



- (51) **International Patent Classification:**
H04L 29/06 (2006.01) H04L 9/08 (2006.01)
H04L 29/08 (2006.01)
- (21) **International Application Number:**
PCT/US2016/031565
- (22) **International Filing Date:**
10 May 2016 (10.05.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/1173,857 10 June 2015 (10.06.2015) US
14/866,203 25 September 2015 (25.09.2015) US
- (71) **Applicant: MCAFEE, INC.** [US/US]; 2821 Mission College Boulevard, Santa Clara, California 95054-1838 (US).
- (72) **Inventors: SMITH, Ned M.;** 375 SW Delta Drive, Beaverton, Oregon 97006 (US). **HUNT, Simon;** 2120 Sheepshead Drive, Naples, Florida 34102 (US). **SAMBANDAM, Venkata Ramanan;** 180 Elm Ct., Apt. 912, Sunnyvale, California 94086 (US).
- (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).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

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

(54) **Title:** SENTINEL APPLIANCE IN AN INTERNET OF THINGS REALM

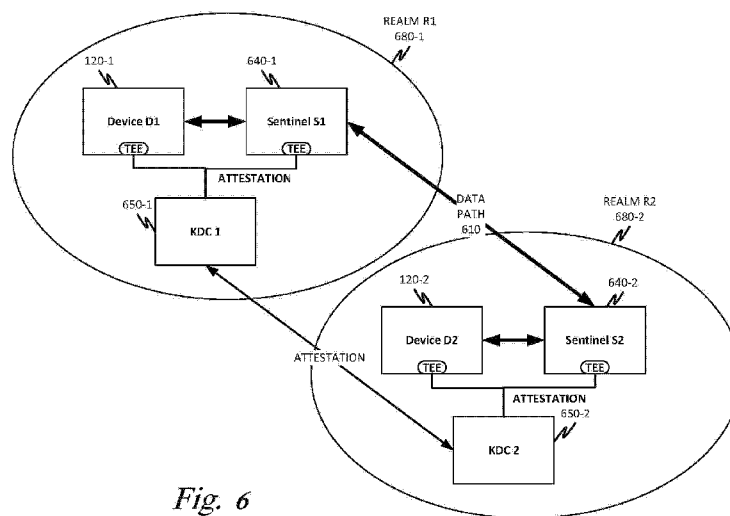


Fig. 6

(57) **Abstract:** In an example, there is disclosed a computing apparatus, comprising : a trusted execution environment (TEE); and a security engine operable to: identify a key negotiation for an encrypted connection between a first device and a second device; request a service appliance key for the key negotiation; receive the service appliance key; and perform a service appliance function on traffic between the first device and the second device. There is also disclosed a method of providing the security engine, and a computer-readable medium having stored thereon executable instructions for providing the security engine.

WO 2016/200526 A1

SENTINEL APPLIANCE IN AN INTERNET OF THINGS REALM

Cross Reference to Related Applications

[0001] This Application claims the benefit of and priority to U.S. Provisional Application No. 62/173,857, titled "Internet of Things Device Security Inspection," filed June 10, 2015 and to U.S. Nonprovisional Utility Application No. 14/866,203, titled "Sentinel Appliance in an Internet of Things Realm", filed September 25, 2015, both of which are incorporated herein by reference.

Field of the specification

[0002] This disclosure relates in general to the field of computer security, and more particularly, though not exclusively to, a system and method for providing a sentinel appliance in an internet-of-things realm.

Background

[0003] The Internet of Things (IoT) is a loosely-defined network of physical objects (things) with embedded computing and communication capabilities, thus allowing the "things" to exchange data with one another. Within the IoT, real world phenomena can be sensed or observed electronically, and outputs from sensors or other data sources may be used as an input to a control system. In some cases, this allows a tighter coupling between the physical world and the virtual space. Each "thing" in the IoT may be uniquely identified with its physical computing platform, and some things are configured to operate within the existing internet infrastructure. Other devices operate over other network topologies, including *ad hoc* networks and direct data connections among others. In general terms, the IoT is considered to be a highly-democratic (sometimes even anarchic) network in which individual devices and networks may have broad autonomy in terms of what they do, how they do it, and how they communicate about it.

Brief Description of the Drawings

[0004] 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 necessarily drawn to scale, and are used for illustration purposes only. Where a scale is shown, explicitly or implicitly, it provides only one illustrative example. In other embodiments, the dimensions of the various features may be arbitrarily increased or reduced for clarity of discussion.

[0005] FIGURE 1 is a block diagram of a monitored network, which may be or comprise an IoT realm or domain, according to one or more examples of the present specification.

[0006] FIGURE 2 is a block diagram of a data source according to one or more examples of the present specification.

[0007] FIGURE 3 is a block diagram of a data aggregator according to one or more examples of the present specification.

[0008] FIGURE 4 is a block diagram of an expert system according to one or more examples of the present specification.

[0009] FIGURE 5 is a block diagram of a trusted execution environment according to one or more examples of the present specification.

[0010] FIGURE 6 is a block diagram of an internet of things according to one or more examples of the present specification.

[0011] FIGURE 7 is a block diagram of a KDC according to one or more examples of the present specification.

[0012] FIGURE 8 is a block diagram of a sentinel device according to one or more examples of the present specification.

[0013] FIGURE 9 is a flow chart of a method according to one or more examples of the present specification.

Summary

[0014] In an example, there is disclosed a computing apparatus, comprising: a trusted execution environment (TEE); and a security engine operable to: identify a key negotiation for an encrypted connection between a first device and a second device; request a service appliance key for the key negotiation; receive the service appliance key; and perform a service appliance function on traffic between the first device and the second device. There is also disclosed a method of providing the security engine, and a computer-readable medium having stored thereon executable instructions for providing the security engine.

Embodiments of the Disclosure

[0015] 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. Different embodiments may have different advantages, and no particular advantage is necessarily required of any embodiment.

[0016] The "Internet of Things" (IoT) is an explosive global aggregation of heterogeneous "smart" and "network-enabled" devices that often each provide a single, specialized function. The IoT includes smart appliances, smart sensors, smart phones, and a plethora of other devices to which "smart" can be (and often is) prepended.

[0017] Because the IoT is not a traditional network, it presents challenges that are sometimes new and unique. For example, in traditional networking, an oligarchy of naming authorities parcel out a relative handful of globally-unique internet protocol (IP) addresses, as in the IPv4 space, which has a theoretical maximum of approximately 4.3 billion unique addresses. This made global IPv4

addresses a relatively dear commodity to be managed centrally by an administrative body. But the cascade of IoT devices makes such a scheme both impractical and undesirable for many. Rather, IoT devices may operate within defined subnetworks using network address translation (NAT), or may self-declare a "Universally Unique Identifier" (UUID), which in one example is a 128-bit integer, and which may be usable as an IPv6 IP address. Such autonomous naming presents both new opportunities and new challenges that users and enterprises are still working to understand and appreciate. In one sense, the IoT may be viewed as a new, wild frontier, where rules are still evolving and where any device can be practically anything that it wants to be. While this provides exciting opportunities for innovators to experiment and try new things, it also provides a sometimes-lawless frontier where devices and their designers may not always be able to rely on traditional security solutions.

[0018] One consequence of the free-wheeling nature of the IoT is that security is essentially up to the system designer of a thing, and as enterprises work to push out faster, cheaper things, sometimes at the expense of better, particularly with respect to security. Recent attention has been brought to attacks such as "Superfish," in which a "man-in-the-middle" (MITM) issued fake Google-branded Secure Socket Layer (SSL) certificates.

[0019] Another difficulty is that IoT devices may be constrained in terms of processing power, memory, storage and power. So they may not be capable of providing their own detailed security, such as deep packet inspection, reputation, network access control (NAC), and security information and event management (SIEM), data loss prevention, antivirus scanning, and anomaly detection for example, during real-time network operations.

[0020] Some existing and emerging network topologies are being constructed using more advanced security, such as datagram transport layer security (DTLS) and transport layer security (TLS) to realize end-to-end secure sessions that are independent of the network transport(s). So in some cases, IoT device traffic cannot be inspected by more powerful upstream monitoring appliances, such as gateways, firewalls, or other security appliances.

[0021] In some IoT network topologies, one or more trustworthy gateways or intermediary devices exist that could be employed as network security monitoring agents. But in some cases, this cannot be accomplished without an infrastructure for establishing shared keys for a trusted intermediary.

[0022] Embodiments of this specification provide methods for appropriately-equipped IoT devices to negotiate an upstream device to decrypt and inspect its network traffic, while still maintaining the overall security and confidentiality of the network connection to the device.

[0023] In an example, four actors are identified to illustrate the IoT network topology:

[0024] A first endpoint device (D1), which may be an IoT "thing" such as a data source.

[0025] A second (remote) endpoint device (D2).

[0026] An intermediary "sentinel" device (S1), which may provide offloaded security or other functions.

[0027] A key management device, such as a key distribution center (KDC1). Optionally, a second KDC may also be used when crossing domain boundaries where D1 and D2 are in separate administrative domains (a.k.a., IoT "realms"). The second KDC (KDC2) is a key management authority for a second realm (R2) while KDC1 is an authority for a first realm (R1).

[0028] In one example, a novel group key management model provides an end-to-end symmetric key exchange protocol, such as the Kerberos system, which may be used to establish a shared key between devices D1 and D2, but where an intermediary (S1) also participates. In an embodiment, an infrastructure such as draft-hardjono-ace-fluffy-OO may be used to generate "mini-tickets," and a request-response protocol with KDC1, wherein S1 is also permitted to receive the D1-D2 shared key.

[0029] For security implementation, certain aspects may be implemented within a trusted execution environment (TEE) to establish an attested and trusted connection to KDC1 such that the KDC can determine whether or not the configuration and capabilities of the endpoint device or sentinel device satisfies a

domain security policy for trustworthy IoT operation. In the case of sentinel device S1 that is specifically responsible for performing network security scanning, deep packet inspection, reputation processing and SIEM, a policy may also require that it have a TEE at least as secure as the TEE on the endpoint device.

[0030] Once trust is established, the KDC can share symmetric keys with secured communication channels such as SSL, TLS, DTLS, or a message encryption technology such as JSON Web Encryption (JWE). Embodiments include sharing a symmetric key (K1) with S1 in addition to the respective endpoint devices D1 and D2.

[0031] In an example, the KDC secures the key (K1) in a TEE. S1 may also use a TEE to protect K1. The TEE may incorporate at least two network connections, (1) an "inside" connection between the S1 TEE and the inside device D1, and (2), a network connection within the TEE. The "outside" network connects the remote device D2 to network monitoring, packet processing and other security scanning, SIEM, and Network Access Control (NAC) functions, where such operations are applied from within a TEE environment. Thus, the risk of leaking the encryption keys (K1) or a compromise from malware attacks on the S1 device is minimized. This includes protection from specific attack scenarios such as SSL bump and unauthorized MITM.

[0032] A further embodiment extends the TEE into the IoT device so that the endpoints of communication are also hardened against attacks on the endpoint devices. For example, KDC1 established that endpoint devices D1 and D2 both possess a valid TEE upon processing a ticket request message. The TEE may then be used to establish a secure session context with which to exchange subsequent key management messages. The secure context establishes the trusted computing properties inherent in the endpoint devices D1 and D2.

[0033] In an embodiment, a second KDC (KDC2) operates within a second realm R2 where the endpoint device D2 is managed by a different domain authority than R1. In this situation, the domain authorities R1 and R2 may negotiate a trust policy such that keys issued through R2 to D2 are of an equivalent trust level to those issued by R1 to D1 and S1. In this scenario, R1 and R2 establish a Diffie-Hellman key agreement protocol that may attest or be

attested by using a TEE-TEE attestation protocol such as the Sign-and-Mac (Sigma) protocol defined by Intel and by Internet Engineering Task Force (IETF) drafts. In particular, an asymmetric key that prevents MITM attackers from posing as a fake R2 may sign the sigma protocol messages. The private key is used to generate the signature and the public key is used to verify the signature. A security monitoring reputation, NAC and SIEM policy may be exchanged between R 1 and R2 to establish a security context in which the issuance of session keys may be common and mutually agreed upon.

[0034] A further embodiment allows for a second realm R2 to establish a security monitoring intermediary S2 whose aim is to apply a security policy aimed at protecting realm R2 according to a policy established by a trusted key management and security policy management device KDC2.

[0035] In yet another embodiment, the KDC function is combined with an endpoint device D1, resources and computation power permitting. In that case, device D1 may implement the functions described for the KDC.

[0036] Advantageously, this specification provides for a trusted key management and security policy management device (KDC1) to establish a shared symmetric key across multiple devices participating in a secure communication that is monitored by one or more trusted intermediaries that are vetted by KDC1 according to a trusted computing policy and an attestation exchange.

[0037] Further advantageously, embodiments establish a common realm-realm security policy and key management protocol whereby different realms R1 and R2 may dynamically agree upon a policy and protocol that allows a device D2 in R2 to interact with a device D1 in R1 following an agreed upon policy.

[0038] Further advantageously, embodiments establish an IoT intermediary or proxy service that may apply any suitable security function, such as monitoring, reputation, NAC, or SIEM by way of non-limiting example. This may otherwise be difficult in certain IoT architectures, as power constraints and rapid time-to-market requirements may otherwise preclude such exotic functions.

[0039] Embodiments also establish a system where endpoint devices (D1, D2) and an upstream device (S1) may apply security monitoring that is “accelerated” from the perspective of devices D1 and D2.

[0040] Embodiments also establish a mutually cooperative security and privacy policy that is common among a single realm R1 or across multiple realms R2, R3 ... Rn, using a KDC representative from each realm.

[0041] There is also provided a system of KDCs that establish provably secure operations using attestation by participant devices (e.g., D1, D2, S1) involving both hardware and software isolated execution and data protection using a TEE, such as Intel SGX and other techniques.

[0042] Advantageously, after a symmetric key exchange between D1, D2 and S1, the system is prepared to apply security scanning and operational integrity functions by S1 to achieve the intended security benefits while removing the threat of a MITM attack by a rogue party masquerading as an S1 device.

[0043] A system and method for providing a sentinel appliance in an IoT realm will now be described with more particular reference to the attached FIGURES. It should be noted that throughout the FIGURES, certain reference numerals may be repeated to indicate that a particular device or block is wholly or substantially consistent across the FIGURES. This is not, however, intended to imply any particular relationship between the various embodiments disclosed. In certain examples, a genus of elements may be referred to by a particular reference numeral (“widget 10”), while individual species or examples of the genus may be referred to by a hyphenated numeral (“first specific widget 10-1” and “second specific widget 10-2”).

[0044] FIGURE 1 is a network level diagram of a monitored network 100 according to one or more examples of the present specification. Monitored network 100 illustrates an example application in which the IoT provides inputs from a plurality of data sources, which are aggregated by a data aggregator. An expert system may then make decisions to drive a controlled system.

[0045] In this example, monitored network 100 includes a plurality of data sources 120 connected to a network 170. Also connected to network 170 is a data

aggregator 110, communicatively coupled to an expert system 130, controlling controlled systems 140. Controlled systems 140 provide feedback 142 to data aggregator 110.

[0046] In one or more examples, data sources 120-1 through 120-N are disclosed. This is to illustrate that the number of data sources 120 may be large and indefinite, and may be in constant fluctuation as new data sources 120 are added to and removed from monitored network 100. Management of data sources 120 may be complicated both by the large number of data sources 120, and by the dynamic nature of monitored network 100. Thus in certain embodiments, it may be impractical for a human administrator to monitor and administer all of the various data sources 120. Furthermore, data sources 120 may not be statically located on network 170. For example, many cars carry data collection devices, and may provide data to network 170 as they hop from node to node on a mobile network. Thus, it may not be practical to predict in advance from which direction data will be coming, or what the nature of the data may be.

[0047] By way of further complication, a plurality of data sources 120 may provide data features of similar or identical types, but in slightly different formats. In one example, each data source 120 is configured to provide a data stream accompanied by a metadata packet identifying the type and source of data. However, there may be no globally enforced or enforceable standard for such metadata packets. In one example, data sources 120 may at least standardize on a delivery format for the metadata, such as XML or a similar standards-compliant data format. In that case, the metadata may have a number of identifiable field names, from which the feature type and source may be inferred. It should also be noted that in some cases data sources 120 may provide features of a compatible type, but in different formats. An example of this is a temperature feature provided by one data source 120-1 in Fahrenheit, and a second feature provided by a second data source 120-2 in Celsius. Similar issues may be encountered in any case where two or more data sources deliver similar features, with one data source providing the feature in metric units and the other data source providing the feature in Imperial or U.S. Customary units.

[0048] This large collection of features of disparate types from different sources, and in different formats, is delivered via network 170 to a data aggregator 110. Data aggregator 110 collects the many features, and in one example attempts to classify the features according to a useful taxonomy. In one case, data aggregator 110 defines a taxonomy having an arbitrary number of classification levels, such as classes, sub-classes, genera, and species. For example, the class of environmental data may include the subclass of temperature, which may include further species of temperatures by location or source. In one example, all environmental features are classified as environmental features, temperature features are classified as temperature features, and temperature features of a common species may be aggregated by data aggregator 110 as being species that may be usefully combined. Data aggregator 110 may then provide one or more outputs to expert system 130.

[0049] Expert system 130 may include one or more devices operable to collect features and to control one or more controlled systems 140. Expert system 130 may make decisions based on lookup tables, computer models, algorithms, or machine learning techniques. Features provided by data aggregator 110 may provide key inputs into the decisions that expert system 130 must make.

[0050] Controlled systems 140 may include a number of real-world systems, such as air-conditioning, environmental systems, security systems, traffic systems, space-based systems, and any other system subject to automated control or data-driven operation. Controlled system 140 may include in certain embodiments facilities to measure the response of controlled systems 140 to inputs from expert system 130. Controlled systems 140 may then provide feedback 142 to data aggregator 110. This may allow data aggregator 110 to measure the effect of combining or cross correlating certain features. In cases where data aggregator 110 determines that combining or cross correlating certain features has minimal impact on controlled systems 140, or in some cases even negative impact on controlled systems 140, data aggregator 110 may elect to unmerge certain features that are not found to be usefully combined.

[0051] In an example, each data source 120 may include an appropriate operating system, such as Microsoft Windows, Linux, Android, Mac OSX, Apple

iOS, Unix, or similar. Some of the foregoing may be more often used on one type of device than another. For example, desktop computers or engineering workstations may be more likely to use one of Microsoft Windows, Linux, Unix, or Mac OSX. Laptop computers, which are usually a portable off-the-shelf device with fewer customization options, may be more likely to run Microsoft Windows or Mac OSX. Mobile devices may be more likely to run Android or iOS. Embedded devices and dedicated appliances may run real-time operating systems such as real-time Linux, QNX, VxWorks, or FreeRTOS. For embedded devices without real-time demands, minimal Linux-based operating systems are currently very popular. However, all of the foregoing examples are intended to be nonlimiting.

[0052] Network 170 may be any suitable network or combination of one or more networks operating on one or more suitable networking protocols, including for example, a local area network, an intranet, a virtual network, a wide area network, a wireless network, a cellular network, Bluetooth connections, or the Internet (optionally accessed via a proxy, virtual machine, or other similar security mechanism) by way of nonlimiting example. Importantly, network 170 need not be an IP-based network, but is broadly intended to encompass any suitable interconnect that allows devices to communicate with one another. This could include direct serial or parallel connections, Bluetooth, infrared communications, packet radio, telephony, or any other suitable communication link.

[0053] Certain functions may also be provided on one or more servers, or one or more "microclouds" in one or more hypervisors. For example, a virtualization environment such as vCenter may provide the ability to define a plurality of "tenants," with each tenant being functionally separate from each other tenant, and each tenant operating as a single-purpose microcloud. Each microcloud may serve a distinctive function, and may include a plurality of virtual machines (VMs) of many different flavors, including agentful and agentless VMs.

[0054] In certain examples, monitored network 100 (or suitable portions thereof) may form an IoT "realm" or "domain," or may be part of a larger realm or domain.

[0055] FIGURE 2 is a block diagram of data source 120 according to one or more examples of the present specification. Data source 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, workstation, server, mainframe, virtual machine (whether emulated or on a "bare-metal" hypervisor), embedded computer, embedded controller, embedded sensor, personal digital assistant, laptop computer, cellular telephone, IP telephone, smart phone, tablet computer, convertible tablet computer, computing appliance, network appliance, receiver, wearable computer, handheld calculator, or any other electronic, microelectronic, or microelectromechanical device for processing and communicating data. Any computing device may be designated as a host on the network. Each computing device may refer to itself as a "local host," while any computing device external to it may be designated as a "remote host."

[0056] Data source 120 includes a processor 210 connected to a memory 220, having stored therein executable instructions for providing an operating system 222 and at least software portions of a data collection engine 224. Other components of data source 120 include a storage 250, network interface 260, and peripheral interface 240. This architecture is provided by way of example only, and is intended to be non-exclusive and non-limiting. Furthermore, the various parts disclosed are intended to be logical divisions only, and need not necessarily represent physically separate hardware and/or software components. Certain computing devices provide main memory 220 and storage 250, for example, in a single physical memory device, and in other cases, memory 220 and/or storage 250 are functionally distributed across many physical devices. In the case of virtual machines or hypervisors, all or part of a function may be provided in the form of software or firmware running over a virtualization layer to provide the disclosed logical function. In other examples, a device such as a network interface 260 may provide only the minimum hardware interfaces necessary to perform its logical operation, and may rely on a software driver to provide additional necessary logic. Thus, each logical block disclosed herein is broadly intended to include one or more logic elements configured and operable for providing the disclosed logical operation of that block. As used throughout this specification, "logic elements" may include hardware, external hardware (digital, analog, or mixed-signal), software, reciprocating software, services, drivers, interfaces,

components, modules, algorithms, sensors, components, firmware, microcode, programmable logic, or objects that can coordinate to achieve a logical operation.

[0057] In an example, processor 210 is communicatively coupled to memory 220 via memory bus 270-3, 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 220 communicates with processor 210 via system bus 270-1 or some other bus. Processor 210 may be communicatively coupled to other devices via a system bus 270-1. 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.

[0058] In various examples, a "processor" may include any combination of logic elements operable to execute instructions, whether loaded from memory, or implemented directly in hardware, including by way of non-limiting example a microprocessor, digital signal processor, field-programmable gate array, graphics processing unit, programmable logic array, application-specific integrated circuit, or virtual machine processor. In certain architectures, a multi-core processor may be provided, in which case processor 210 may be treated as only one core of a multi-core processor, or may be treated as the entire multi-core processor, as appropriate. In some embodiments, one or more co-processor may also be provided for specialized or support functions.

[0059] Processor 210 may be connected to memory 220 in a DMA configuration via DMA bus 270-3. To simplify this disclosure, memory 220 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

220 may comprise a relatively low-latency volatile main memory, while storage 250 may comprise a relatively higher-latency non-volatile memory. However, memory 220 and storage 250 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.

[0060] Storage 250 may be any species of memory 220, or may be a separate device. Storage 250 may include one or more non-transitory computer-readable mediums, including by way of non-limiting example, 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 250 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 222 and software portions of data collection engine 224. Many other configurations are also possible, and are intended to be encompassed within the broad scope of this specification.

[0061] Network interface 260 may be provided to communicatively couple data source 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 be used 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, direct parallel or serial connection, packet radio,

or any other appropriate architecture or system that facilitates communications in a network or telephonic environment.

[0062] Data collection engine 224, in one example, is operable to carry out computer-implemented methods as described in this specification. Data collection engine 224 may include one or more tangible non-transitory computer-readable mediums having stored thereon executable instructions operable to instruct a processor to provide a data collection engine 224. As used throughout this specification, an "engine" includes any combination of one or more logic elements, of similar or dissimilar species, operable for and configured to perform one or more methods provided by the engine. Thus, data collection engine 224 may comprise one or more logic elements configured to provide methods as disclosed in this specification. In some cases, data collection engine 224 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, data collection engine 224 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 certain embodiments, daemon processes may run with elevated privileges in a "driver space," or in ring 0, 1, or 2 in a protection ring architecture. It should also be noted that data collection engine 224 may also include other hardware and software, including configuration files, registry entries, and interactive or user-mode software by way of non-limiting example.

[0063] In one example, data collection engine 224 includes executable instructions stored on a non-transitory medium operable to perform a method according to this specification. At an appropriate time, such as upon booting data source 120 or upon a command from operating system 222 or a user, processor 210 may retrieve a copy of the instructions from storage 250 and load it into memory 220. Processor 210 may then iteratively execute the instructions of data collection engine 224 to provide the desired method.

[0064] Peripheral interface 240 may be configured to interface with any auxiliary device that connects to data source 120 but that is not necessarily a part of the core architecture of data source 120. A peripheral may be operable to provide extended functionality to data source 120, and may or may not be wholly dependent on data source 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, data ports (e.g., serial, parallel, USB, Firewire, or similar), network controllers, optical media, external storage, sensors, transducers, actuators, controllers, data acquisition buses, cameras, microphones, speakers, or external storage by way of non-limiting example.

[0065] In one example, peripherals include display adapter 242, audio driver 244, and input/output (I/O) driver 246. Display adapter 242 may be configured to provide a human-readable visual output, such as a command-line interface (CLI) or graphical desktop such as Microsoft Windows, Apple OSX desktop, or a Unix/Linux X Window System-based desktop. Display adapter 242 may provide output in any suitable format, such as a coaxial output, composite video, component video, VGA, or digital outputs such as DVI or HDMI, by way of nonlimiting example. In some examples, display adapter 242 may include a hardware graphics card, which may have its own memory and its own graphics processing unit (GPU). Audio driver 244 may provide an interface for audible sounds, and may include in some examples a hardware sound card. Sound output may be provided in analog (such as a 3.5mm stereo jack), component ("RCA") stereo, or in a digital audio format such as S/PDIF, AES3, AES47, HDMI, USB, Bluetooth or Wi-Fi audio, by way of non-limiting example.

[0066] In an example, peripherals include one or more sensors 290, which may be configured and operable to collect data about real-world phenomena and to process the data into a digital form. In one operative example, data collection engine 224 collects data from sensor 290 via peripheral interface 240. The collected data may then be stored in storage 250 and/or sent over network interface 260.

[0067] FIGURE 3 is a block diagram of a data aggregator 110 according to one or more examples of the present specification. Data aggregator 110 may be any suitable computing device, as described in connection with FIGURE 2. In general, the definitions and examples of FIGURE 2 may be considered as equally applicable to FIGURE 3, unless specifically stated otherwise.

[0068] Data aggregator 110 includes a processor 310 connected to a memory 320, having stored therein executable instructions for providing an operating system 322 and at least software portions of an aggregation engine 324. Other components of data aggregator 110 include a storage 350, network interface 360, and peripheral interface 340. As described in FIGURE 2, each logical block may be provided by one or more similar or dissimilar logic elements.

[0069] 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. Processor 310 may be communicatively coupled to other devices via a system bus 370-1.

[0070] Processor 310 may be connected to memory 320 in a DMA configuration via DMA bus 370-3, or via any other suitable memory configuration. As discussed in FIGURE 2, memory 320 may include one or more logic elements of any suitable type.

[0071] Storage 350 may be any species of memory 320, or may be a separate device, as described in connection with storage 250 of FIGURE 2. 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 aggregation engine 324.

[0072] Network interface 360 may be provided to communicatively couple data aggregator 110 to a wired or wireless network, and may include one or more logic elements as described in FIGURE 2.

[0073] Aggregation engine 324 is an engine as described in FIGURE 2 and, in one example, includes one or more logic elements operable to carry out computer-implemented methods as described in this specification. Software portions of aggregation engine 324 may run as a daemon process.

[0074] Aggregation engine 324 may include one or more non-transitory computer-readable mediums having stored thereon executable instructions operable to instruct a processor to provide a security engine. At an appropriate time, such as upon booting data aggregator 110 or upon a command from operating system 322 or a user or security administrator, processor 310 may retrieve a copy of aggregation engine 324 (or software portions thereof) from storage 350 and load it into memory 320. Processor 310 may then iteratively execute the instructions of aggregation engine 324 to provide the desired method. Operationally, aggregation engine 324 may be configured to collect and classify data provided by data sources 120.

[0075] Peripheral interface 340 may be configured to interface with any auxiliary device that connects to data aggregator 110 but that is not necessarily a part of the core architecture of data aggregator 110. Peripherals may include, by way of non-limiting examples, any of the peripherals disclosed in FIGURE 2. In some cases, data aggregator 110 may include fewer peripherals than data source 120, reflecting that it may be more focused on providing processing services rather than interfacing directly with users.

[0076] FIGURE 4 is a block diagram of an expert system 130 according to one or more examples of the present specification. Expert system 130 may be any suitable computing device, as described in connection with FIGURE 2. In general, the definitions and examples of FIGURE 2 may be considered as equally applicable to FIGURE 4, unless specifically stated otherwise.

[0077] Expert system 130 includes a processor 410 connected to a memory 420, having stored therein executable instructions for providing an operating system 422 and at least software portions of an expert system engine 424. Other components of expert system 130 include a storage 450, network interface 480, and peripheral interface 440. As described in FIGURE 2, each logical block may be provided by one or more similar or dissimilar logic elements.

[0078] 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.

[0079] Processor 410 may be connected to memory 420 in a DMA configuration via DMA bus 470-3, or via any other suitable memory configuration. As discussed in FIGURE 2, memory 420 may include one or more logic elements of any suitable type.

[0080] Storage 450 may be any species of memory 420, or may be a separate device, as described in connection with storage 250 of FIGURE 2. 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 expert system engine 424.

[0081] Network interface 480 may be provided to communicatively couple expert system 130 to a wired or wireless network, and may include one or more logic elements as described in FIGURE 2.

[0082] Expert system engine 424 is an engine as described in FIGURE 2 and, in one example, includes one or more logic elements operable to carry out computer-implemented methods as described in this specification. Software portions of expert system engine 424 may run as a daemon process.

[0083] Expert system engine 424 may include one or more non-transitory computer-readable mediums having stored thereon executable instructions operable to instruct a processor to provide a security engine. At an appropriate time, such as upon booting expert system 130 or upon a command from operating system 422 or a user or security administrator, processor 410 may retrieve a copy of expert system engine 424 (or software portions thereof) from storage 450 and load it into memory 420. Processor 410 may then iteratively execute the instructions of expert system engine 424 to provide the desired method. Operationally, expert system engine 424 may be configured to receive aggregated data from data aggregator 110 and to make decisions about how to control controlled system 140.

[0084] Peripheral interface 440 may be configured to interface with any auxiliary device that connects to expert system 130 but that is not necessarily a part of the core architecture of expert system 130. Peripherals may include, by way of non-limiting examples, any of the peripherals disclosed in FIGURE 2. In

some cases, expert system 130 may include fewer peripherals than data source 120, reflecting that it may be more focused on providing processing services rather than interfacing directly with users.

[0085] FIGURE 5 is a block diagram of a trusted execution environment (TEE) 500 according to one or more examples of the present specification.

[0086] In the example of FIGURE 5, memory 220 is addressable by n -bits, ranging in address from 0 to $2^n - 1$. Within memory 220 are an OS 222, enclave 540, application stack 520, and application code 530.

[0087] In this example, enclave 540 is a specially-designated portion of memory 220 that cannot be entered into or exited from except via special instructions, such as Intel® SGX or similar. Enclave 540 is provided as an example of a secure environment which, in conjunction with a secure processing engine 510, forms a trusted execution environment (TEE) computing device 200. A TEE 3500 is a combination of hardware, software, and/or memory allocation that provides the ability to securely execute instructions without interference from outside processes, in a verifiable way. By way of example, TEE 3500 may include memory enclave 540 or some other protected memory area, and a secure processing engine 510, which includes hardware, software, and instructions for accessing and operating on enclave 540. Non-limiting examples of solutions that either are or that can provide a TEE include Intel® SGX, ARM TrustZone, AMD Platform Security Processor, Kinibi, securiTEE, OP-TEE, TLK, T6, Open TEE, and SierraTEE, CSE, VT-x, MemCore, Canary Island, Docker, and Smack. Thus, it should be noted that in an example, secure processing engine 510 may be a user-mode application that operates via a trusted execution framework 524 within enclave 540. TEF 524 may be a framework or set of APIs that provide methods, classes, and wrappers that ease a programmer's access to TEE 500. TEE 500 may also conceptually include processor instructions that secure processing engine 510 and trusted execution framework 524 require to operate within enclave 540.

[0088] Secure processing engine 510 and trusted execution framework 524 may together form a trusted computing base (TCB), which is a set of programs or computational units that are trusted to be secure. Conceptually, it may be advantageous to keep TCB relatively small so that there are fewer attack vectors

for malware objects or for negligent software. Thus, for example, operating system 222 may be excluded from TCB, in addition to the regular application stack 520 and application code 530.

[0089] In certain systems, computing devices equipped with the Intel Software Guard Extension (SGX) or equivalent instructions may be capable of providing an enclave 540. It should be noted however, that many other examples of TEEs are available, and TEE 500 is provided only as one example thereof. Other secure environments may include, by way of nonlimiting example, a virtual machine, sandbox, testbed, test machine, or other similar device or method for providing a TEE 500.

[0090] In an example, enclave 540 provides a protected memory area that cannot be accessed or manipulated by ordinary computer instructions. Enclave 540 is described with particular reference to an Intel® SGX enclave by way of example, but it is intended that enclave 540 encompass any secure processing area with suitable properties, regardless of whether it is called an “enclave.”

[0091] One feature of an enclave is that once an enclave region 540 of memory 220 is defined, as illustrated, a program pointer cannot enter or exit enclave 540 without the use of special enclave instructions or directives, such as those provided by Intel® SGX architecture. For example, SGX processors provide the ENCLU[EENTER], ENCLU[ERESUME], and ENCLU[EEXIT]. These are the only instructions that may legitimately enter into or exit from enclave 540.

[0092] Thus, once enclave 540 is defined in memory 220, a program executing within enclave 540 may be safely verified to not operate outside of its bounds. This security feature means that secure processing engine 510 is verifiably local to enclave 540. Thus, when an untrusted packet provides its content to be rendered with secure processing engine 510 of enclave 540, the result of the rendering is verified as secure.

[0093] Enclave 540 may also digitally sign its output, which provides a verifiable means of ensuring that content has not been tampered with or modified since being rendered by secure processing engine 510. A digital signature provided by enclave 540 is unique to enclave 540 and is unique to the hardware of the device hosting enclave 540.

[0094] FIGURE 6 is a block diagram of an IoT network according to one or more examples of the present specification. In this example, a realm R1 680-1 communicates with a realm R2 680-2. As used herein, a "realm" includes a domain of interpretation for security policies, monitoring, and key management. Realm R1 includes one or more devices 120-1, which in an example each have a TEE 500. A sentinel device S1 640-1 may be configured to perform a service appliance function, such as a security function, and also includes a TEE 500. A key distribution center KDC1 650-1 provides security functionality and key management.

[0095] In this example, D1 120-1 needs to communicate with D2 120-2. D2 120-2 resides in a different realm, R2 680-2. R2 680-2 similarly includes a sentinel S2 640-2, and a KDC2 650-2.

[0096] KDC1 650-1 and KDC2 650-2 may be configured to communicate with one another as necessary, for example to perform attestation and dynamic inter-realm security policy and key management.

[0097] In this example, the communication path between D1 and D2 follows the form D1 <-> S1 <-> S2 <-> D2.

[0098] In an example, realms R1 680-1 and R2 680-2 may cooperate to agree upon a security monitoring policy applied by S1 and S2 respectively. A key management function allows KDC1 650-1 and KDC2 650-2 to issue shared symmetric keys to each of a plurality of participating devices (D1, D2, and possibly others) and sentinels (S1, S2, and possibly others) such that an end-to-end encryption scheme may be intermediated by trusted intermediaries that are identifiable and distinguished from untrusted intermediaries otherwise known as MITM attackers.

[0099] KDC2 650-2 attests a trusted execution environment (TEE) of the sentinels (S1, S2) before extending or otherwise delegating the monitoring functions. This may ensure, for example, that the security features of the sentinels are at least as secure as the security features of the devices (D1, D2), or that they otherwise meet minimum security requirements for the network. Security monitoring functions may include deep packet inspection, network and session

monitoring, reputation processing, security incident event management (SIEM), and network access control (NAC) by way of nonlimiting example.

[0100] In an example, KDC 1 650-1 may negotiate with KDC2 650-2 to establish a cross-domain security monitoring and key management policy where a symmetric key shared with local device D1 120-2 and local sentinel S1 640-1 may also be shared with a remote device D2 120-2 and a remote sentinel S2 640-2. The sharing of keys is conditioned on an expectation that the shared key (K1) is protected and otherwise not misused. A trusted execution environment (TEE) is used by both S1 640-1 and S2 640-2 to establish a basis for trustworthy operation of the security processing functions.

[0101] An attestation protocol such as Sigma may be used by sentinels 640 and devices 120 to prove the existence of a hardened environment, including a TEE that can be used to protect shared keys and to perform security monitoring functions.

[0102] The configuration and topology of sentinels within a realm need not be restricted to a single physical device performing all security operations, but may be distributed across a plurality of devices each performing a subset of possible security monitoring, reputation, deep packet inspections, NAC, SIEM, or other functions. Each shares key K1 to perform encryption and re-encryption as needed to perform its respective function.

[0103] A device or sentinel with sufficient security protections and trust may be deemed by a realm owner to be satisfactory such that the KDC function may be hosted by that device. Thus, in some embodiments, the respective roles of "device," "sentinel," and "KDC" may be combined in a single device, each function may be provided by a separate device, or functions may be distributed across multiple devices according to the physical and logical constraints of the system.

[0104] FIGURE 7 is a block diagram of a KDC 650 according to one or more examples of the present specification.

[0105] KDC 650 may be any suitable computing device, as described in connection with FIGURE 2. In general, the definitions and examples of FIGURE 2

may be considered as equally applicable to FIGURE 7, unless specifically stated otherwise.

[0106] KDC 650 includes a processor 710 connected to a memory 720, having stored therein executable instructions for providing an operating system 722 and at least software portions of an key management engine 724. Other components of KDC 650 include a storage 750, network interface 760. As described in FIGURE 2, each logical block may be provided by one or more similar or dissimilar logic elements.

[0107] In an example, processor 710 is communicatively coupled to memory 720 via memory bus 770-3, which may be for example a direct memory access (DMA) bus. Processor 710 may be communicatively coupled to other devices via a system bus 770-1.

[0108] Processor 710 may be connected to memory 720 in a DMA configuration via DMA bus 770-3, or via any other suitable memory configuration. As discussed in FIGURE 2, memory 720 may include one or more logic elements of any suitable type.

[0109] Storage 750 may be any species of memory 720, or may be a separate device, as described in connection with storage 250 of FIGURE 2. Storage 750 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 722 and software portions of key management engine 724.

[0110] Network interface 760 may be provided to communicatively couple KDC 650 to a wired or wireless network, and may include one or more logic elements as described in FIGURE 2.

[0111] Key management engine 724 is an engine as described in FIGURE 2 and, in one example, includes one or more logic elements operable to carry out computer-implemented methods as described in this specification. Software portions of key management engine 724 may run as a daemon process.

[0112] Key management engine 724 may include one or more non-transitory computer-readable mediums having stored thereon executable instructions operable to instruct a processor to provide a security engine. At an

appropriate time, such as upon booting KDC 650 or upon a command from operating system 722 or a user or security administrator, processor 710 may retrieve a copy of key management engine 724 (or software portions thereof) from storage 750 and load it into memory 720. Processor 710 may then iteratively execute the instructions of key management engine 724 to provide the desired method. Operationally, key management engine 724 may be configured to receive aggregated data from data aggregator 110 and to make decisions about how to control controlled system 140.

[0113] FIGURE 8 is a block diagram of a sentinel device 640 according to one or more examples of the present specification. Expert system 130 may be any suitable computing device, as described in connection with FIGURE 2. In general, the definitions and examples of FIGURE 2 may be considered as equally applicable to FIGURE 8, unless specifically stated otherwise.

[0114] Expert system 130 includes a processor 810 connected to a memory 820, having stored therein executable instructions for providing an operating system 822 and at least software portions of a service appliance engine 824. Other components of sentinel 640 include a storage 850, network interface 860. As described in FIGURE 2, each logical block may be provided by one or more similar or dissimilar logic elements.

[0115] In an example, processor 810 is communicatively coupled to memory 820 via memory bus 870-3, which may be for example a direct memory access (DMA) bus. Processor 810 may be communicatively coupled to other devices via a system bus 870-1.

[0116] Processor 810 may be connected to memory 820 in a DMA configuration via DMA bus 870-3, or via any other suitable memory configuration. As discussed in FIGURE 2, memory 820 may include one or more logic elements of any suitable type.

[0117] Storage 850 may be any species of memory 820, or may be a separate device, as described in connection with storage 250 of FIGURE 2. Storage 850 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 822 and software portions of service appliance engine 824.

[0118] Network interface 860 may be provided to communicatively couple sentinel 640 to a wired or wireless network, and may include one or more logic elements as described in FIGURE 2.

[0119] Service appliance engine 824 is an engine as described in FIGURE 2 and, in one example, includes one or more logic elements operable to carry out computer-implemented methods as described in this specification. Software portions of service appliance engine 824 may run as a daemon process.

[0120] Service appliance engine 824 may include one or more non-transitory computer-readable mediums having stored thereon executable instructions operable to instruct a processor to provide a security engine. At an appropriate time, such as upon booting sentinel 640 or upon a command from operating system 822 or a user or security administrator, processor 810 may retrieve a copy of service appliance engine 824 (or software portions thereof) from storage 850 and load it into memory 820. Processor 810 may then iteratively execute the instructions of service appliance engine 824 to provide the desired method. Operationally, service appliance engine 824 may be configured to receive aggregated data from data aggregator 110 and to make decisions about how to control controlled system 140.

[0121] FIGURE 9 is a flow chart of a method 900 according to one or more examples of the present specification. In an example, method 900 establishes the suitability of each participant in a secure transaction, monitored in an end-to-end session between IoT devices D1 and D2, where one or more sentinels spanning one or more administrative realms have domain specific security policies that are managed by respective KDCs.

[0122] Starting in block 900, in block 910, D1 requests access to D2 using KDC1.

[0123] In block 920, KDC1 establishes a shared session key K1 context. A topic T1 may also be assigned to the context.

[0124] In block 930, KDC1 (or any other suitable network broker) notifies S1, D2, and optionally S2, and KDC2 of the topic T1, by which D1 and D2 may securely interact.

[0125] Decision block 940 checks whether KDC2 is in charge of the second realm R2. If yes, then in block 942, each of D1 and S1 may request the topic key K1 from KDC1 and D2. S2 may request it from KDC2.

[0126] In block 944, KDC2 connects with KDC1 to negotiate a common security policy for interactions between D1 and D2. Control then passes to block 960.

[0127] Returning to block 940, if KDC2 is not in charge of a second realm R2, then in block 950, each of D1, D2, S1, and S2 may request the topic key K1 from KDC1.

[0128] In block 960, KDC1 or KDC2 may attest D1, S1, and D2, S2 respectively to establish the TEE properties of each.

[0129] In decision block 970, KDC1 and KDC2 may check whether the TEE and other security capabilities of each device match the minimum security requirements, as determined by policy. If not, then the sentinels may not be used within the data flow, and in block 999 the method is done.

[0130] On the other hand, if the devices meet the minimum security requirements, then in block 980, the appropriate KDC may issue a "mini ticket" or other encrypted response over K1 for each participant device (D1, S1, D2, S2).

[0131] In block 990, D1 opens a secure connection to D2 with S1 and S2 as intermediary sentinels. Data may then be exchanged as necessary.

[0132] In block 999, the method is done.

[0133] After this method is complete, sentinel S1 has shared the key K1 that allows access to the plaintext traffic within a TEE such that security processes can mitigate malware, protect against protocol exploits, perform traffic analysis, reputation, NAC, and SIEM, provide data leakage protection (monitoring for exposure of sensitive data), or any other suitable security service, within the context of an end-to-end data exchange.

[0134] As inspections are completed inside a TEE, this prevents exposure of keys and data to unauthorized or untrusted devices and intermediaries.

[0135] The KDC device can define the minimum security requirements for participant devices. For example the KDC device can establish trust in a subset of

sentinel devices such as those having secure boot or trusted update capability and where a TEE is operating.

[0136] With symmetric key cyphers, sentinel devices may be able to inspect both the incoming and outgoing traffic between devices D1 and D2 without D1 and D2 specifically being aware of or authenticating to any of the sentinel devices. This transparency is achieved using the KDC, which establishes a security context (topic) through which the collection of participants cooperate to achieve the security objective.

[0137] Use of asymmetric keys may also be allowed in cases where intermediaries should be explicitly authenticated. In such situations, the intermediary is issued an asymmetric certificate by a domain authority and the shared symmetric key K1 is signed by the certificate private key issued to the sending sentinel or device. The upstream device uses its realm trust anchor to verify the sender's certificate. The pre-shared key K1 may then be used to establish temporal keys for hop-hop data protection.

[0138] 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.

[0139] 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.

[0140] 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.

[0141] 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.'

[0142] 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.

[0143] 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.

[0144] 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.

[0145] 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 (pre-AIA) or paragraph (f) of the same section (post-AIA), 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

[0146] There is disclosed in one example, a computing apparatus, comprising: a trusted execution environment (TEE); and a security engine operable to: identify a key negotiation for an encrypted connection between a first device and a second device; request a service appliance key for the key negotiation; receive the service appliance key; and perform a service appliance function on traffic between the first device and the second device.

[0147] There is also disclosed an example, wherein the first device is in a first realm and the second device is in a second realm.

[0148] There is also disclosed an example, wherein the security engine further comprises logic to establish a connection between the first device and a sentinel device logically disposed between the computing apparatus and the second device.

[0149] There is also disclosed an example, wherein the security engine is operable to route traffic from the first device to the sentinel device.

[0150] There is also disclosed an example, wherein the security engine is operable to perform routing at an application layer.

[0151] There is also disclosed an example, wherein the security engine is operable to perform internet protocol routing.

[0152] There is also disclosed an example, wherein the computing apparatus further comprises a service appliance engine operable to perform a service appliance security monitoring function.

[0153] There is also disclosed an example, wherein the security monitoring function is selected from the group consisting of network monitoring, data loss prevention, packet processing, security scanning, antivirus, firewall, deep packet inspection, reputation services, security information and event monitoring, and network access control.

[0154] There is also disclosed an example, wherein the security engine is operable to establish a secure connection with the service appliance key.

[0155] There is also disclosed an example, wherein the service appliance key is different from a key for the encrypted connection between the first device and second device.

[0156] There is also disclosed an example, wherein the security engine is operable to receive a service appliance ticket comprising the service appliance key and a second key different from the service appliance key.

[0157] There is also disclosed an example, wherein the encrypted connection between the first device and the second device is encrypted within the TEE.

[0158] There is also disclosed in an example, a computing apparatus, comprising: a key store; and a key management engine, operable to: receive a negotiation between a first device and a second device; allocate a shared key for the negotiation; receive a service appliance key request from a sentinel device; issue a service appliance ticket to the sentinel device.

[0159] There is also disclosed an example, wherein the key management engine is further configured to perform an attestation of the security capability of a TEE of the sentinel device, comprising verifying that a trusted execution environment (TEE) of the sentinel device is at least as capable as a TEE of the first device.

[0160] There is further disclosed an example of one or more tangible, non-transitory computer-readable storage mediums having stored thereon executable instructions for instructing one or more processors for providing a security engine and/or key management engine operable for performing any or all of the operations of the preceding examples.

[0161] There is further disclosed an example of a method of providing a security engine and/or key management engine comprising performing any or all of the operations of the preceding examples.

[0162] There is further disclosed an example of an apparatus comprising means for performing the method.

[0163] There is further disclosed an example wherein the means comprise a processor and a memory.

[0164] There is further disclosed an example wherein the means comprise one or more tangible, non-transitory computer-readable storage mediums.

[0165] There is further disclosed an example wherein the apparatus is a computing device.

Claims:

1. A computing apparatus, comprising:
 - a trusted execution environment (TEE); and
 - a security engine operable to:
 - identify a key negotiation for an encrypted connection between a first device and a second device;
 - request a service appliance key for the key negotiation;
 - receive the service appliance key; and
 - perform a service appliance function on traffic between the first device and the second device.
2. The computing apparatus of claim 1, wherein the first device is in a first realm and the second device is in a second realm.
3. The computing apparatus of claim 1, wherein the security engine further comprises logic to establish a connection between the first device and a sentinel device logically disposed between the computing apparatus and the second device.
4. The computing apparatus of claim 3, wherein the security engine is operable to route traffic from the first device to the sentinel device.
5. The computing apparatus of claim 4, wherein the security engine is operable to perform routing at an application layer.
6. The computing apparatus of claim 4, wherein the security engine is operable to perform internet protocol routing.

7. The computing apparatus of any of claims 1 – 6, wherein the computing apparatus further comprises a service appliance engine operable to perform a service appliance security monitoring function.

8. The computing apparatus of claim 7, wherein the security monitoring function is selected from the group consisting of network monitoring, data loss prevention, packet processing, security scanning, antivirus, firewall, deep packet inspection, reputation services, security information and event monitoring, network access control, and anomaly prevention.

9. The computing apparatus of any of claims 1 – 6, wherein the security engine is operable to establish a secure connection with the service appliance key.

10. The computing apparatus of claim 9, wherein the service appliance key is different from a key for the encrypted connection between the first device and second device.

11. The computing apparatus of any of claims 1 – 6, wherein the security engine is operable to receive a service appliance ticket comprising the service appliance key and a second key different from the service appliance key.

12. The computing apparatus of any of claims 1 – 6, wherein the encrypted connection between the first device and the second device is encrypted within the TEE.

13. One or more tangible, non-transitory computer readable storage mediums having stored thereon executable instructions for providing a security engine within a trusted execution environment (TEE), wherein the security engine is operable to:

identify a key negotiation for an encrypted connection between a first device and a second device;

request a service appliance key for the key negotiation;

receive the service appliance key; and

perform a service appliance function on traffic between the first device and the second device.

14. The one or more tangible, non-transitory computer readable mediums of claim 13, wherein the first device is in a first realm and the second device is in a second realm.

15. The one or more tangible, non-transitory computer readable mediums of claim 13, wherein the security engine further comprises logic to establish a connection between the first device and a sentinel device logically disposed between the computing apparatus and the second device.

16. The one or more tangible, non-transitory computer readable mediums of claim 15, wherein the security engine is operable to route traffic from the first device to the second device.

17. The one or more tangible, non-transitory computer readable mediums of claim 16, wherein the security engine is operable to perform routing at an application layer.

18. The one or more tangible, non-transitory computer readable mediums of claim 16, wherein the security engine is operable to perform internet protocol routing.

19. The one or more tangible, non-transitory computer readable mediums of any of claims 13 – 18, wherein the computing apparatus further comprises a service appliance engine operable to perform a service appliance security monitoring function.

20. The one or more tangible, non-transitory computer readable mediums of claim 19, wherein the security monitoring function is selected from the group consisting of network monitoring, data loss prevention, packet processing, security scanning, antivirus, firewall, deep packet inspection, reputation

services, security information and event monitoring, network access control, and anomaly prevention.

21. The one or more tangible, non-transitory computer readable mediums of any of claims 13 – 18, wherein the security engine is operable to establish a secure connection with the service appliance key.

22. The one or more tangible, non-transitory computer readable mediums of claim 20, wherein the service appliance key is different from a key for the encrypted connection between the first device and second device.

23. The one or more tangible, non-transitory computer readable mediums of any of claims 13 – 18, wherein the encrypted connection between the first device and the second device is encrypted within the TEE.

24. A computing apparatus, comprising:

a key store; and

a key management engine, operable to:

receive a negotiation between a first device and a second device;

allocate a shared key for the negotiation;

receive a service appliance key request from a sentinel device;

issue a service appliance ticket to the sentinel device.

25. The computing apparatus of claim 24, wherein the key management engine is further configured to perform an attestation of the security capability of a TEE of the sentinel device, comprising verifying that a trusted execution environment (TEE) of the sentinel device is at least as capable as a TEE of the first device.




Fig. 1




Fig. 2




Fig. 3




Fig. 4




Fig. 5




Fig. 6




Fig. 7




Fig. 8

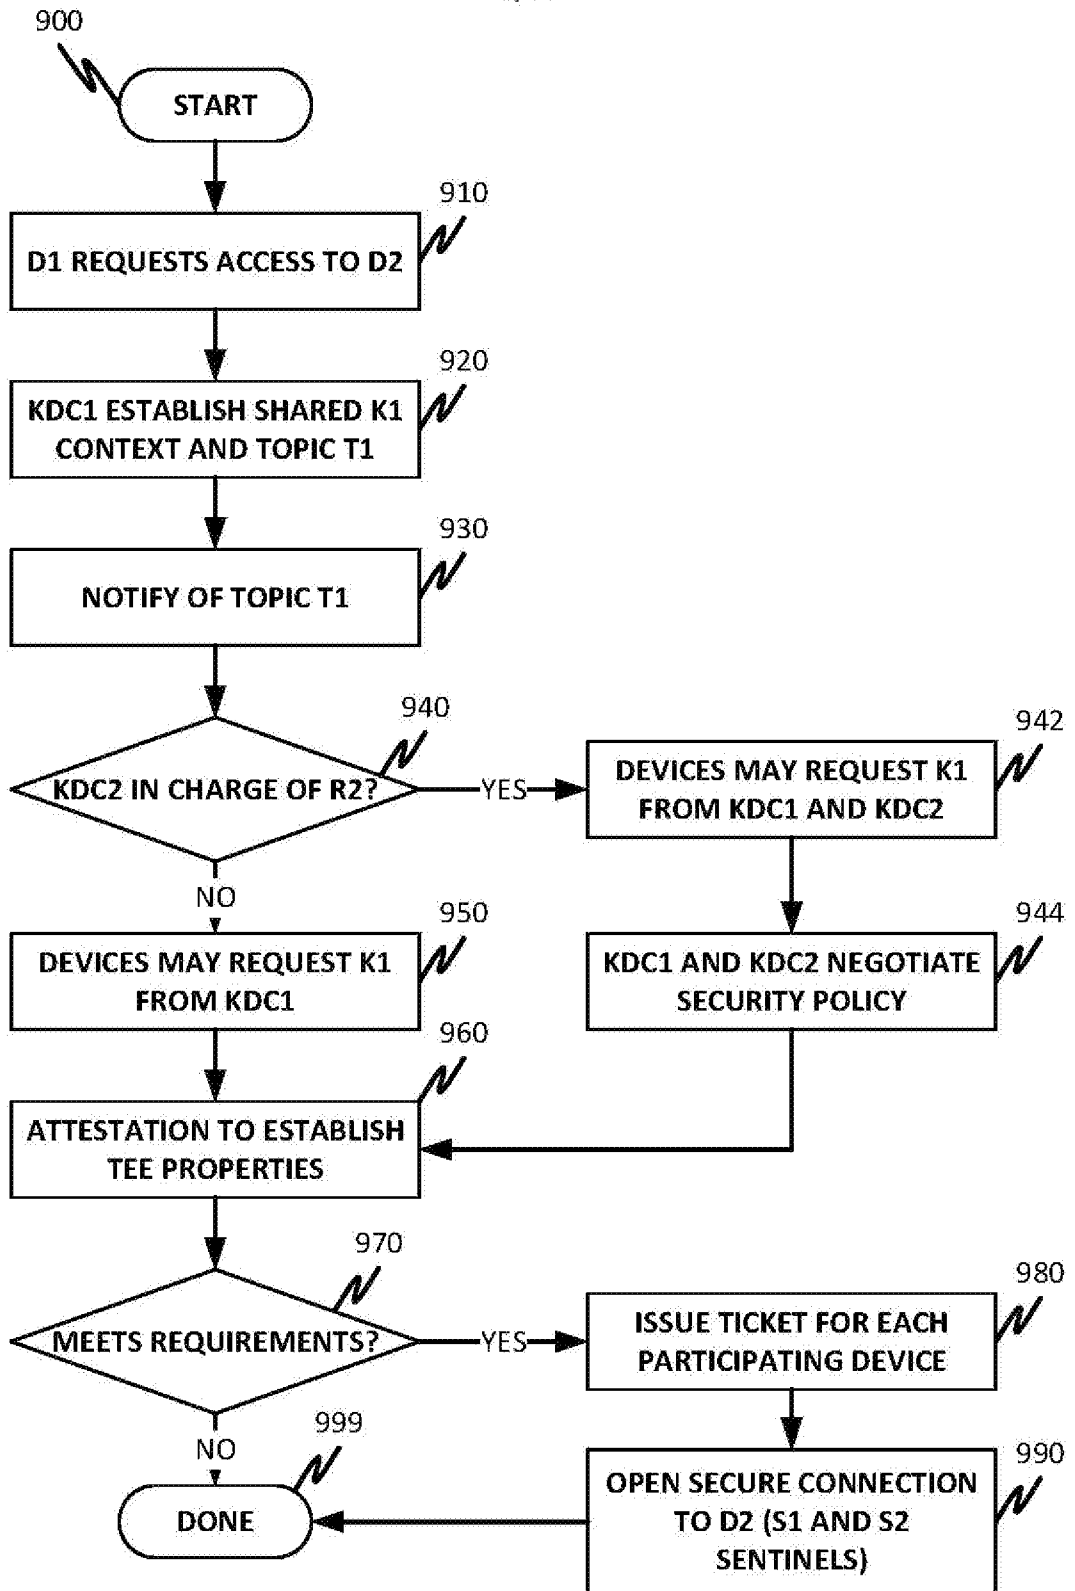


Fig. 9

A. CLASSIFICATION OF SUBJECT MATTER**H04L 29/06(2006.01)i, H04L 29/08(2006.01)i, H04L 9/08(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
H04L 29/06; H04W 12/02; G06F 21/30; H04K 1/00; G06F 21/00; H04L 29/08; H04L 9/08Documentation 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: secure connection, key, service, sentinel device, trusted execution environment (TEE)**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	WO 2014-194858 A1 (QUANTUM COMMUNICATION TECHNOLOGY CO., LTD., ANHUI et al.) 11 December 2014 See page 8, lines 17-25, page 9, lines 11-18, page 11, lines 16-28, page 12, lines 1-21, page 15, lines 5-23, claims 1-2 and figures 4-5.	1-25
Y	WO 2014-209322 A1 (INTEL CORPORATION) 31 December 2014 See page 4, lines 26-29, page 6, lines 3-6, page 10, lines 4-14, page 11, lines 8-29, page 22, lines 13-20, claim 1 and figures 1-2.	1-25
A	WO 2014-165747 A1 (INTERDIGITAL PATENT HOLDINGS, INC.) 09 October 2014 See paragraphs [0053]-[0054], claims 1-16 and figure 7A.	1-25
A	US 2012-0304310 A1 (JAMES BLAISDELL) 29 November 2012 See claims 1-7 and figure 1B.	1-25
A	US 2007-0154016 A1 (MADJID F. NAKHJIRI et al.) 05 July 2007 See paragraphs [0057]-[0060] 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

"&" document member of the same patent family

Date of the actual completion of the international search

08 August 2016 (08.08.2016)

Date of mailing of the international search report

08 August 2016 (08.08.2016)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

LEE, EUN KYU

Telephone No. +82-42-481-3580



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/031565

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2014-194858 A1	11/12/2014	CN 104243143 A EP 3007478 A1 US 2016-0127127 A1	24/12/2014 13/04/2016 05/05/2016
WO 2014-209322 A1	31/12/2014	CN 105247528 A EP 3014507 A1 KR 10-2016-0004353 A US 2016-0127351 A1	13/01/2016 04/05/2016 12/01/2016 05/05/2016
WO 2014-165747 A1	09/10/2014	CN 105103578 A EP 2982148 A1 KR 10-2015-0139602 A TW 201511513 A US 2016-0065362 A1	25/11/2015 10/02/2016 11/12/2015 16/03/2015 03/03/2016
US 2012-0304310 A1	29/11/2012	US 8769305 B2	01/07/2014
US 2007-0154016 A1	05/07/2007	CN 101356759 A EP 1972089 A2 WO 2007-081588 A2 WO 2007-081588 A3	28/01/2009 24/09/2008 19/07/2007 27/12/2007