensible plugins;
- **[0098]** more flexible certificate and handling or even a full-fledged public key infrastructure (PKI);
- **[0099]** additional decoupling between the identifying/routing functions and the exchange of context and state-related data.

[0100] According to various embodiments, the federation techniques herein can also be extended to exchange information regarding connected assets and/or the workforce of the organizations. For instance, assume that warehouse operator **310** has deployed a number of AMRs or automated forklifts in its warehouse. In such cases, the connector deployed to warehouse operator **310** may also interface with their corresponding systems, to provide the manufacturer of those devices (and any other authorized participant of the federation) information regarding these devices. Similarly, personnel information could also be captured and sent, through the use of the appropriate connector and plugins. For instance, warehouse operator **310** may send to transporter **308** information regarding the precise location of the item, identification of the robot or person transporting the item within the warehouse, and an expected time that the item will be ready for pickup by transporter **308**.

[0101] Warehouse operators are increasingly leasing high-value equipment, such as AMRs and the like. Since such systems require data network communication in order to operate, AMRs or unmanned forklifts represent third-party devices that require trusted access to the warehouse network. Thus, identity federation and visibility service **302** may be used to dynamically identify these assets and enable trustworthy communications between these assets and the management systems of their corresponding service providers, which may be integrated into the federation as additional participants.

[0102] For instance, warehouse operator **310** may exchange data gathered through a Wi-Fi technology connector with remote Forklift or Robotics Management systems running in a remote location, such as a cloud-hosted compute facility. However, it may exchange data gathered only through the RFID connector with other parties upstream in the supply chain, including the manufacturer/seller of the items tagged with RFID.

[0103] As noted above, items of different nature (e.g., parcels, boxes, pallets, entire containers, etc.) will move across domains (i.e., change hands) several times as they are transited a supply chain, for example, from a shipper via a carrier to a destination. In order to process and handle these items, individual organizations along the one or more supply chains that (temporarily) retain custody of the items need to be able to identify the items. For the purposes of this identification, organizations may, as described herein above, implement barcodes, RFID transponders, QR codes, BLE, active 5G (or other communication protocol) tags along with various tracking methods. Each organization that controls the items, however brief, oftentimes uses different technologies and/or different identifier systems. For example, a shipper may identify its shipment of items using barcodes, but the barcodes may be meaningless to other organizations that it received the goods from (e.g., a manufacturer) or organizations that it sends the goods to (e.g., a third-party warehouse). These other organizations are oftentimes unable to recognize the barcodes and, therefore, are unable to ascertain any contextual information that the shipping organization may have associated to individual barcodes. Even more, the other organizations may "relabel" the goods using its own system of identifiers, which are very different than that of the shipping organization, for the items once they items are in their custody. This "relabeling" of identifiers may occur several times for the shipment of goods, as it moves across domains of a supply chain.

[0104] In an example from the perspective of a domain (e.g., organization) that is taking custody of items from another domain, where the goods may have been labeled with an identifier (e.g., ID1) that is used by an enterprise system. The domain taking custody of the goods may:

[0105] scan and reuse the identifier (e.g., ID1) used by the enterprise system (from the domain the goods come from), which is challenging due to the lack of interoperability of identifiers across domains/enterprise systems;

[0106] scan and log the identifier (e.g., ID1) as well as generate a new identifier (e.g., ID2), which is used by another enterprise system of the domain (e.g., organization) that is taking custody of the goods—effectively ignoring the identifier used by the enterprise system (from the domain the goods come from);

[0107] ignore the identifier (e.g., ID1) and relabel the goods by generating a new identifier (e.g., ID2), which is used by another enterprise system of the domain (e.g., organization) that is taking custody of the goods; or

[0108] ignore the identifier, aggregate/separate the goods into different quantifies, then generate a new identifier (e.g., ID2) that is associated with the aggregated/separated goods, where the other enterprise system of the domain (e.g., organization) that is taking custody of the goods uses the new identifier (ID2).

Additionally, particular identification technologies used by the two organizations may be different. For example, a sending organization may use passive RFID transponders while a receiving organization uses barcodes. These technologies may also change over time (e.g., an organization may adopt BLE for a subset of items after a period of time).

[0109] In this example, then, multiple challenges are faced when goods, products, etc. move across domains from one organization to another. That is, in part due to the previously mentioned challenges of "relabeling" identifiers to other identifiers (e.g., cross-technology issues, changes in technologies used, etc.), identification and traceability of goods remains fragmented across organizations. In other words, organizations may lose "visibility" of particular goods, items, etc. Without end-to-end visibility across domains, automation of operations, informed decision-making for operators of supply chain, etc. are strictly bound to individual, per-domain bases.

Relabeling-Resistant Item Identifiers Used Across Domains

[0110] The features provided by the example identity federation and visibility service (for supply chain partners), according to various embodiments described above, enables previously siloed systems of a supply chain to share information with one another. The techniques herein further introduce an identification binding process **249** to provide for to relabeling-resistant item identifier used across domains.

[0111] By tracking and identifying items across different domains (e.g., manufacturer, carrier, shipper, etc.) of a supply chain, in combination with the above-described example identity federation and visibility service, the techniques herein bind different identifiers (from different domains) together for individual (or sets of) items or goods. Doing so avoids problems or issues that arise due to individual organizations/domains "relabeling" items or goods (i.e., in a "relabeling-resistant" manner). In addition, automatic detection of identification system(s) at a particular domain may be implemented, so that the example identity federation and visibility service can automatically configure and obtain access to item, goods, etc. information of the particular domain. Of note, the techniques described herein do not require an overall convention or a priori agreement on technologies, standards, etc. for identifiers across domains. Similarly, no peer-to-peer integration among identification system(s) and/or backend system(s) of particular domains is required. Accordingly, having visibility across an entire supply chain may be achieved; specifically, physical assets may be identified and located across the supply chain, even when particular domains of the supply chain apply different tracking technologies and/or identifiers.

[0112] In particular, one or more connectors (e.g., described above with respect to connector generator **408**), when installed at a corresponding domain or organization of a supply chain, may be configured to automatically access and identifying identifiers used by a particular domain. For instance, the one or more connectors may be implemented at tracking identifier systems of a number of different domains/organizations, including warehouses, manufacturers, carriers and other members of a supply chain. These one or more connectors may further be configured to automate distribution of this identifier information to other organizations in the supply chain, where this identifier information may be used to bind different identifiers (from different domains) together for individual (or sets of) items or goods.

[0113] Illustratively, the techniques described herein may be performed by hardware, software, and/or firmware, such as in accordance with the identification binding process **249** (in combination with identity federation and visibility process **248**), which may include computer executable instructions executed by the processor **220** (or independent processor of interfaces **210**) to perform functions relating to the techniques described herein.

[0114] Specifically, according to various embodiments, a device makes a determination that an asset identification system has been deployed at an organization. The device automatically deploys, based on the determination, a connector for the asset identification system to capture labeling information indicative of transport of one or more physical assets via the organization. The device receives, from the connector, a binding request tag notification indicative of transport of a particular physical asset of the one or more physical assets via the organization. The device associates the binding request tag notification with one or more identifiers for the particular physical asset used by one or more other organizations.

[0115] Operationally, FIG. **5** illustrates an example architecture **500** for an identification binding system **502** (that implements the identification binding process **249**) for an identity federation and visibility system/service **504** is shown. Identity federation and visibility system/service **504** may, as shown, be in secure communication with a first domain **506** (e.g., a shipper) and a second domain **508** (e.g., a third-party warehouse) that may each be understood as organizations that are each separately part of an overall supply chain, where each domain may use different operational technologies, also known as, asset identification systems. It is to be understood that each of the domains may comprise, for example, a seaport terminal, an airport, a transport vehicle, a warehouse, or a distribution center.

[0116] That is, each of these domains may correspondingly implement a first enterprise system **510** and a second enterprise system **512** that are used for tracking assets, goods, items, etc. within, respectively, first domain **506** and second domain **508**. These enterprise systems may, as

described above, comprise one or more of an ERP, WMS, WES, TMS, etc. Further, it is to be understood that each of these enterprise systems may integrate, potentially different, identifying technologies, such as barcodes, RFID transponders, QR codes, BLE, etc. for identifying assets, good, items, etc. as they move within each of the domains. Particularly, within first domain **506** a first identification system **514** may be used, while within second domain **508** a second identification system **516**. First identification system **514** and second identification system **516** may generate respectively, a first set of data capture workflow **518** and a second set of data capture workflow **520**, which comprise information, data, etc. that are obtained by respective identification systems (e.g., first identification system **514** or second identification system **516**). The data capture workflow may include information indicative of identifiers (e.g., labeling information) scanned, obtained, etc. in the "field" that respective domains make up (e.g., a factory floor, a physical transport line, etc.).

[0117] In an example, first domain **506** and second domain **508** may, during the course of their operation, use identifiers (e.g., tags) for the purpose of tracking assets, goods, items, etc. within each domain. As shown, within first domain **506**, assets **522** may be assigned a set of identifiers **524**. When these items are "passed" on for example, from first domain **506** to second domain **508**, "relabeling" of assets **522** may occur such that second domain **508** uses another set of identifiers **526** that are incompatible or not interoperable with first domain **506**.

[0118] As shown, identity federation and visibility system/service **504** may be in connection with a first connector **528** of first domain **506** and a second connector **530** of second domain **508**. Both of these connectors may be deployed in the respective domains, for example, to support first enterprise system **510** and/or second enterprise system **512**. Notably, these may configure and maintain various identifying/tracking technologies. First domain **506** and second domain **508** may, during the course of their operation, use identifiers for the purpose of tracking assets, goods, items, etc. within each domain. Further, first connector **528** and second connector **530** are configured to discover types, functionality, etc. of identification, tracking, etc. technologies used inside within a domain. This may be implemented by, for example, monitoring and analyzing data from the data capture workflows, directly from OT fields, or through a combination of both. More specifically, the connectors detect implementation of a new identification system, either actively through the use of probing or passively. It is contemplated that some identification systems may advertise their capabilities actively (e.g., to the connectors), thereby enabling a push model as opposed to a discovery only mode. When a connector discovers an identification system, the connector may automatically initiate deployment and configuration of the corresponding plugins, as described herein, that enable communication with identity federation and visibility system/service **504** and, thus, identification binding system **502**.

[0119] Identification binding system **502** enables bonding of different identifiers and the association of the identifiers to specific assets (e.g., goods, items, products, etc.) throughout a supply chain, even across domains/organizations. In other words, identification binding system **502** allows various domains (e.g., first domain **506** and second domain **508**) to open, maintain, and close identity bindings for items (e.g., asset **522**) that are transited across a supply chain. That is, identification binding system **502**, in combination with identity federation and visibility system/service **504**, enables visibility of assets across domains (even if domains/organizations along a supply chain implement their own identifiers for the assets).

[0120] FIGS. **6**A-**6**B illustrates an example flow diagram **600** for an identification binding system for an identity federation and visibility system/service, in particular, how "relabeling" of one or more assets along a supply chain may be addressed. As shown in FIG. **6**A, a first domain **602** and a second domain **604** may implement its own labels for a particular asset **606** (e.g., an item, pallet, etc. that may be itself an aggregate of multiple items, pallets, etc.) labeling method. At step **608**, first domain **602** may add an identifier, for example, "ID1" to particular asset **606**, where particular asset **606** is to be shipped, sent, transported, etc. to second domain **604**. At step **610**, first domain **602** may scan the identifier ID1 and store a binding (item, ID1) in a particular enterprise system implemented by first domain **602** (e.g., first enterprise system **510**).

[0121] At step **612**, first domain **602** may generate an identity binding request for a domain pair between first domain **602** and second domain **604** (D1, D2), which indicates that identifiers for particular asset **606** are to be bound. This may be instrumented by sending a binding request to an identification binding system **614** (e.g., identification binding system **502**). It is contemplated that the binding request may be processed and forwarded through an identity federation and visibility system/service. Such binding request for the domain pair (D1, D2) may be initiated by a connector deployed at first domain **602**, by an enterprise system of first domain **602**, a client associated to first domain **602**. First domain **602** may open the binding request for a single item/identifier or for several items/identifiers simultaneously within a single request. The binding request may include metadata, information, etc. associated to the items themselves, such as an item type (e.g., a description of the items), item dimension and/or weight, a place where the item was shipped from, an indication of the identifier used within first domain **602**, etc. At step **616**, first domain **602** may ship the labeled/tagged item to second domain **604**.

[0122] As shown, at step **618**, second domain **604** may have several options for identifying particular asset **606** item within its own systems, procedures, etc. For instance, at step **620**, second domain **604** may scan the identifier, ID1, that has been physically added by first domain **602** to particular asset **606**. Alternatively, at step **622**, second domain **604** may, according to its own systems, procedures, etc., add a new identifier (e.g., ID2). In either case, second domain **604** may, at step **624**, scan and store that new identifier and the binding (item, ID2) in an enterprise system implemented by second domain **604**. It is contemplated that second domain **604** may, at step **626**, scan and store and also associate additional identifiers to particular asset **606**.

[0123] After these steps, at step **628**, second domain **604**, using a connection with an identity federation and visibility system/service, notify of identification binding system **614** of steps the second domain **604** has performed. Such notification may be sent for a single item/identifier; alternatively, second domain **604** may aggregate several items/identifiers simultaneously within a single notification. Second domain **604** may also include metadata, information, etc. in the notification regarding particular asset **606**. More specifi-

cally, the notification sent by second domain **604** may include one or more identifiers (e.g., only ID1, only ID2, both ID1 and ID2, or even additional identifiers) along with metadata linked to particular asset **606** associated to those identifiers.

[0124] At step **630** as shown in FIG. **6**B, identification binding system **614** may parse the input data sent by second domain **604** through a corresponding connector and, at step **632**, obtain identifier bindings associated with first domain **602** and second domain **604**. After this step, at step **634**, identification binding system **614** may apply two flows, based on information (in the binding request) received from second domain **604** in determining whether to close the binding request.

[0125] In one scenario, where second domain **604** notifies identification binding system **614** that second domain **604** will simply reuse identifier ID1 (assigned by first domain **602** to particular asset **606**). A connector of second domain **604** may retrieve metadata provided by first domain **602** when the connector generated the identity binding request (e.g., a description of the item, its dimensions and weight, its provenance, etc.) then push the metadata to an enterprise system of second domain **604**. This metadata may be appended to a record already stored in the enterprise system of second domain **604** for the identifier ID1, thereby allowing second domain **604** to associate a specific, now known, item to the identifier ID1 (e.g., particular asset **606**). In this case, the identity binding request opened by first domain **602** for the domain pair (D1, D2) may be closed, at step **636**, directly, and identification binding system **614** may notify, at step **638**, a connector at first domain **602** (by leveraging an identity federation and visibility system/service and a connector at first domain **602**). In such scenario, commonality and consistency of identifiers (i.e., the reuse of identifier ID1) across stakeholders or domains is supported.

[0126] In another scenario, second domain **604** may notify the identification binding system **614** that second domain **604** relabeled particular asset **606**, by using the new identifier, ID2, (which is assigned by second domain **604**). In this scenario, the identity binding request opened for the domain pair (D1, D2) may not be closed immediately, as a binding algorithm might be, at step **640**, executed to analyze and bond the metadata received as input from second domain **604** with the metadata provided by first domain **602** when it opened the identity binding request for the pair domain (D1, D2) in identification binding system **614**.

[0127] It is contemplated that this may be especially relevant when second domain **604** has not scanned the identifier ID1 (used by identification binding system **502**). In such a case, the notification sent to the identification binding system **614** from second domain **604** may come from the enterprise system of second domain **604**. This notification may include the identifier ID2 along with metadata associated to particular asset **606**, such as a description of the item, dimensions and weight of the item, a provenance of the item (e.g., shipped by first domain **602** from a particular location), etc. The binding algorithm may use or combine techniques such as string matching, semantic matching, fuzzy name matching, or others, to find a match, at step **642**, based on the metadata provided by both domains (first domain **602** and second domain **604**). In some embodiments, connectors may be endowed with data normalization capability, enabling them to structure the metadata provided by the domains before sending them to an identity federation and visibility system/service, which in turn, will forward the structured metadata to the identification binding system **614**. In some cases, data structuring and normalization capability might be embedded directly in an interface between the connector and the enterprise systems (e.g., at an API level). In other embodiments, the metadata provided by the domains as well as its structure might be standardized, thereby simplifying tasks of the binding algorithm. In any case, once the binding algorithm finds a match, the identification binding system **614** may notify an enterprise system of first domain **602** (by leveraging an identity federation and visibility system/service and a connector at first domain **602**).

[0128] In summary, when particular asset **606** is shipped from first domain **602** to second domain **604**, first domain **602** may send a binding request for the two domains to the identification binding system **614**. In turn, if a new label/ID is assigned to particular asset **606** at second domain **604** (i.e., "relabeled"), second domain **604** communicate with identification binding system **614** service to bind the plurality of IDs for particular asset **606**. Consequently, if first domain **602** issues a request for an update regarding particular asset **606**, identification binding system **614** may cross-reference the different IDs to provide an update to first domain **602**.

[0129] It is contemplated, that identification binding system **614** and corresponding notifications might be subject to one or more policies. That is, first domain **602** may be configured to indicate a set of identifiers and/or identifying technologies that should be subject to such bindings, and second domain **604** might exercise control on whether to provide such visibility to first domain **602**. In some embodiments, first domain **602** might open a binding request for several domains (e.g., D1, D2, D3, etc.) simultaneously, such that an identity federation and visibility system/service may enable notifications across several organizations (e.g., both upstream and downstream to its position on a supply chain). This may be performed on a peer-to-peer basis (e.g., one-to-one) or at once, concurrently, to multiple parties (e.g., one-to-n).

[0130] Furthermore, as described above herein, it is to be understood that deployment of various connectors at, for example, first domain **602** and second domain **604** may be required. That is, different domains may use different identification systems (e.g., RFID, barcodes, QR codes, etc.) and even deploy new identification systems over time. For instance, a warehouse may have recently deployed a new RFID scanning system. In such a case, a connector for that warehouse may detect this deployment and automatically configure access to the RFID scanning system for identity federation and visibility system/service **504** and identification binding system **502**. If the warehouse begins to identify items internally using RFID tags, the connector may enable the closing of bindings in identification binding system **502** (i.e., after the items have been relabeled (retagged) by the warehouse using RFID technology). Thus, identifiers may be bound across domains as described above, even when the identification technologies in the different domains evolve over time.

[0131] FIG. **7** illustrates an example simplified procedure for relabeling-resistant item identification/identification binding service for item identity federation and visibility as a service, in accordance with one or more embodiments described herein. In various embodiments, a non-generic, specifically configured device (e.g., device **200**) may per-

form procedure **700** by executing stored instructions (e.g., process **248** and process **249**), to provide relabeling-resistant item identification/identification binding service for item identity federation and visibility as a service. The procedure **700** may start at step **705**, and continues to step **710**, where, as described in greater detail above, a device may make a determination that an asset identification system has been deployed at an organization. For instance, the device may determine that the organization has implemented a particular tracking technology for physical assets. In some embodiments, the asset identification system may be based on radio-frequency identification, Bluetooth, barcodes, or quick response codes. As would be appreciated, the asset identification system may be different than one previously deployed at the organization.

[0132] At step **715**, as detailed above, the device may automatically deploy, based on the determination, a connector for the asset identification system to capture labeling information indicative of transport of one or more physical assets via the organization. In some embodiments, the connector may normalize information regarding the one or more physical assets. Further, as would be appreciated, in some embodiments, the connector may be deployed at a seaport terminal, an airport, a transport vehicle, a warehouse, or a distribution center.

[0133] At step **720**, the device may receive, from the connector, a binding request tag notification indicative of transport of a particular physical asset of the one or more physical assets via the organization. In some embodiments, the binding request tag notification may comprise an indication of physical asset type, physical asset dimension and/or weight, an origin of the one or more physical assets, or an identifier type used by the asset identification system. Further, as would be appreciated, in some embodiments, the connector may send the binding request tag notification to a particular connector of a different asset identification system.

[0134] At step **725**, as detailed above, the device may associate the binding request tag notification with one or more identifiers for the particular physical asset used by one or more other organizations. In some embodiments, the one or more other organizations may comprise one or more asset identification systems different than the asset identification system. That is, the one or more other organizations may track the physical assets using different identifiers and/or different technologies. In some embodiments, associating the binding request tag notification with the one or more identifiers for the particular physical asset used by the one or more other organizations may be based on a binding algorithm. For instance, the binding algorithm may comprise one or more of string matching, semantic matching, or fuzzy name matching. Procedure **700** then ends at step **730**.

[0135] It should be noted that while certain steps within procedure **700** may be optional as described above, the steps shown in FIG. **7** are merely examples for illustration, and certain other steps may be included or excluded as desired. Further, while a particular order of the steps is shown, this ordering is merely illustrative, and any suitable arrangement of the steps may be utilized without departing from the scope of the embodiments herein.

[0136] The techniques described herein, therefore, allow for relabeling-resistant item identifier used across domains. One of the main advantages of the techniques herein is that it the techniques described herein do not require an overall convention or a priori agreement on technologies, standards, etc. for identifiers across domains of a supply chain. Similarly, no peer-to-peer integration among identification system(s) and/or backend system(s) of particular domains is required. In addition, the techniques herein provide improved visibility of assets and inventory across different domains of a supply chain. Further, the techniques herein enable automatic detection of identification system(s) at a particular domain, so that an identity federation and visibility service(s), within a supply chain, may automatically configure and obtain access to item, goods, etc. information of the particular domain.

[0137] While there have been shown and described illustrative embodiments for relabeling-resistant item identification/identification binding service with a supply chain, it is to be understood that various other adaptations and modifications may be made within the intent and scope of the embodiments herein. For example, while specific types of identification technologies are described herein (e.g., RFID, BLE, etc.), the techniques herein are not limited as such and can be applied to any form of identification technology. Further, while the techniques herein are described primarily with respect to federating item identities and providing item visibility as well as relabeling-resistant item identification across a supply chain, the techniques herein are not limited as such and can also be extended to federating workforce identity and visibility, connected asset identity and visibility, and the like.

[0138] The foregoing description has been directed to specific embodiments. It will be apparent, however, that other variations and modifications may be made to the described embodiments, with the attainment of some or all of their advantages. For instance, it is expressly contemplated that the components and/or elements described herein can be implemented as software being stored on a tangible (non-transitory) computer-readable medium (e.g., disks/CDs/RAM/EEPROM/etc.) having program instructions executing on a computer, hardware, firmware, or a combination thereof. Accordingly, this description is to be taken only by way of example and not to otherwise limit the scope of the embodiments herein. Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true intent and scope of the embodiments herein.

What is claimed is:

1. A method, comprising:

making, at a device, a determination that an asset identification system has been deployed at an organization;

automatically deploying, by the device and based on the determination, a connector for the asset identification system to capture labeling information indicative of transport of one or more physical assets via the organization;

receiving, by the device and from the connector, a binding request tag notification indicative of transport of a particular physical asset of the one or more physical assets via the organization; and

associating, by the device, the binding request tag notification with one or more identifiers for the particular physical asset used by one or more other organizations.

2. The method as in claim **1**, wherein the asset identification system is based on radio-frequency identification, Bluetooth, barcodes, or quick response codes.

3. The method as in claim 1, wherein the one or more other organizations comprise one or more asset identification systems different than the asset identification system.

4. The method as in claim 1, wherein the binding request tag notification comprises an indication of physical asset type, physical asset dimension and/or weight, an origin of the one or more physical assets, or an identifier type used by the asset identification system.

5. The method as in claim 1, wherein associating, by the device, the binding request tag notification with the one or more identifiers for the particular physical asset used by the one or more other organizations is based on a binding algorithm.

6. The method as in claim 5, wherein the binding algorithm comprises one or more of string matching, semantic matching, or fuzzy name matching.

7. The method as in claim 1, wherein the connector normalizes information regarding the one or more physical assets.

8. The method as in claim 1, wherein the connector sends the binding request tag notification to a particular connector of a different asset identification system.

9. The method as in claim 1, wherein the connector is deployed at a seaport terminal, an airport, a transport vehicle, a warehouse, or a distribution center.

10. The method as in claim 1, wherein the asset identification system is different than one previously deployed at the organization.

11. An apparatus, comprising:

one or more interfaces;

a processor coupled to the one or more interfaces and configured to execute one or more processes; and

a memory configured to store a process that is executable by the processor, the process when executed configured to:

make a determination that an asset identification system has been deployed at an organization;

automatically deploy, based on the determination, a connector for the asset identification system to capture labeling information indicative of transport of one or more physical assets via the organization;

receive, from the connector, a binding request tag notification indicative of transport of a particular physical asset of the one or more physical assets via the organization; and

associate the binding request tag notification with one or more identifiers for the particular physical asset used by one or more other organizations.

12. The apparatus as in claim 11, wherein the asset identification system is based on radio-frequency identification, Bluetooth, barcodes, or quick response codes.

13. The apparatus as in claim 11, wherein the one or more other organizations comprise one or more asset identification systems different than the asset identification system.

14. The apparatus as in claim 11, wherein the binding request tag notification comprises an indication of physical asset type, physical asset dimension and/or weight, an origin of the one or more physical assets, or an identifier type used by the asset identification system.

15. The apparatus as in claim 11, wherein to associate the binding request tag notification with the one or more identifiers for the particular physical asset used by the one or more other organizations is based on a binding algorithm.

16. The apparatus as in claim 15, wherein the binding algorithm comprises one or more of string matching, semantic matching, or fuzzy name matching.

17. The apparatus as in claim 11, wherein the connector normalizes information regarding the one or more physical assets.

18. The apparatus as in claim 11, wherein the connector sends the binding request tag notification to a particular connector of a different asset identification system.

19. The apparatus as in claim 11, wherein the connector is deployed at a seaport terminal, an airport, a transport vehicle, a warehouse, or a distribution center.

20. A tangible, non-transitory, computer-readable medium storing program instructions that cause a device to execute a process comprising:

making, by the device, a determination that an asset identification system has been deployed at an organization;

automatically deploying, based on the determination, a connector for the asset identification system to capture labeling information indicative of transport of one or more physical assets via the organization;

receiving, from the connector, a binding request tag notification indicative of transport of a particular physical asset of the one or more physical assets via the organization; and

associating the binding request tag notification with one or more identifiers for the particular physical asset used by one or more other organizations.

* * * * *