



- (51) International Patent Classification:  
*H04L 29/02* (2006.01)
- (21) International Application Number:  
PCT/US2014/057935
- (22) International Filing Date:  
28 September 2014 (28.09.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
61/884,006 28 September 2013 (28.09.2013) US
- (71) Applicant: MCAFEE INC. [US/US]; 2821 Mission College Boulevard, Santa Clara, California 95054 (US).
- (72) Inventor; and  
(71) Applicant (for US only): ARKIN, Ofir [IL/IL]; 15/3 Hasitvanit, 49745 Petach Tikva (IL).
- (74) Agent: CRANDALL, Sean C.; Patent Capital Group, c/o CPA Global, P.O. Box 52050, Minneapolis, Minnesota 55402 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: LOCATION SERVICES ON A DATA EXCHANGE LAYER




FIG. 2

(57) Abstract: In an example, a security-connected platform is provided on a data exchange layer (DXL), which provides messaging on a publish-subscribe model. The DXL provides a plurality of DXL endpoints connected via DXL brokers. DXL endpoints communicate via DXL messages divided by topics. An asset management engine may be provided to use DXL messaging to query DXL endpoints so as to logically and physically map a network.

## LOCATION SERVICES ON A DATA EXCHANGE LAYER

CROSS-REFERENCE TO RELATED APPLICATIONS

**[0001]** This application claims priority to U.S. Provisional Application 61/884,006, filed 28 September 2013, titled "SECURITY-CONNECTED PLATFORM," which is incorporated by reference in its entirety. Co-pending PCT Application PCT/US2013/076570, titled "CONTEXT-AWARE NETWORK ON A DATA EXCHANGE LAYER," filed December 19, 2013, is also incorporated herein by reference.

**[0002]** Co-pending PCT International Application Serial No. PCT/US2014/057934, filed on September 28, 2014 and entitled "SECURITY-CONNECTED FRAMEWORK" is also incorporated herein by reference.

FIELD OF THE DISCLOSURE

**[0003]** This application relates to the field of enterprise security, and more particularly to a security-connected framework for a data exchange layer.

BACKGROUND

**[0004]** An enterprise service bus (ESB) is a software-based network architecture that provides a medium of data exchange over a service-oriented architecture. In some embodiments, ESB is a special case of a client-server software architecture in which clients may route messages through the server.

**[0005]** Software, binaries, executables, advertising, web pages, documents, macros, executable objects, and other data provided to users (collectively "executable objects") may include security flaws and privacy leaks that are subject to exploitation by malware. As used throughout this Specification, malicious software ("malware") may include a virus, Trojan, zombie, rootkit, backdoor, worm, spyware, adware, ransomware, dialer, payload, malicious browser helper object, cookie, logger, or similar application or part of an application designed to take a potentially-unwanted action, including by way of non-limiting example, data destruction, covert data collection, covert communication, browser hijacking, network proxy hijacking or redirection, covert tracking, data logging, keylogging, excessive or deliberate barriers to removal, contact harvesting, unwanted use of premium services, and unauthorized self-propagation. In some cases, malware may also include legitimate software that includes

inadvertent security flaws that cause or enable malware behavior. "Malware behavior" is defined as any behavior that qualifies an application as malware or grayware. Some existing systems are configured to identify and block malware, for example by maintaining databases of known malware.

**[0006]** In addition to executable objects, computing devices may encounter static objects, which are not intended to change the operating state of a computer. As a class, executable objects and static objects may be referred to simply as "objects." An enterprise security concern is the classification of objects' malware status.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

**[0007]** The present disclosure is best understood from the following detailed description when read with the accompanying figures. It is emphasized that, in accordance with the standard practice in the industry, various features are not drawn to scale and are used for illustration purposes only. In fact, the dimensions of the various features may be arbitrarily increased or reduced for clarity of discussion.

**[0008]** FIGURE 1 is a block diagram of a context-aware network according to one or more examples of the present Specification.

**[0009]** FIGURE 2 is a block diagram of a data exchange layer according to one or more examples of the present Specification.

**[0010]** FIGURE 3 is a block diagram of a DXL endpoint according to one or more examples of the present Specification.

**[0011]** FIGURE 4 is a block diagram of a server according to one or more examples of the present Specification.

**[0012]** FIGURE 5 is a block diagram of a producer-consumer architecture according to one or more examples of the present Specification.

**[0013]** FIGURE 6 is a block diagram of a logical network map according to one or more examples of the present Specification.

**[0014]** FIGURE 7 is a block diagram of a method according to one or more examples of the present Specification.

**[0015]** FIGURE 8 is a block diagram of a method according to one or more examples of the present Specification.

**[0016]** FIGURE 9 is a block diagram of a method according to one or more examples of the present Specification.

**[0017]** FIGURE 10 is a block diagram of a method according to one or more examples of the present Specification.

**[0018]** FIGURE 11 is a block diagram of a method according to one or more examples of the present Specification.

## DETAILED DESCRIPTION OF THE EMBODIMENTS

### OVERVIEW

**[0019]** In an example, a security-connected platform is provided on a data exchange layer (DXL), which provides messaging on a publish-subscribe model. The DXL provides a plurality of DXL endpoints connected via DXL brokers. DXL endpoints communicate via DXL messages divided by topics. An asset management engine may be provided to use DXL messaging to query DXL endpoints so as to logically and physically map a network.

### EXAMPLE EMBODIMENTS OF THE DISCLOSURE

**[0020]** The following disclosure provides many different embodiments, or examples, for implementing different features of the present disclosure. Specific examples of components and arrangements are described below to simplify the present disclosure. These are, of course, merely examples and are not intended to be limiting. Further, the present disclosure may repeat reference numerals and/or letters in the various examples. This repetition is for the purpose of simplicity and clarity and does not in itself dictate a relationship between the various embodiments and/or configurations discussed.

**[0021]** Different embodiments may have different advantages, and no particular advantage is necessarily required of any embodiment.

**[0022]** In an example, the present Specification provides a data exchange layer (DXL), which may operate on a lightweight messaging-based communications infrastructure such as ESB and be configured to allow endpoints to share contextual data. The DXL is built on top of messaging technology, such as an ESB, which enables multiple different use cases and capabilities (context-aware and adaptive security by sharing context between multiple different products allowing for security components to operate as one to immediately share

relevant data between endpoint, gateway, and other security products enabling security intelligence and adaptive security (immediately change behavior according to information you are getting); superior command and control of endpoints and many other use cases).

**[0023]** DXL messaging enable real-time, bi-directional communications infrastructure allowing products and solutions to integrate with each other using a single application programming interface (API). Each device may share essentially any data it likes over the DXL fabric, while the devices themselves are only loosely coupled, and need not operate on common or standardized protocols.

**[0024]** Examples of DXL message types include publish-subscribe notifications, query-response, and push notification. Devices may also share events, including security-related events, contextual information (about the node, the identity of its user, the applications used, the location for which they are found, etc.), commands, tasks, and policies, by way of non-limiting example.

**[0025]** In one example, client devices may variously be classified as either “producers” or “consumers” of messages, wherein producers provide relevant security messages, and consumers receive those messages. A device’s role as either a producer or consumer need not be static, and in general any device can be either a producer or a consumer, depending on context. A devices role as producer or consumer may also change over time or circumstances. Thus, one device may be both a producer and consumer of certain message topics; a second device may be a producer for some topics but not for other topics; a third device may be a consumer for certain topics and a producer for other topics.

**[0026]** As long as clients are subscribed to one or more topics (wherein each topic includes a different class or species of message) a communication path exists to the consumers. Connection initiation is performed by clients.

**[0027]** As used throughout this Specification, a “DXL message” includes any message communicated over a DXL ESB. Each DXL message includes at least one “topic”; a topic, representing a subject or classification of the message. Topics may include threat intelligence, context, and events by way of non-limiting example.

**[0028]** By design, each DXL endpoint is configured to subscribe to at least one DXL message topic (at a minimum, a private destination topic for routing DXL messages to that

DXL endpoint). This enables bi-directional communications, over the DXL, between a management platform, for example, and one or more of its managed clients.

**[0029]** FIGURE 1 is a network-level block diagram of a context-aware network 100 with DXL capability. According to this example, a plurality of DXL endpoints 120 is connected to a DXL enterprise service bus (ESB) 130 (FIGURE 2). DXL ESB 130 is an example of a DXL fabric, and may be provided on top of an existing network, such as a local area network (LAN). DXL ESB 130 need not be a particular physical bus, or even reside on a physical network. Rather, DXL ESB 130 may span multiple physical network and subnetworks. Conceptually, DXL ESB 130 is simply the “fabric” over which DXL endpoints 120 share DXL messages, wherein each DXL message includes a topic, and only DXL endpoints 120 that subscribe to that topic receive and/or act on messages for that topic.

**[0030]** DXL endpoint 120 may be any suitable computing device. In the example of FIGURE 1, DXL endpoints are generically represented as computing devices. In one example, DXL endpoint 120-1 may be an embedded device, such as a network security sensor. DXL endpoint 120-2 may be a virtual machine. DXL endpoint 120-3 may be a laptop or notebook computer. It should also be noted that DXL endpoints are not limited to end-user or client devices.

**[0031]** Domain services controller (DSC) 180 may be considered a special case of a DXL endpoints 120. Other network elements may include a domain master, a joint threat intelligence (JTI) server, default gateway, proxy, and threat intelligence service. Any network element may join DXL ESB 130 as a DXL endpoint.

**[0032]** Network elements configured to operate on or with DXL ESB 130 may be referred to as “DXL endpoints.” These may include, in an example, DXL endpoints 120, and DXL broker 110.

**[0033]** DXL broker 110 may be configured to provide DXL messaging services over DXL ESB 130, such as maintaining DXL routing tables and delivering messages.

**[0034]** DSC 180 may be configured to provide domain services, such as auxiliary and support services, antivirus services, and/or command and control services. In one particular example, DSC 180 may include an asset management engine and/or physical network topology engine, which may be configured to logically and physically map and track network

assets as described herein. Additionally, or in another example, DSC 180 may include security management software (SMS), which may provide network command and control functions.

**[0035]** In an example, DXL endpoints connect to DXL ESB 130. Each DXL endpoint is assigned a distinct identity, and authenticates itself to DXL ESB 130 upon startup, for example via a certificate or other secure token. DXL endpoints may establish one-to-one communications via DXL ESB 130, for example by sending a DXL message addressed to a DXL endpoint with a specific identity. This enables DXL endpoints to communicate with each other without having to establish a point-to-point network connection. In an example, this is analogous to a person-to-person phone call.

**[0036]** In another example, DXL may provide a publish-subscribe framework in which certain DXL endpoints “subscribe” to messages of a certain type. When a DXL endpoint “publishes” a message of that type on DXL ESB 130, all subscribers may process the message, while non-subscribers may safely ignore it. In an example, this is analogous to a podcast subscription service. In yet another example, DXL may provide a request-response framework.

**[0037]** FIGURE 2 is a network diagram disclosing a DXL ESB 130 on a context aware network 100 according to one or more examples of the present Specification. In this example, DXL broker 110-1 may be designated as the “hub,” while DXL broker is 110-2, 110-3, 110-4, and 110-5 may be designated as “spokes.” in an example, all DXL traffic that passes through spoke will be forwarded to the hub, which will distribute the traffic to other spokes. Designation of a DXL broker 110 as the hub may be accomplished via any suitable means, such as selecting the hub based on MAC ID, IP address, or network proximity to a domain master.

**[0038]** If DXL broker 110-1 goes offline, another hub may be at least temporarily needed. In that case, another hub may be elected. When DXL broker 110-1 comes back online, it may resume its duties as a hub, or may act as a spoke, depending on network topology and design considerations.

**[0039]** In another example, spokes may form a temporary mesh network upon using effectively to DXL broker 110-1. In yet other embodiments, DXL broker is 110 may be configured to operate full time in a mesh configuration.

**[0040]** Additional extensibility may be provided by bridging DXL ESB 130 across disparate networks, enabling data to be exchanged over larger networks, including the Internet.

**[0041]** DXL brokers 110 may be configured to enhance the efficiency of DXL ESB 130. For example, each DXL broker 110 may maintain a list of subscribed topics for each connected DXL endpoint 120. The DXL broker 110 then itself subscribes to those topics. Each DXL broker 110 also maintains a list of which topics each other DXL broker 110 subscribes to. When a DXL broker 110 receives a message from any of its DXL endpoints 120, the DXL broker 110 determines which other brokers are subscribed to the topic of the DXL message, and forwards the message only to those brokers.

**[0042]** FIGURE 3 is a block diagram of client device 120 according to one or more examples of the present Specification. Client device 120 may be any suitable computing device. In various embodiments, a “computing device” may be or comprise, by way of non-limiting example, a computer, embedded computer, embedded controller, embedded sensor, personal digital assistant (PDA), laptop computer, cellular telephone, IP telephone, smart phone, tablet computer, convertible tablet computer, handheld calculator, or any other electronic, microelectronic, or microelectromechanical device for processing and communicating data

**[0043]** Client device 120 includes a processor 310 connected to a memory 320, having stored therein executable instructions for providing an operating system 322 and DXL client 324. Other components of client device 120 include a storage 350, network interface 360, and peripheral interface 340.

**[0044]** In an example, processor 310 is communicatively coupled to memory 320 via memory bus 370-3, which may be for example a direct memory access (DMA) bus by way of example, though other memory architectures are possible, including ones in which memory 320 communicates with processor 310 via system bus 370-1 or some other bus. Processor 310 may be communicatively coupled to other devices via a system bus 370-1. As used throughout this Specification, a “bus” includes any wired or wireless interconnection line, network, connection, bundle, single bus, multiple buses, crossbar network, single-stage network, multistage network or other conduction medium operable to carry data, signals, or power between parts of a computing device, or between computing devices. It should be

noted that these uses are disclosed by way of non-limiting example only, and that some embodiments may omit one or more of the foregoing buses, while others may employ additional or different buses.

**[0045]** In various examples, a “processor” may include any combination of hardware, software, or firmware providing programmable logic, including by way of non-limiting example a microprocessor, digital signal processor, field-programmable gate array, programmable logic array, application-specific integrated circuit, or virtual machine processor.

**[0046]** Processor 310 may be connected to memory 320 in a DMA configuration via DMA bus 370-3. To simplify this disclosure, memory 320 is disclosed as a single logical block, but in a physical embodiment may include one or more blocks of any suitable volatile or non-volatile memory technology or technologies, including for example DDR RAM, SRAM, DRAM, cache, L1 or L2 memory, on-chip memory, registers, flash, ROM, optical media, virtual memory regions, magnetic or tape memory, or similar. In certain embodiments, memory 320 may comprise a relatively low-latency volatile main memory, while storage 350 may comprise a relatively higher-latency non-volatile memory. However, memory 320 and storage 350 need not be physically separate devices, and in some examples may represent simply a logical separation of function. It should also be noted that although DMA is disclosed by way of non-limiting example, DMA is not the only protocol consistent with this Specification, and that other memory architectures are available.

**[0047]** Storage 350 may be any species of memory 320, or may be a separate device, such as a hard drive, solid-state drive, external storage, redundant array of independent disks (RAID), network-attached storage, optical storage, tape drive, backup system, cloud storage, or any combination of the foregoing. Storage 350 may be, or may include therein, a database or databases or data stored in other configurations, and may include a stored copy of operational software such as operating system 322 and software portions of DXL client 324. Many other configurations are also possible, and are intended to be encompassed within the broad scope of this Specification.

**[0048]** Network interface 360 may be provided to communicatively couple client device 120 to a wired or wireless network. A “network,” as used throughout this Specification, may include any communicative platform operable to exchange data or

information within or between computing devices, including by way of non-limiting example, an ad-hoc local network, an internet architecture providing computing devices with the ability to electronically interact, a plain old telephone system (POTS), which computing devices could use to perform transactions in which they may be assisted by human operators or in which they may manually key data into a telephone or other suitable electronic equipment, any packet data network (PDN) offering a communications interface or exchange between any two nodes in a system, or any local area network (LAN), metropolitan area network (MAN), wide area network (WAN), wireless local area network (WLAN), virtual private network (VPN), intranet, or any other appropriate architecture or system that facilitates communications in a network or telephonic environment.

**[0049]** Network interface 360 may be configured to communicatively couple client device 120 to a DXL broker 110.

**[0050]** DXL client 324 and runtime agent 326 are examples of “engines.” As used throughout this agent, an “engine” includes one or more logic elements, including hardware software, and/or firmware, configured to perform or operable for performing the function of the engine. An engine may be self-contained on a single device, or may span multiple devices. Furthermore, a single device may include a plurality of engines. For ease of discussion, the engines disclosed herein are shown, by way of example only, as a software subroutine running from memory. In one example, an engine is a utility or program that provides its host device with the necessary APIs and interfaces for performing a particular function. An engine may be, in other examples, embodied in hardware, software, firmware, or some combination thereof. For example, in some cases, an engine may include a special integrated circuit designed to carry out a method or a part thereof, and may also include software instructions operable to instruct a processor to perform the method. In some cases, the engine may run as a “daemon” process. A “daemon” may include any program or series of executable instructions, whether implemented in hardware, software, firmware, or any combination thereof, that runs as a background process, a terminate-and-stay-resident program, a service, system extension, control panel, bootup procedure, BIOS subroutine, or any similar program that operates without direct user interaction. In other examples, an engine may include other hardware and software, including interactive or user-mode software, which may be provided

in conjunction with, in addition to, or instead of the engine to perform methods according to this Specification.

**[0051]** In one example, an engine comprises executable instructions stored on a non-transitory computer-readable medium operable, when executed, for performing the methods of the engine. At an appropriate time, such as upon booting the host device, or upon a command from an operating system or a user, the processor may retrieve from storage a copy of the engine load it into memory. The processor may then iteratively execute the instructions of the engine.

**[0052]** Peripheral interface 340 may be configured to interface with any auxiliary device that connects to client device 120 but that is not necessarily a part of the core architecture of client device 120. A peripheral may be operable to provide extended functionality to client device 120, and may or may not be wholly dependent on client device 120. In some cases, a peripheral may be a computing device in its own right. Peripherals may include input and output devices such as displays, terminals, printers, keyboards, mice, modems, network controllers, sensors, transducers, actuators, controllers, data acquisition buses, cameras, microphones, speakers, or external storage by way of non-limiting example.

**[0053]** FIGURE 4 is a block diagram of a server 400 according to one or more examples of the present Specification. Server 400 may be any suitable computing device, as described in connection with FIGURE 3. In general, the definitions and examples of FIGURE 3 may be considered as equally applicable to FIGURE 3, unless specifically stated otherwise. However, a server 400 may usefully and conceptually be used to refer to a class of computing devices that provide server functions in client-server architectures. Thus, in certain embodiments, DSC 180, DXL broker 110, and other devices described herein may all be various embodiments of servers 400. They are discussed here together for purposes of clarity, and because servers 400 may share many common elements. However, it should be particularly noted that a number of “engines” are described herein. Each engine may include one or more hardware, software, and/or firmware logic elements configured to provide server functions. The number of server engines, and which server engines to include, are determined by the server’s function.

**[0054]** Server 400 includes a processor 410 connected to a memory 420, having stored therein executable instructions for providing an operating system 422, and one or

more of a physical network topology engine 424, a runtime agent 326, and an asset management engine 428. Other components of server 400 include a storage 450, network interface 460, and peripheral interface 440.

**[0055]** In an example, processor 410 is communicatively coupled to memory 420 via memory bus 470-3, which may be for example a direct memory access (DMA) bus. Processor 410 may be communicatively coupled to other devices via a system bus 470-1.

**[0056]** Processor 410 may be connected to memory 420 in a DMA configuration via DMA bus 470-3.

**[0057]** Storage 450 may be any species of memory 420, or may be a separate device, as described in connection with storage 350 of FIGURE 3. Storage 450 may be, or may include therein, a database or databases or data stored in other configurations, and may include a stored copy of operational software such as operating system 422 and software portions of physical network topology engine 424, DXL services engine 426, runtime agent 326, and asset management engine 428. Many other configurations are also possible, and are intended to be encompassed within the broad scope of this Specification.

**[0058]** Network interface 460 may be provided to communicatively couple Server 400 to a wired or wireless network.

**[0059]** Physical network topology engine 424, DXL services engine 426, runtime agent 326, and asset management engine 428 are embodiments of engines, as described in FIGURE 3 above.

**[0060]** Asset management engine 428 and physical network topology engine 424 may carry out methods as described in FIGURES 6 – 11 below. In one example, asset management engine 428 includes or interfaces with a McAfee® Asset Manager (MAM). In some cases, asset management engine 428 may also provide location services, as described in more detail in connection with FIGURES 6 – 11.

**[0061]** Peripheral interface 440 may be configured to interface with any auxiliary device that connects to Server 400 but that is not necessarily a part of the core architecture of Server 400. A peripheral may be operable to provide extended functionality to Server 400, and may or may not be wholly dependent on Server 400. In some cases, a peripheral may be a computing device in its own right. Peripherals may include, by way of non-limiting

examples, any of the devices discussed in connection with peripheral interface 340 of FIGURE 3.

**[0062]** FIGURE 5 is a network-level block diagram of a data exchange layer according to one or more examples of the present Specification. In the example of FIGURE 5, five DXL brokers 110 provide services to DXL endpoints 120.

**[0063]** Specifically, DXL endpoints 120-12 and 120-14 are connected to DXL broker 110-1. DXL endpoints 120-32 and 120-34 are connected to DXL broker 110-3. DXL endpoints 120-42 and 120-44 are connected to DXL broker 110-4. DXL endpoints 120-52 and 120-54 are connected to DXL broker 110-5. DXL endpoint 120-22 is connected to DXL broker 110-2.

**[0064]** FIGURES 6 – 11 provide an example of location services provided by the security-connected framework of the present Specification.

**[0065]** As used in this Specification, a “logical location” includes a collection of one or more interconnected routers and/or switches, a collection of one or more IP subnets, or other topologies, which may map to one or more geo-locations. For example, an enterprise network may include a plurality of geographical locations. A single enterprise network may include devices spread across a large variety of geo-locations. In one example, an enterprise includes offices and/or data centers in Santa Clara, California; Plano, Texas; Bangalore, India; Tel Aviv, Israel; and Tokyo, Japan. In this case, there may be a need for a DXL device in Tel Aviv to communicate with a DXL device in Plano.

**[0066]** Location context may play a role in such transactions, and is part of a broader collection of situational and environmental information that plays a unique role in satisfying multiple use cases. These include, by way of non-limiting example:

**[0067]** Gathering complete understanding of the true scope and size of an organization (in this example, spanning from Santa Clara westward to Tokyo).

**[0068]** Dividing geo-locations and network locations into a plurality of security zones.

**[0069]** Ensuring completeness of an operation.

**[0070]** Detecting network locations a security appliance does not protect and/or cover (e.g. IP subnets or devices that do not include runtime security agent).

**[0071]** Supporting location-dependent product configurations.

**[0072]** Resolving situations where conflicting IP addressing exists.

**[0073]** Providing network-based services according to geographical or logical location (e.g. providing location-based services).

**[0074]** Enabling products operating in the same geo-location to form a collaboration.

**[0075]** In certain embodiments, properly representing a location requires, at a minimum, acquiring a list of all routers operating at the location along with the following information per router:

**[0076]** An interface list, including the configuration of each interface.

**[0077]** A routing table.

**[0078]** A content-addressable memory (CAM) table, including media access control (MAC) address-to-virtual local area network (VLAN) mapping, and port-of-arrival mapping. The CAM table may also include address resolution protocol (ARP) information, which maps MAC addresses to IP addresses.

**[0079]** With this information, DSC 180 can compile a list of all IP subnets in use at a location, and can verify that the list of routers is indeed part of the same location.

**[0080]** Because there may be conflicting IP address allocations within the same network, representing an IP subnet by its IP address will not necessarily uniquely identify one subnet from another. An alternative method of identifying subnets is to correlate the subnet's IP address to the MAC address of the router interface set as the default gateway for the IP subnet. This coupling creates a unique identifier that resolves the conflicting IP addressing representation issue. For example:

$$MAC\ Address(x)_{Router(z)} :: IP\ Subnet(y) \text{ Where } MAC\ Address(x) \in Router(z)$$

**[0081]** Common configurations may include the following:

**[0082]** A router's interface may serve as the default gateway for multiple IP subnets

**[0083]** A router may be configured using multiple interfaces acting as default gateways for multiple different IP subnets

**[0084]** Multiple routers may operate in the same location

**[0085]** An IP subnet may be defined in-between two or more routers

**[0086]** In one example of the present Specification, a logical location is specified as follows:

```

Location =
{
    (MAC Address(1)Router(1)::IP Subnet(1) ... n) ... MAC Address(n)Router(1)::IP
    Subnet(n)) U
    (MAC Address(1)Router(z)::IP Subnet(1) ... MAC Address(n)Router(z)::IP
    Subnet(n)) U
    (MAC Address(x)Router(1), MAC Address(y)Router(k), ... MAC Address(z)Router(z)::IP
    Subnet(j))
}

```

**[0087]** FIGURE 6 is a network diagram of a location ( $n$ ) 600 according to one or more examples of the present Specification. Location ( $n$ ) 600 includes five subnets, namely subnet N 610, subnet Y 620, subnet X 630, subnet Z 640, and subnet T 650.

**[0088]** In this example, routers  $R_1$  and  $R_2$  may be configured to act as default gateways to their respective subnets.

**[0089]** Router  $R_1$  provides multiple interfaces ( $j, k, i$ ), serving as the default gateways for IP subnets  $X, Z$ , and  $T$  respectively.

**[0090]** Switch  $S_4$  interfaces with router  $R_1$  on interface  $R_{1j}$  and controls subnet  $Z$ , including endpoints  $Z_1$  and  $Z_2$ .

**[0091]** Switches  $S_1$  and  $S_3$  control subnet  $X$ . Switch  $S_1$  connects to router  $R_1$  on interface  $R_{1i}$ . Subnet  $X$  includes endpoints  $X_1, X_2, X_3$ , and  $X_4$  connect to switch  $S_1$ . Endpoints  $X_5, X_6$ , and  $X_7$  connect to switch  $S_3$ , which also interfaces to switch  $S_1$ .

**[0092]** Subnet  $T$  is logically defined on interface  $R_{1k}$ . No endpoints exist on this subnet. However, router  $R_2$  is communicatively coupled to router  $R_1$  on interface  $R_{2k}$ .

**[0093]** Router  $R_2$  also includes logical interfaces  $i$  and  $k$ , connecting to subnets  $Y$  and  $T$  respectively. Connection  $R_{2j}$  is not defined or configured in this FIGURE.

**[0094]** Subnet  $Y$  includes switch  $S_2$ , which communicatively couples to interface  $R_{2i}$ . Endpoints  $Y_1, Y_2$ , and  $Y_3$  are communicatively coupled to switch  $S_2$ .

**[0095]** Switch  $S_2$  also communicatively couples to subnet  $N$  via switch  $S_5$ . Note that in this example, subnet  $N$  does not have a direct or single-hop link to either router  $R_1$  or  $R_2$  on any logical interface.

**[0096]** Non-limiting examples of conditions on this network include:

**[0097]** Two interconnected routers, namely  $R_1$  and  $R_2$ .

**[0098]** Interface  $R_{2i}$  services both subnet  $Y$  and IP subnet  $N$ .

**[0099]** Interfaces  $R_{1k}$  and  $R_{2k}$  share a single subnet, namely subnet  $T$ .

**[0100]** Location (n) may therefore be represented as follows:

```
Location (n) =  
{  
MAC Address (R1i)::IP Subnet (X), MAC Address (R1j)::IP Subnet (Z), MAC  
Address (R2i)::IP Subnet (Y), MAC Address (R2i)::IP Subnet (N), (MAC Address  
(R1k), MAC Address (R2k))::IP Subnet (T)  
}
```

**[0101]** A runtime agent 324 may be configured in each device to collect and report relevant information. In some cases, a plugin may be provided to runtime agent 324 for collecting such information. Information gathered from runtime agents 324 running on a plurality of devices may enable an enterprise to create a comprehensive and accurate list of all IP subnets in the enterprise. An enterprise security administrator can then manually map IP subnets to geo-locations.

**[0102]** Advantageously, this addresses some inherent difficulties in mapping IP subnets to geo-locations, including:

**[0103]** Many—sometimes thousands or more—IP subnets may be used within a network. This makes manual mapping difficult.

**[0104]** Manually associating IP subnets to geo-locations can be error prone.

**[0105]** A strong understanding of the network topology cannot be assumed because poor understanding of the network topology may be the very problem that needs to be addressed by mapping

**[0106]** Distinguishing between internal IP subnets to external IP subnets requires detailed knowledge of the IT infrastructure and the physical topology of the network.

**[0107]** Thus, in some cases, automated mapping may be preferable to manual mapping, particularly for large enterprises. In one example, a network device uses the SNMP protocol to remotely interrogate all routers and/or switches to determine the connectivity relationship that exists between networking devices operating in the same location building a mapping of all locations part of a network.

**[0108]** An asset management engine 428 may be used to build and maintain the physical network topology of a network. The output provided by asset management engine 428 is a list of discovered locations, and their physical network topology disclosing the connectivity relationship between routers and switches operating in the same location (e.g.

layer-2 and layer-3 connectivity information). In one example, the output provided by asset management engine 428 marks the borders between internal to external locations.

**[0109]** Asset management engine 428 may include a physical network topology engine to discover and maintain the logical location of a device, the IP address of the switch, and the port a device is connected to. These data may be provided for each authorized device operating on DXL ESB 130. Asset management engine 428 may gather information in real-time or near real-time as a device joins the network. This enables asset management engine 428 to continuously maintain an accurate view of the physical network topology of the network.

**[0110]** In some cases, the virtual map that asset management engine 428 builds may not align perfectly with the actual physical network topology, particularly when some devices do not include a runtime agent 324 or when routers cannot be queried. In that case, manual mapping may be used to correlate and align the virtual map to the actual map.

**[0111]** When distributed applications provide multiple service instances (such as multiple policy orchestration servers, multiple security information and event management (SIEM) receivers that collect logs from different parts of the network). There is a need to map between the different service instances (one or many), per distributed application, and the location(s) each instance is to serve. One use case is automatic configuration of services per a device's location, in which case asset management engine 428 holds the relevant information.

**[0112]** In one example, mapping between locations and services is a manual process wherein a user associates a location with one or more instances of different services. Asset management engine 428 may provide an appropriate graphical user interface for performing this function.

**[0113]** FIGURE 7 is a block diagram of a method of an asset management engine 428 providing location services for managed DXL clients 110.

**[0114]** Before any DXL operations take place, DXL endpoint 120 may join a network. For example, DXL endpoint 120 may boot up and communicate with DHCP server 720, which assigns to DXL endpoint 120 an IP address and a default gateway, such as default gateway 660. DHCP server 720 may also designate DNS server 710 as the proper source for DNS lookups.

**[0115]** In operation 1, DXL endpoint 120 queries DNS server 710 to try to resolve, for example, FABRIC.MY\_DOMAIN (or any other unique domain name assigned).

**[0116]** In operation 2, DNS server 710 responds to the query with the IP address of one or more DXL brokers 110 that may be configured to service DXL endpoint 120.

**[0117]** In operation 3, DXL endpoint 120 uses the resolved IP address to try to connect to DXL ESB 130 via DXL broker 110. If DNS server 710 provided more than one DXL broker 110, DXL endpoint 120 may try to connect to each one in turn until it successfully makes a connection. DXL endpoint 120 is now joined to the fabric of DXL ESB 130.

**[0118]** In operation 4, after successfully connecting to DXL ESB 130, DXL endpoint 120 subscribes to a message topic such as LOCATION\_SERVICES to gain access to location services provided by asset management engine 428. The topic of this message may be the private topic for the DXL endpoint 120.

**[0119]** In operation 5, DXL endpoint 120 initiates a query-response communications with asset management engine 428 over the DXL. As part of the query, DXL endpoint 120 provides its MAC address, along with the MAC address of the router acting as its default gateway.

**[0120]** In operation 6, asset management engine 428 uses the MAC address of the default gateway router to identify the location or subnetwork that DXL endpoint 120 belongs to. Asset management engine 428 then crafts a DXL response message directed specifically to be routed back to DXL endpoint 120. The body of the DXL message may include the current location of DXL endpoint 120, along with a list of services available at that location. Asset management engine 428 may also update its internal database to include location information for DXL endpoint 120.

**[0121]** Upon receiving the response message, DXL endpoint 120 may use the information to automatically configure available products and services including, but not limited to, the DXL broker it should be connected to.

**[0122]** In an alternative embodiment, DXL endpoint 120 may be configured by a management system, such as a policy orchestrator, with the IP address of a DXL broker(s) it should be connecting to starting the process.

**[0123]** FIGURE 8 is a block diagram of a method of mapping a logical location according to one or more examples of the present Specification.

**[0124]** In the example of FIGURE 7, DXL endpoint 120 learned only its location relative to its layer-3 connectivity (e.g., the router it is using as default gateway).

**[0125]** In the example of FIGURE 8, a physical network topology engine 810 may also provide layer-2 connectivity data. In this case, DXL endpoint 120 is connected to a logical network 870, including a default gateway 120 and one or more switches 680.

**[0126]** In operation 1, asset management engine 428 detects a new device connected to the network (as discussed above), and publishes a `DEVICE_DETECTED` message, including the MAC address of the DXL endpoint 120 and the MAC address of the router acting as its default gateway. Asset management engine 428 may also build a list of switches that DXL endpoint 120 may be connected to.

**[0127]** In operation 2, using this information, physical network topology engine 810 queries each switch on the path to the router acting as the default gateway for DXL endpoint 120. This yields the switch IP address and port number that DXL endpoint 120 is connected to.

**[0128]** In operation 3, physical network topology engine 810 uses information from operation 2 to build the logical location of DXL endpoint 120.

**[0129]** In operation 4, physical network topology engine 810 publishes a `DEVICE_LOGICAL_LOCATION_UPDATE` message on DXL ESB 130.

**[0130]** In operation 4, asset management engine 428 consumes this message and updates its records.

**[0131]** A block diagram example of the derived path is seen in FIGURE 9. Switches  $S_3$  and  $S_1$  are on the layer-3 path, which interface  $R_{1i}$  of router  $R_1$  serves. To find the location of DXL endpoint 120 using interface  $R_{1i}$  as its default gateway, physical network topology engine 810 queries switches  $S_3$  and  $S_1$  for their knowledge of the MAC address of DXL endpoint 120, which in this case is endpoint  $X_3$ . Using the received answers, and having the physical network topology knowledge, physical network topology engine 810 deduces that DXL endpoint 120 is connected to switch  $S_1$  via port  $R_{1i}$ . A method for accomplishing this is described in U.S. Patent 8,811,224, which is incorporated herein by reference.

**[0132]** FIGURE 10 is a block diagram of a method of determining both logical location and physical location according to one or more examples of the present Specification.

**[0133]** In operation 1, DXL endpoint 120 performs a DNS lookup query to `FABRIC.MY_DOMAIN`, directed to DNS server 710.

**[0134]** In operation 2, DNS server 710 responds to the DNS query with the IP address of DXL broker 110, which is to operate as the DXL broker DXL endpoint 120 is to initially connect to.

**[0135]** In operation 3, DXL endpoint 120 connects to DXL ESB 130 via DXL broker 110.

**[0136]** In operation 4, DXL endpoint 120 subscribes to LOCATION\_SERVICES messages.

**[0137]** In operation 5, DXL endpoint 120 queries asset management engine 428 and requests location and services information.

**[0138]** In operation 6, asset management engine 428 uses the attributes contained within the request of operation 5 to return location and services information to DXL endpoint 120.

**[0139]** In operation 7, DXL endpoint 120 uses the response to automatically configure available products.

**[0140]** In operation 8, physical network topology engine 810 of asset management engine 428 performs a device logical location update after determining, in operation 8(A), the logical location of DXL endpoint 120. This determination is according to the disclosure of FIGURE 8.

**[0141]** FIGURE 11 is a block diagram of a method of maintaining location information according to one or more examples of the present Specification. In the example of FIGURE 11, a runtime agent plugin 1110 is provided.

**[0142]** In operation 1, runtime agent plugin 1110 publishes a DEVICE\_DETECTED event when a new device is attached to DXL ESB 130 or to a subnet that runtime agent plugin 1110 is operating on. The message may include the MAC address and the IP address of the detected device, along with the MAC address of the router servicing this IP subnet.

**[0143]** This message is to be consumed by asset management engine 428. Thus, in operation 2, asset management engine 428 uses the MAC address of the gateway router and the IP subnet address to pinpoint the location on which the new device is operating and to create an association between the location and the device.

**[0144]** In operation 3, physical network topology engine 810 of asset management engine 428 determines the logical location of the new device according to method disclosed in FIGURE 8.

**[0145]** In operation 4, asset management engine 428 updates its internal database with the logical location of the new device.

**[0146]** The foregoing outlines features of several embodiments so that those skilled in the art may better understand the aspects of the present disclosure. Those skilled in the art should appreciate that they may readily use the present disclosure as a basis for designing or modifying other processes and structures for carrying out the same purposes and/or achieving the same advantages of the embodiments introduced herein. Those skilled in the art should also realize that such equivalent constructions do not depart from the spirit and scope of the present disclosure, and that they may make various changes, substitutions, and alterations herein without departing from the spirit and scope of the present disclosure.

**[0147]** The particular embodiments of the present disclosure may readily include a system on chip (SOC) central processing unit (CPU) package. An SOC represents an integrated circuit (IC) that integrates components of a computer or other electronic system into a single chip. It may contain digital, analog, mixed-signal, and radio frequency functions: all of which may be provided on a single chip substrate. Other embodiments may include a multi-chip-module (MCM), with a plurality of chips located within a single electronic package and configured to interact closely with each other through the electronic package. In various other embodiments, the digital signal processing functionalities may be implemented in one or more silicon cores in Application Specific Integrated Circuits (ASICs), Field Programmable Gate Arrays (FPGAs), and other semiconductor chips.

**[0148]** In example implementations, at least some portions of the processing activities outlined herein may also be implemented in software. In some embodiments, one or more of these features may be implemented in hardware provided external to the elements of the disclosed figures, or consolidated in any appropriate manner to achieve the intended functionality. The various components may include software (or reciprocating software) that can coordinate in order to achieve the operations as outlined herein. In still other embodiments, these elements may include any suitable algorithms, hardware, software, components, modules, interfaces, or objects that facilitate the operations thereof.

**[0149]** Additionally, some of the components associated with described microprocessors may be removed, or otherwise consolidated. In a general sense, the arrangements depicted in the figures may be more logical in their representations, whereas a

physical architecture may include various permutations, combinations, and/or hybrids of these elements. It is imperative to note that countless possible design configurations can be used to achieve the operational objectives outlined herein. Accordingly, the associated infrastructure has a myriad of substitute arrangements, design choices, device possibilities, hardware configurations, software implementations, equipment options, etc.

**[0150]** Any suitably configured processor component can execute any type of instructions associated with the data to achieve the operations detailed herein. Any processor disclosed herein could transform an element or an article (for example, data) from one state or thing to another state or thing. In another example, some activities outlined herein may be implemented with fixed logic or programmable logic (for example, software and/or computer instructions executed by a processor) and the elements identified herein could be some type of a programmable processor, programmable digital logic (for example, a field programmable gate array (FPGA), an erasable programmable read only memory (EPROM), an electrically erasable programmable read only memory (EEPROM)), an ASIC that includes digital logic, software, code, electronic instructions, flash memory, optical disks, CD-ROMs, DVD ROMs, magnetic or optical cards, other types of machine-readable mediums suitable for storing electronic instructions, or any suitable combination thereof. In operation, processors may store information in any suitable type of non-transitory storage medium (for example, random access memory (RAM), read only memory (ROM), field programmable gate array (FPGA), erasable programmable read only memory (EPROM), electrically erasable programmable ROM (EEPROM), etc.), software, hardware, or in any other suitable component, device, element, or object where appropriate and based on particular needs. Further, the information being tracked, sent, received, or stored in a processor could be provided in any database, register, table, cache, queue, control list, or storage structure, based on particular needs and implementations, all of which could be referenced in any suitable timeframe. Any of the memory items discussed herein should be construed as being encompassed within the broad term 'memory.' Similarly, any of the potential processing elements, modules, and machines described herein should be construed as being encompassed within the broad term 'microprocessor' or 'processor.' Furthermore, in various embodiments, the processors, memories, network cards, buses, storage devices, related peripherals, and other hardware elements described herein may be realized by a processor,

memory, and other related devices configured by software or firmware to emulate or virtualize the functions of those hardware elements.

**[0151]** Computer program logic implementing all or part of the functionality described herein is embodied in various forms, including, but in no way limited to, a source code form, a computer executable form, and various intermediate forms (for example, forms generated by an assembler, compiler, linker, or locator). In an example, source code includes a series of computer program instructions implemented in various programming languages, such as an object code, an assembly language, or a high-level language such as OpenCL, Fortran, C, C++, JAVA, or HTML for use with various operating systems or operating environments. The source code may define and use various data structures and communication messages. The source code may be in a computer executable form (e.g., via an interpreter), or the source code may be converted (e.g., via a translator, assembler, or compiler) into a computer executable form.

**[0152]** In the discussions of the embodiments above, the capacitors, buffers, graphics elements, interconnect boards, clocks, DDRs, camera sensors, dividers, inductors, resistors, amplifiers, switches, digital core, transistors, and/or other components can readily be replaced, substituted, or otherwise modified in order to accommodate particular circuitry needs. Moreover, it should be noted that the use of complementary electronic devices, hardware, non-transitory software, etc. offers an equally viable option for implementing the teachings of the present disclosure.

**[0153]** In one example embodiment, any number of electrical circuits of the FIGURES may be implemented on a board of an associated electronic device. The board can be a general circuit board that can hold various components of the internal electronic system of the electronic device and, further, provide connectors for other peripherals. More specifically, the board can provide the electrical connections by which the other components of the system can communicate electrically. Any suitable processors (inclusive of digital signal processors, microprocessors, supporting chipsets, etc.), memory elements, etc. can be suitably coupled to the board based on particular configuration needs, processing demands, computer designs, etc. Other components such as external storage, additional sensors, controllers for audio/video display, and peripheral devices may be attached to the board as plug-in cards, via cables, or integrated into the board itself. In another example embodiment, the electrical circuits of the FIGURES may be implemented as stand-alone modules (e.g., a

device with associated components and circuitry configured to perform a specific application or function) or implemented as plug-in modules into application specific hardware of electronic devices.

**[0154]** Note that with the numerous examples provided herein, interaction may be described in terms of two, three, four, or more electrical components. However, this has been done for purposes of clarity and example only. It should be appreciated that the system can be consolidated in any suitable manner. Along similar design alternatives, any of the illustrated components, modules, and elements of the FIGURES may be combined in various possible configurations, all of which are clearly within the broad scope of this Specification. In certain cases, it may be easier to describe one or more of the functionalities of a given set of flows by only referencing a limited number of electrical elements. It should be appreciated that the electrical circuits of the FIGURES and its teachings are readily scalable and can accommodate a large number of components, as well as more complicated/sophisticated arrangements and configurations. Accordingly, the examples provided should not limit the scope or inhibit the broad teachings of the electrical circuits as potentially applied to a myriad of other architectures.

**[0155]** Numerous other changes, substitutions, variations, alterations, and modifications may be ascertained to one skilled in the art and it is intended that the present disclosure encompass all such changes, substitutions, variations, alterations, and modifications as falling within the scope of the appended claims. In order to assist the United States Patent and Trademark Office (USPTO) and, additionally, any readers of any patent issued on this application in interpreting the claims appended hereto, Applicant wishes to note that the Applicant: (a) does not intend any of the appended claims to invoke paragraph six (6) of 35 U.S.C. section 112 as it exists on the date of the filing hereof unless the words “means for” or “steps for” are specifically used in the particular claims; and (b) does not intend, by any statement in the specification, to limit this disclosure in any way that is not otherwise reflected in the appended claims.

#### EXAMPLE IMPLEMENTATIONS

**[0156]** There is disclosed in an example 1, a domain services controller apparatus for a data exchange layer (DXL), comprising: a network interface operable for communicatively

coupling the apparatus to DXL enterprise service bus (ESB); one or more logic elements comprising a DXL services engine operable to provide an application programming interface for communicating with DXL endpoints via the DXL ESB; and one or more logic elements comprising an asset management engine operable for: subscribing to a DXL location services topic; receiving a DXL location services query from a DXL endpoint via the DXL ESB; and providing network location data to the DXL endpoint via the DXL ESB.

**[0157]** There is disclosed in an example 2, the apparatus of example 1, wherein providing network location data comprises providing logical network location data.

**[0158]** There is disclosed in an example 3, the apparatus of example 1, wherein providing network location data comprises providing physical network location data.

**[0159]** There is disclosed in an example 4, the apparatus of example 1, wherein providing network location data comprises: receiving information about the DXL endpoint via a DXL query; determining a logical location of the DXL client based on the information about the DXL client; and publishing a DXL response comprising the logical location.

**[0160]** There is disclosed in an example 5, the apparatus of example 4, wherein a topic of the DXL response is a private destination topic of the DXL client.

**[0161]** There is disclosed in an example 6, the apparatus of example 1, wherein the asset management engine further comprises a physical network topology engine operable for mapping physical network locations.

**[0162]** There is disclosed in an example 7, the apparatus of example 1, wherein the asset management engine is further operable to interoperate with a physical network topology engine operable for mapping physical network locations.

**[0163]** There is disclosed in an example 8, the apparatus of any of examples 1 – 7, wherein the asset management engine is further operable for: subscribing to a DXL device detected topic; receiving a DXL device detected message comprising an identifier for a detected device and an identifier for a broker for the detected device; building a list of switches to the detected device; and building a logical location of the detected device.

**[0164]** There is disclosed in an example 9, the apparatus of example 8, wherein the asset management engine is further operable for publishing the logical location of the detected device as a detected device DXL message.

**[0165]** There is disclosed in an example 10, the apparatus of any of examples 1 – 7, wherein the asset management engine is operable for providing both logical and physical location services.

**[0166]** There is disclosed in an example 11, the apparatus of any of examples 1 – 7, wherein the asset management engine is operable for receiving manual user input for updating the location services.

**[0167]** There is disclosed in an example 12, the apparatus of any of examples 1 – 7, wherein the apparatus is further configured to act as a producer in a DXL producer-consumer architecture.

**[0168]** There is disclosed in an example 13, the apparatus of any of examples 1 – 7, wherein the apparatus is further configured to act as a producer of DXL location services messages.

**[0169]** There is disclosed in an example 14, one or more computer-readable mediums having stored thereon instructions for instructing a processor for providing a data exchange layer (DXL) asset management engine operable for: subscribing to a DXL location services topic; receiving a DXL location services query from a DXL endpoint via a DXL enterprise service bus (ESB); and providing network location data to the DXL endpoint via the DXL ESB.

**[0170]** There is disclosed in an example 15, the one or more computer-readable mediums of example 14, wherein providing network location data comprises providing logical network location data.

**[0171]** There is disclosed in an example 16, the one or more computer-readable mediums of example 14, wherein providing network location data comprises providing physical network location data.

**[0172]** There is disclosed in an example 17, the one or more computer-readable mediums of example 14, wherein providing network location data comprises: receiving information about the DXL endpoint via a DXL query; determining a logical location of the DXL client based on the information about the DXL client; and publishing a DXL response comprising the logical location.

**[0173]** There is disclosed in an example 18, the one or more computer-readable mediums of example 17, wherein a topic of the DXL response is a private destination topic of the DXL client.

**[0174]** There is disclosed in an example 19, the one or more computer-readable mediums of example 14, wherein the asset management engine further comprises a physical network topology engine operable for mapping physical network locations.

**[0175]** There is disclosed in an example 20, the one or more computer-readable mediums of example 14, wherein the asset management engine is further operable to interoperate with a physical network topology engine operable for mapping physical network locations.

**[0176]** There is disclosed in an example 21, the one or more computer-readable mediums of any of examples 14 – 20, wherein the asset management engine is further operable for: subscribing to a DXL device detected topic; receiving a DXL device detected message comprising an identifier for a detected device and an identifier for a broker for the detected device; building a list of switches to the detected device; and building a logical location of the detected device.

**[0177]** There is disclosed in an example 22, the one or more computer-readable mediums of example 21, wherein the asset management engine is further operable for publishing the logical location of the detected device as a detected device DXL message.

**[0178]** There is disclosed in an example 23, the one or more computer-readable mediums of any of examples 14 – 20, wherein the asset management engine is operable for receiving manual user input for updating the location services.

**[0179]** There is disclosed in an example 24, a computer-implemented method of location services on a data exchange layer (DXL) comprising: subscribing to a DXL location services topic; receiving a DXL location services query from a DXL endpoint via a DXL enterprise service bus (ESB); and providing network location data to the DXL endpoint via the DXL ESB.

**[0180]** There is disclosed in an example 25, the computer-implemented method of example 24, further comprising: subscribing to a DXL device detected topic; receiving a DXL device detected message comprising an identifier for a detected device and an identifier for a broker for the detected device; building a list of switches to the detected device; and building a logical location of the detected device.

**[0181]** There is disclosed in an example 26, a method comprising the performing the instructions disclosed in any of examples 14 – 23.

**[0182]** There is disclosed in example 27, an apparatus comprising means for performing the method of example 26.

**[0183]** There is disclosed in example 28, the apparatus of claim 27, wherein the apparatus comprises a processor and memory.

**[0184]** There is disclosed in example 29, the apparatus of claim 28, wherein the apparatus further comprises a computer-readable medium having stored thereon software instructions for performing the method of example 26.

CLAIMS

WHAT IS CLAIMED IS:

1. A domain services controller apparatus for a data exchange layer (DXL), comprising:  
  
a network interface operable for communicatively coupling the apparatus to DXL enterprise service bus (ESB);  
  
one or more logic elements comprising a DXL services engine operable to provide an application programming interface for communicating with DXL endpoints via the DXL ESB;  
  
and  
  
one or more logic elements comprising an asset management engine operable for:  
  
subscribing to a DXL location services topic;  
  
receiving a DXL location services query from a DXL endpoint via the DXL ESB; and  
  
providing network location data to the DXL endpoint via the DXL ESB.
2. The apparatus of claim 1, wherein providing network location data comprises providing logical network location data.
3. The apparatus of claim 1, wherein providing network location data comprises providing physical network location data.
4. The apparatus of claim 1, wherein providing network location data comprises:  
  
receiving information about the DXL endpoint via a DXL query;  
  
determining a logical location of the DXL client based on the information about the DXL client; and  
  
publishing a DXL response comprising the logical location.
5. The apparatus of claim 4, wherein a topic of the DXL response is a private destination topic of the DXL client.

6. The apparatus of claim 1, wherein the asset management engine further comprises a physical network topology engine operable for mapping physical network locations.

7. The apparatus of claim 1, wherein the asset management engine is further operable to interoperate with a physical network topology engine operable for mapping physical network locations.

8. The apparatus of any of claims 1 – 7, wherein the asset management engine is further operable for:

subscribing to a DXL device detected topic;

receiving a DXL device detected message comprising an identifier for a detected device and an identifier for a broker for the detected device;

building a list of switches to the detected device; and

building a logical location of the detected device.

9. The apparatus of claim 8, wherein the asset management engine is further operable for publishing the logical location of the detected device as a detected device DXL message.

10. The apparatus of any of claims 1 – 7, wherein the asset management engine is operable for providing both logical and physical location services.

11. The apparatus of any of claims 1 – 7, wherein the asset management engine is operable for receiving manual user input for updating the location services.

12. The apparatus of any of claims 1 – 7, wherein the apparatus is further configured to act as a producer in a DXL producer-consumer architecture.

13. The apparatus of any of claims 1 – 7, wherein the apparatus is further configured to act as a producer of DXL location services messages.

14. One or more computer-readable mediums having stored thereon instructions for instructing a processor for providing a data exchange layer (DXL) asset management engine operable for:

subscribing to a DXL location services topic;

receiving a DXL location services query from a DXL endpoint via a DXL enterprise service bus (ESB); and

providing network location data to the DXL endpoint via the DXL ESB.

15. The one or more computer-readable mediums of claim 14, wherein providing network location data comprises providing logical network location data.

16. The one or more computer-readable mediums of claim 14, wherein providing network location data comprises providing physical network location data.

17. The one or more computer-readable mediums of claim 14, wherein providing network location data comprises:

receiving information about the DXL endpoint via a DXL query;

determining a logical location of the DXL client based on the information about the DXL client; and

publishing a DXL response comprising the logical location.

18. The one or more computer-readable mediums of claim 17, wherein a topic of the DXL response is a private destination topic of the DXL client.

19. The one or more computer-readable mediums of claim 14, wherein the asset management engine further comprises a physical network topology engine operable for mapping physical network locations.

20. The one or more computer-readable mediums of claim 14, wherein the asset management engine is further operable to interoperate with a physical network topology engine operable for mapping physical network locations.

21. The one or more computer-readable mediums of any of claims 14 – 20, wherein the asset management engine is further operable for:

- subscribing to a DXL device detected topic;

- receiving a DXL device detected message comprising an identifier for a detected device and an identifier for a broker for the detected device;

- building a list of switches to the detected device; and

- building a logical location of the detected device.

22. The one or more computer-readable mediums of claim 21, wherein the asset management engine is further operable for publishing the logical location of the detected device as a detected device DXL message.

23. The one or more computer-readable mediums of any of claims 14 – 20, wherein the asset management engine is operable for receiving manual user input for updating the location services.

24. A computer-implemented method of location services on a data exchange layer (DXL) comprising:

- subscribing to a DXL location services topic;

- receiving a DXL location services query from a DXL endpoint via a DXL enterprise service bus (ESB); and

- providing network location data to the DXL endpoint via the DXL ESB.

25. The computer-implemented method of claim 24, further comprising:

subscribing to a DXL device detected topic;

receiving a DXL device detected message comprising an identifier for a detected device and an identifier for a broker for the detected device;

building a list of switches to the detected device; and

building a logical location of the detected device.

1/10




FIG. 1






FIG. 2

3/10






FIG. 5




FIG. 6




FIG. 7

7/10




FIG. 8




FIG. 9




FIG. 10

10/10




FIG. 11

## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2014/057935****A. CLASSIFICATION OF SUBJECT MATTER****H04L 29/02(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**Minimum documentation searched (classification system followed by classification symbols)  
H04L 29/02; G06F 15/173; G06F 17/30; H04L 12/28; G06F 21/00Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
Korean utility models and applications for utility models  
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
eKOMPASS(KIPO internal) & Keywords: ESB(Enterprise Service Bus), service, interface, endpoint, network, DXL, location**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                                                                                                                                                                                               | Relevant to claim No. |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| A         | US 2010-0088326 A1 (DAVID D. CHEN et al.) 08 April 2010<br>See paragraphs [0030], [0037]; claim 7; and figure 2.                                                                                                                                                                                                                                                                                                                                 | 1-25                  |
| A         | US 7814226 B2 (PAUL B. PATRICK) 12 October 2010<br>See column 7, lines 1-20; and figure 4.                                                                                                                                                                                                                                                                                                                                                       | 1-25                  |
| A         | US 2011-0088099 A1 (KATHY VILLANI et al.) 14 April 2011<br>See paragraphs [0004]-[0005], [0066]-[0067]; and figure 4.                                                                                                                                                                                                                                                                                                                            | 1-25                  |
| A         | ALAEDDIN KALANTARI et al. `Enabling Security Requirements for Enterprise Service-Oriented Architecture` ACEEE : Int. J. on Recent Trends in Engineering and Technology, Vol. 6, No. 1, 31 November 2011<br>( <a href="http://searchdl.org/public/journals/2011/IJRTET/6/1/243.pdf">http://searchdl.org/public/journals/2011/IJRTET/6/1/243.pdf</a> )<br>See page 76, right column, lines 21-42; page 79, left column, lines 12-42; and figure 2. | 1-25                  |
| A         | HEATHER HINTON, MARYANN HONDO, DR. BETH HUTCHISON. `Security Patterns within a Service-Oriented Architecture` 31 November 2005<br>( <a href="http://cdn.ttgtmedia.com/searchSOA/downloads/SecuritySOA_(2).pdf">http://cdn.ttgtmedia.com/searchSOA/downloads/SecuritySOA_(2).pdf</a> )<br>See page 10, lines 13-41; and figure 4.                                                                                                                 | 1-25                  |

☐ Further documents are listed in the continuation of Box C.☒ See patent family annex.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&amp;" document member of the same patent family

Date of the actual completion of the international search

08 January 2015 (08.01.2015)

Date of mailing of the international search report

**12 January 2015 (12.01.2015)**

Name and mailing address of the ISA/KR

International Application Division  
Korean Intellectual Property Office  
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,  
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

KIM, Seong Woo

Telephone No. +82-42-481-3348



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2014/057935**

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date |
|-------------------------------------------|---------------------|----------------------------|---------------------|
| US 2010-0088326 A1                        | 08/04/2010          | US 8140582 B2              | 20/03/2012          |
| US 7814226 B2                             | 12/10/2010          | US 2008-0069124 A1         | 20/03/2008          |
|                                           |                     | WO 2008-036777 A2          | 27/03/2008          |
|                                           |                     | WO 2008-036777 A3          | 04/09/2008          |
| US 2011-0088099 A1                        | 14/04/2011          | US 8555407 B2              | 08/10/2013          |