



- (51) International Patent Classification:
H04L 29/06 (2006.01) *H04L 9/32* (2006.01)
- (21) International Application Number:
PCT/US2019/044301
- (22) International Filing Date:
31 July 2019 (31.07.2019)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/719,341 17 August 2018 (17.08.2018) US
16/291,874 04 March 2019 (04.03.2019) US
- (71) Applicant: **T-MOBILE USA, INC.** [US/US]; 12920 SE
38th Street, Bellevue, Washington 98006-1350 (US).
- (72) Inventor: **MALIK, Sireen**; c/o T-Mobile USA, Inc., 12920
SE 38th Street, Bellevue, Washington 98006-1350 (US).
- (74) Agent: **PECK, Robert C.** et al.; Lee & Hayes, PC, 601
West Riverside, Ste 1400, Spokane, Washington 99201
(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, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, 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.

(54) Title: VERIFYING NETWORK SUBSYSTEM INTEGRITY WITH BLOCKCHAIN




FIG. 1

(57) **Abstract:** In a telecommunication network, individual network slices are provided for various uses and/or for various enterprise customers. A network subsystem such as a network slice comprises multiple components, such as routers, applications, virtual network functions, etc. Each component of the subsystem generates and provides a digital signature, such as a hash, based on state properties of the component that have been designated as being invariant. The signatures from the multiple components are then combined and hashed to form a subsystem signature. A chronological sequence of subsystem signatures is saved in a distributed ledger, which may use blockchain technology to protect against after-the-fact modifications to the saved signatures. A network threat may be detected by detecting situations in which the subsystem signature of a particular subsystem changes over time. A saved blockchain, containing the sequence of signatures, can be provided as historical evidence of network integrity.

(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))*

VERIFYING NETWORK SUBSYSTEM INTEGRITY WITH BLOCKCHAIN

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This patent application claims priority to U.S. Utility patent application with
5 Serial No. 16/291,874, filed March 4, 2019, which claims priority to U.S. Provisional
Patent Application No. 62/719,341 filed on August 17, 2018, and titled “Network-
Component Monitoring for Security”. Application Serial Nos. 16/291,874 and
62/719,341 are fully incorporated herein by reference.

BACKGROUND

10 [0002] Network slicing is a technology that has been proposed to enhance
flexibility in wireless communication networks. Network slicing allows multiple
logical networks to operate using a common underlying infrastructure. Network slicing
can reduce costs of deployments, increase the efficiencies of physical infrastructures,
and provide flexibility in the types of services that may be offered to different users.
15 Dedicated network slices can be provided for enterprise customers of a network
operator, with the slices being customized to the needs of the customers.

[0003] A slice-based communication network can have multiple software-defined
core networks, referred to as slices, each of which may be optimized for a different type
of service. For example, an LTE (Long-Term Evolution) core network can be provided
20 for 4G (4rd-Generation) LTE services and an IMS (IP Multimedia Subsystem) core
network can be provided for multimedia messaging. Generally, network slices can be
provided for many different purposes, such as Internet browsing, voice
communications, industrial control, home automation, IoT (Internet of Things), etc.
Furthermore, different slices may be configured to provide different Qualities of
25 Service (QoS). For example, one core network may be designed for low latency while
another core network might be designed to provide fast streaming.

[0004] Threat detection is an important part of any communication network,
including slice-based networks. Threat detection systems attempt to detect malicious
attacks that are designed to gain unauthorized access to private, secret, or otherwise
30 non-public information. Malware is increasingly obfuscated or otherwise disguised in
an effort to avoid detection by security software. Determining whether a program is

malware or is exhibiting malicious behavior can thus be very time-consuming and resource-intensive. As network components become more sophisticated, they may be targeted more often by malware.

BRIEF DESCRIPTION OF THE DRAWINGS

5 [0005] The detailed description is described with reference to the accompanying figures. In the figures, the left-most digit(s) of a reference number identifies the figure in which the reference number first appears. The use of the same reference numbers in different figures indicates similar or identical components or features.

[0006] FIG. 1 is a block diagram illustrating techniques for detecting malicious
10 activities performed against a network subsystem such as a network slice.

[0007] FIG. 2 is a block diagram illustrating the use of a blockchain for storing an immutable record of network subsystem integrity.

[0008] FIG. 3 is a block diagram illustrating the use of a distributed blockchain ledger for storing an immutable record of network subsystem integrity.

15 [0009] FIG. 4 is a block diagram illustrating the storage of multiple blockchains, respectively indicating historical network integrity of different network subsystems.

[0010] FIG. 5 is a flow diagram illustrating an example method of detecting malicious activities performed against a network subsystem such as a network slice.

[0011] FIGS. 6, 7, and 8 are flow diagrams illustrating example methods of
20 providing proof of historical subsystem integrity.

[0012] FIG. 9 is a block diagram of an example computing device that may be configured in accordance with the described techniques.

DETAILED DESCRIPTION

[0013] This disclosure describes, in part, techniques for detecting unauthorized
25 modifications to components of a network, e.g., physical devices connected to a network or network functions implemented on physical devices (e.g., running in a virtual machine).

[0014] In a described embodiment, different types of state properties are periodically collected from multiple components of a network slice, including physical
30 components, virtual components, and/or software components. The collected state properties, referred to herein as designated state properties, relate to component characteristics that are intended or expected to remain unchanged through the lifecycle

of the network slice. As one example, a cryptographic hash of a component's firmware may be designated as a state property that is intended to remain unchanged throughout the lifetime of a network slice. As another example, a list of users that have been granted administrative roles may be designated as a state property that should remain unchanged.

[0015] In various examples, the state of each component is represented by a signature, e.g., a cryptographic hash value computed over the state of that component. Reported signatures from the various components of the network slice are aggregated and a digital signature of the aggregated signatures is computed. Over time, this signature, referred to herein as a slice signature, is calculated over and over using the same technique to create a chronological sequence of slice signatures.

[0016] In various examples, each component is deployed in a known-good (or believed-good) state. That state is associated with a particular slice signature. The slice signature is monitored to ensure that it remains constant. If the signature varies at any time, a potential threat is reported and/or threat mitigation procedures are initiated. This can permit detecting intrusion of malware into network components without requiring deep packet inspection or other forms of behavioral detection, and without requiring antivirus systems running at each network component. Behavioral detection and antivirus software both consume resources at network components, so detecting malware without either of those can increase network capacity and throughput, reduce packet processing latency at network components, or otherwise improve network performance.

[0017] In various examples, slice signatures are generated periodically, e.g., approximately hourly or on a different time schedule. Each signature is stored as or in a transaction in a distributed ledger in association with information identifying the network slice, such as a fully-qualified domain name representing the network slice. A transaction can also include a timestamp indicating the time the signature was determined or the time the transaction was recorded.

[0018] A distributed ledger is designed to replicate and synchronize data between multiple databases. In some systems, the databases may be geographically distributed over a large geographic area, which protects the data from loss. The databases are configured to synchronize with each other using peer-to-peer communications and consensus-based data synchronization between multiple distributed data stores.

Consensus-based data synchronization algorithms are designed so that a majority or some percentage of the databases must agree upon any changes before the changes are committed. In some peer-to-peer ledgers, there may be no central database administrator, and therefore no single point that can be attacked to introduce unauthorized changes.

[0019] In some embodiments, data within a distributed ledger may be stored using blockchain technology, in which data is recorded in blocks that are cryptographically linked, in the sequence of their creation or addition, to prevent after-the-fact changes of the data. Data stored in this manner is considered to be immutable because once stored, it would be very difficult to alter the data. Accordingly, it can serve as a historical record of network slice integrity.

[0020] The distributed ledger may be either private or public. In a private distributed ledger, also referred to as a permissioned ledger, multiple databases are administered by a single entity and/or by other specifically authorized entities. In a public distributed ledger, also referred to as a permissionless ledger, the databases can be administered by any entity. In either case, consensus algorithms ensure that no individual entity can alter the saved data without the participation and agreement of some number of other entities. Because of consensus-based synchronization, an attack on the stored data needs to involve more than a given percentage of the distributed databases.

[0021] Storing slice signatures in a blockchain protects against various attack scenarios. As an example, an adversary may wish to change a stored signature in the distributed ledger, which would permit the adversary to install malware in a network component without being detected. However, the use of a blockchain forces the adversary to recalculate a large portion of the blockchain in order to change even a single signature. Furthermore, the peer-to-peer distributed ledger forces an adversary to change the blockchain at more than half of the peers (or another value specific to the type of consensus algorithm) in order for the falsified signature to be accepted into the distributed ledger.

[0022] The ability to provide an auditable, immutable history that demonstrates network slice integrity may be a valuable feature for which enterprise customers are willing to pay.

[0023] The techniques described herein may allow the programmatic detection of malicious activity. Various examples herein can be performed without human judgement regarding whether a program or data block is in fact malicious. Techniques herein can permit more readily identifying potential threats, e.g., in the context of a cloud security service or security appliance used in a telecommunication network.

[0024] Security actions taken in response to detecting a change in a slice signature may include removing a network component from the network, blocking packet flows, blocking execution of processes, terminating processes, blocking DNS lookups or network connections, terminating lookups or connections already in progress, deleting, renaming, or quarantining files, and so forth.

[0025] For purposes of discussion, the introduction above has been in the context of a telecommunication system that implements network slicing. However, the described techniques can be used in any environment having multiple components that are capable of providing state information or that are capable of providing signatures of state properties. Generally, the described techniques can be used to provide an immutable integrity record for any system or subsystem of physical, virtual, and/or software components, not limited to components of network slices.

[0026] FIG. 1 illustrates a technique for ensuring the integrity of multiple network slices or other subsystems of networkable components. The techniques described here can be implemented by an operator of a telecommunication network, such as a cellular communication network, to detect potential attacks and to provide an immutable, auditable history of network slice integrity.

[0027] FIG. 1 shows a network slice 102 as an example of multiple network slices 104 that may exist at any given time within a telecommunication network. The network slices 104 may correspond to or comprise multiple dedicated core networks, as an example. Different core networks may be provided for different types of services, for different types of customers, for different types of traffic, to provide different levels of QoS (Quality of Service), etc. Various of the network slices 104 may also implement other parts of a telecommunication system, such as radio access networks, IP Multimedia Subsystems, and so forth.

[0028] In certain embodiments, the network slices 104 may comprise software-defined networks, and may at times be referred to as “virtual” network slices. In certain situations, network slices may be dynamically created and destroyed in response to

increasing and decreasing demands for services or in response to other variable needs of enterprise customers.

[0029] The network slice 102 comprises multiple components 106, including physical components as well as software and virtual components that may be implemented by or based on the physical components. The components 106 may include, without limitation:

5 routers;
switches;
gateways;
10 Internet-of-things devices;
virtual network functions;
containers;
software applications;
software modules;
15 virtual machines;
radios;
radio access networks;
Internet Multimedia Subsystems (IMSs);
etc.

20 [0030] The components 106 may be located in geographically diverse locations and may communicate with each other through available data networks, which may include both wired and wireless data networks. In many cases, the components communicate using Internet Protocol (IP) packet-based messaging.

[0031] The network slice 102 may be created by instantiating virtual components, using a hardware infrastructure that is common to multiple components and/or network slices. The components 106 are initially configured as needed to perform their designated functions within the network slice 102. Configuration may include installing software, configuring settings, initializing data, allocating or assigning resources, defining functions, setting permissions, and so forth. Instantiation and configuration may be performed in some cases as an automated process, based on predefined parameters or code/data images.

[0032] Each component 106 has certain state properties that remain constant during operation. Any change to these state properties represents a potential violation of

network integrity. As an example, a checksum of a component's operating system code might be expected to be unchanging during operation of the network slice. As another example, the identity of users having administrator privileges may be expected to be unchanging. Invariant state properties may include checksums, signatures, hashes,
5 lists, parameters, settings, and so forth. The actual state properties that are to remain constant over time may vary from one component to another.

[0033] Certain types of component states may be represented by digital signatures. For example, firmware state may be represented as a signature of the firmware code, where the signature is generated in such a way that any change to the firmware code
10 will result in a different signature. Lists, tables, data objects, software applications, software modules, and other types of data may also be represented as signatures.

[0034] For purposes of this disclosure, the term "signature" is used to indicate a bit string that (a) is calculated from or otherwise generated as a function of input data, (b) remains constant when the input data remains constant, and (c) changes when the input
15 data changes. Signatures may include message digests, checksums, digital fingerprints, hash values, and other bit strings that may be calculated in various ways from the input data. A signature is typically much smaller than the input data. In some cases, signatures may have a fixed length, regardless of the size of the input data.

[0035] A hash value is a bit string that is calculated using a mathematical one-way
20 function, referred to as a hash function, so that the original input data cannot be derived from the hash value. When the term "hash" is used as a verb, it refers to the creation of a hash value by applying a hash function. The term "hash" may also be used as a noun to reference a hash value.

[0036] In the example configuration shown by FIG. 1, a threat monitor 108 is
25 configured to receive a component signature 110 from each of the slice components 106. The threat monitor 108 is shown here as a representation of logical functionality that may be performed by any of various entities of a communication network and/or by entities associated with the communication network. In practice, the functionality attributed herein to the threat monitor 108 may be performed by one or multiple
30 elements of a network infrastructure, including both special-purpose and general-purpose hardware and software. As an example, the functions of the threat monitor 108 may be performed by the same entity that supervises and manages the network slices 104.

[0037] Each component signature 110 represents the current state of a respective component 106. To create a component signature 110, various state properties of the component as described above are identified and designated as being invariant. The designated invariant properties are collected and aggregated in a repeatable way to
5 create a data object that is then hashed to create the component signature 110.

[0038] The term “invariant” is used herein to indicate properties that do not change during normal operations, but which may change in response to abnormal circumstances such as when an unauthorized entity is attempting to gain access to a component or network slice.

10 [0039] In some embodiments, a component 106 may itself calculate and provide its corresponding component signature 110. For example, a daemon (not shown) may run within the slice 102 or within each component 106 of the slice 102 to gather state properties and to create the component signatures 110.

[0040] In other embodiments, the threat monitor 108 may query the network slice
15 102 and/or the components 106 of the network slice for state properties. The threat monitor 108 may obtain state properties for a component 106 and calculate the corresponding component signature 110 based on the obtained state properties.

[0041] The threat monitor 108 receives or generates the component signatures 110, concatenates them or otherwise combines them to form a single data object, and hashes
20 the data object to create a slice signature 112. The slice signature 112 corresponds to the current state of the slice 102, and is expected to remain invariant.

[0042] In some embodiments, the slice signature may comprise a Merkle tree root hash of the component signatures 110.

[0043] For each slice 102, a chronological sequence of slice signatures 112 may be
25 generated and stored. That is, the threat monitor 108 may periodically receive the component signatures 110 from the respective components 106 (or create the signatures 110 from state properties received from the components 106) and calculate a new slice signature 112 based on the component signatures 110.

[0044] The threat monitor 108 monitors the slice signature 112 to detect changes in
30 the slice signature 112 over time. Any change in the slice signature 112 signifies a change in one or more designated invariant state properties of the slice 102, which in turn signifies a potential threat, intrusion, or other violation of slice integrity. Upon detecting a change, the threat monitor 108 may notify a client 114 such as a threat

management system maintained by the operator of the telecommunication network. In response, the client 114 may generate a notification and/or initiate threat mitigation procedures. The threat monitor 108 may have an API and/or other means for communicating with a client 114 and with other network entities.

- 5 [0045] In some cases, the client 114 may periodically query the threat monitor 108 to obtain the most recent slice signature 112 and may itself determine whether there have been any changes in the slice signature 112. In some cases, the client 114 may query or signal the threat monitor 108 to initiate the collection of the component signatures 110 and the creation of the slice signature 112. That is, the process of
10 collecting the component signatures 110 and generating the slice signature 112 may be driven by the client 114.

- [0046] In some embodiments, the slice signature 112 may be provided to a distributed ledger system 116 for long-term storage. The distributed ledger system 116 includes multiple data stores or databases 118 that are geographically distributed over
15 a large area, such as across a country or across multiple countries. The distributed ledger system 116 is designed to replicate and store identical data in each of the databases 118.

- [0047] In some examples, the distributed ledger system 116 may be implemented using cloud computing. Cloud computing allows for computing resources to be
20 provided as services rather than as a deliverable product. For example, in a cloud computing environment, resources such as computing power, software, information, or network connectivity are provided (for example, through a rental agreement) over a network, such as the Internet.

- [0048] In other examples, the distributed ledger system 116 may be implemented
25 using a plurality of computing appliances operated by or on behalf of a particular user, group, or organization. For example, a corporation may deploy an appliance per office site, per division, or per geographic region (e.g., ten or 8–12 peers per country, spread geographically). In some examples, a computing appliance can be used in addition to, or instead of, a cloud service.

- 30 [0049] For each slice 102, the threat monitor 108 may store a chronological sequence of slice signatures 112 that have been generated throughout the life of the slice 102, as a record of whether or not the designated invariant state properties of the slice 102 have changed since creation of the slice. Each slice signature 112 may be

stored along with a timestamp indicating the time at which the slice signature 112 was obtained or generated. In some embodiments, each newly generated slice signature 112 may be stored in the distributed ledger system 116 as is it generated, resulting in a continuously growing list of signatures.

5 [0050] In some cases, the distributed ledger 116 may comprise multiple geographically-distributed servers 120, each of which manages a corresponding one of the databases 118. The distributed ledger 116 may be configured to operate without a central authority and to instead use peer-to-peer communications between the servers 120 to replicate and synchronize data across the databases 118. Synchronization may
10 be performed using a consensus algorithm, such that any proposed addition, deletion, or modification to the databases 118 must be agreed upon by a predetermined percentage of the servers 120 before being committed. For example, a proposed addition of data to the distributed ledger 116 might be committed only after a majority of the servers 120 agree upon the proposed change. This makes it difficult for an
15 attacker to alter the stored data, because the attacker would need to alter the data in a majority of the databases 118 before the change would be accepted.

[0051] FIG. 2 illustrates how a chronological sequence of slice signatures 112 may be stored as a list of linked hash values to create a blockchain 202. The blockchain 202 comprises multiple data blocks 204. Each data block 204 is a data object containing a
20 slice signature 112 and a corresponding timestamp 206. As new slice signatures 112 are received, respectively corresponding new blocks 204 are added to the blockchain 202. In addition to the slice signature 112 and corresponding timestamp 206, each block 204 has a hash value 208, which is generated from and links to the previous block 204 in the blockchain 202. FIG. 2 shows multiple blocks 204 in the sequence of their
25 creation from bottom to top. The uppermost block 204 of the blockchain 202 is the most recently added block.

[0052] During normal operation, all of the signatures 112 stored by the blockchain 202 should be equal to each other. If they are not, there may have been a violation of network slice integrity. In this case, the timestamp at point within the blockchain where
30 the signature takes on a new value indicates the time at which the potential violation occurred.

[0053] As part of an attack, an attacker might attempt to alter the blockchain 202 to prevent an alarm from being raised in response to tampering with components of the

network slice 102. For example, the attacker may attempt set all of the signatures 112 to the same value. This would be difficult, however, because the chain would need to be recalculated for every block subsequent to where the change was introduced. Furthermore, this would need to be accomplished on a majority of the servers 120.

5 Otherwise, the consensus-based synchronization algorithms of the distributed ledger would prevent the attempted change from being committed.

[0054] FIG. 3 illustrates how the blockchain 202 may be stored within the distributed ledger system 116. FIG. 3 also illustrates how newly received or created slice signatures are added to a stored blockchain.

10 [0055] FIG. 3 shows a blockchain interface 302 as a general representation of functionality that might be implemented either by the distributed ledger 116, by the threat monitor 108, by one of the servers 120, or by some other component.

[0056] To store a new slice signature 304, the blockchain interface 302 obtains the existing blockchain 306 from the servers 120. In some cases, the functions of the
15 blockchain interface 302 may be performed by one of the servers 120, which may hold its own copy of the blockchain 306.

[0057] The blockchain interface 302 calculates a hash of the previous, most recent block of the blockchain 306. The blockchain interface 302 then uses the newly created hash, the new slice signature 304, and a corresponding timestamp to create a modified
20 blockchain 308, and sends the modified blockchain 308 to the servers 120. A peer-to-peer consensus procedure is then implemented to distribute the modified blockchain to all of the servers 120. This procedure may utilize an appropriate consensus algorithm, such as, for example, Byzantine Fault Tolerance, Proof-Based Consensus, Voting-Based Consensus, Crash-Based Fault tolerance, etc. Assuming that the servers 120
25 reach consensus on the data modification, the modified blockchain 308 is replicated and committed by all of the servers 120.

[0058] FIG. 4 shows how historical integrity data may be maintained for later use in verifying past network slice integrity. When the distributed ledger 116 is implemented using blockchain technology, the distributed ledger 116 stores multiple
30 blockchains 402, corresponding respectively to different network slices 104 of an operator's network. In some embodiments, the distributed ledger 116 may have an application programming interface 404 that can be accessed by the client 114 to obtain blockchains for inspection and for proof of network integrity. For example, the client

114 may be a customer application configured to obtain one or more blockchains that have integrity records for one or more network slices that are operated for the benefit of the customer. The client 114 may also comprise a threat management system as mentioned above, which may access and analyze the blockchains to detect integrity violations.

[0059] In some embodiments, the distributed ledger 116 may be a public blockchain system that can be accessed by any interested party, with the integrity of the data being ensured by the consensus synchronization between the distributed servers and databases.

[0060] In some cases, the client 114 may have read-only access to the distributed ledger 116 so that it can obtain and analyze stored data directly from the distributed ledger system 116 rather than querying the threat monitor 108 for information.

[0061] FIG. 5 illustrates an example method 500 of detecting potential threats to the integrity of a network subsystem such as a network slice of a telecommunication network. The method 500 works by monitoring signatures that have been calculated from designated invariant state properties of the components of the network subsystem. The method 500 may be performed by a monitoring system or component, such as the threat monitor 108 described above, or by any other component.

[0062] An action 502 comprises receiving designated invariant state data from the components of the network subsystem, where the designated invariant state data is based at least in part on one or more designated invariant properties of each component. An action 504 comprises calculating a component signature of the designated invariant state data of each of the components.

[0063] The term “designated invariant” is used here to refer to data, properties, information, etc. of a component that do not change during normal operation of the component. Designated invariant state data includes state properties that should remain constant. It is assumed that any change in a designated invariant state property signifies some type of problem such as a threat to the integrity of the component. In other words, a change in any designated invariant state property could be a result of unauthorized activity involving the component.

[0064] The following are examples of state properties, any of which might be designated as invariant in certain cases:

content of a software image;

- content of a firmware image;
- content of a file;
- stored data;
- a setting or configuration parameter;
- 5 whether an application or other software component is in memory and/or executing;
- a permission, such as for permitted administrators and users;
- whether a patch or update has been applied;
- content of a code image of a patch or update;
- 10 whether a file is present;
- whether a file is absent;
- whether a resource is available or sufficient;
- etc.

[0065] State data and state properties may include hashes or other signatures of the state properties above or of any other designated invariant data. For example, an MD5 (message digest 5) hash of file or firmware image might be used as an invariant property, where the firmware image itself is intended to be invariant.

[0066] In some cases, the designated invariant state data may comprise a component signature corresponding to each component of the subsystem, where the component signature is generated based on a state of the component at a corresponding time.

[0067] The actions 502 and 504 may in some cases be performed by the components themselves. For example, each component may have a background daemon that collects designated invariant state properties and hashes them to produce a component signature.

[0068] In other cases, a software function such as the threat monitor 108 may collect the designated invariant state properties from the components and calculate the component signatures based on the collected invariant state properties.

[0069] In some examples, network components can report the designated invariant state signatures periodically, e.g., hourly. In some examples, network components can report state signatures on request from a management server, such as the threat monitor 108, or at predetermined times. For example, the components in a slice primarily used in a particular time zone can be configured to report their signatures at times that are

during the business day in that time zone. This can permit the threat monitor 108 to obtain current signature information or change/no-change indications during normal business hours. In some examples, network components can be configured to stagger their reporting times to provide a steadier (less peaky) flow of transactions. For example, each network component can be configured to report at intervals of H hours ($H \geq 1$), at approximately M minutes past the hour ($M \geq 0$). Different network components can be given different M values to spread the network load and increase the transaction capacity of the distributed ledger.

[0070] In some examples, network components can report signatures in response to detection of a signature change, even if a periodic-reporting schedule would not normally call for a signature to be reported at that time. This can provide more rapid detection of malicious activity, which can in turn permit more effective mitigation.

[0071] The action 504 of calculating the component signatures may be performed in different ways. For example, the action 504 may comprise calculating a cryptographic hash of the designated invariant state data, such as by applying any of multiple available cryptographic hash functions. MD5 (message digest 5), SHA-1 (secure hash algorithm 1, and SHA-2 (secure hash algorithm 2) are examples of suitable hash functions. The component signatures may alternatively comprise checksums of the designated invariant state data.

[0072] After calculating a component signature corresponding to each of the subsystem components, an action 506 is performed of calculating a subsystem signature. The subsystem signature is a signature representing the state of the monitored subsystem. The subsystem signature can be created by concatenating the individual component signatures and applying a hash function to the resulting concatenated bit string or other data object. Again, any of multiple available cryptographic hash functions may be used in the action 506, including MD5, SHA-1, and SHA-2 hash functions. In some cases, the subsystem signature may comprise a Merkle tree root hash of the individual component signatures.

[0073] An action 508 comprises storing the subsystem signature. As described above, the subsystem signature may be stored in a distributed ledger system, such as the distributed ledger system 116. In some cases, the distributed ledger system may comprise a private or public blockchain ledger system, in which new subsystem

signatures are appended to existing blockchains as described above with reference to FIGS. 2 and 3.

[0074] An action 510 comprises determining whether the most recently generated subsystem signature is different than the most recent previous signature. If there has
5 been a change in the value of the subsystem signature, an action 512 is performed of initiating an integrity threat response, such as generating a notification or alerting a threat detection system. If the signature has not changed, the preceding actions are repeated. Any change between consecutive first and second subsystem signatures generated by the action 506 is treated as a potential threat.

10 [0075] The actions 502, 504, 506, 508, and 510 may be performed repeatedly or periodically to calculate and store a chronological sequence of multiple signatures for each subsystem, based on state data received over time from the components of the subsystem. Each signature is based on a set of the designated invariant state data corresponding to a particular time. In the case of a dynamically created network slice,
15 the stored signatures may correspond to the entire lifecycle of the network slice. The stored chronological sequence of signatures is preserved as historical evidence of the integrity of the monitored subsystem.

[0076] FIG. 6 illustrates an example method 600 of providing integrity assurance to customers that are associated with network subsystems such as network slices. An
20 action 602 comprises receiving a customer request for verification of network subsystem integrity. For example, a customer component might communicate with an application programming interface (API) of a network provider to request historical records regarding the integrity of a specified network subsystem. In response, an action 604 may be performed of providing a data object, such as one of the blockchains 402,
25 corresponding to the network subsystem. The blockchain is provided as historically immutable verification of the subsystem integrity. The customer can audit the blockchain by verifying the hash of each block, from the beginning block to the most recent block, and confirming that all the stored signatures are the same.

[0077] FIG. 7 illustrates another example method 700 of providing integrity
30 assurance to a customer that is associated with a network subsystem such as a network slice. An action 702 comprises receiving a customer request for verification of network subsystem integrity. For example, a customer might communicate with an API of a network provider to request an audit of historical records regarding the integrity of a

network slice. In response, an action 704 is performed of accessing and analyzing the blockchain corresponding to the network subsystem. For example, the action 704 may comprise verifying the hash of each block of the blockchain, from the beginning block to the most recent block, and checking that all of the stored signatures are the same.

5 [0078] An action 706 may comprise reporting the results of the analysis to the requesting customer, such as by providing an integrity report indicating whether or not the integrity of the network subsystem has been violated during its lifetime. In cases where there has been an integrity violation, the action 706 may comprising indicating the time at which the violation occurred, based on timestamps associated with the stored
10 signatures.

[0079] FIG. 8 illustrates yet another example method 800 of providing integrity assurance to a customer that is associated with a network subsystem such as a network slice. An action 802 comprises receiving a customer request for verification of network subsystem integrity. For example, a customer might communicate with an API of a
15 network provider to request an audit of historical records regarding the integrity of a network slice. In response, an action 804 is performed of providing a network address of or link to of the blockchain corresponding to the network subsystem to the requesting entity. An action 806, which may be performed by the network-based computer system of the customer, may comprise accessing and analyzing the blockchain to verify the
20 historical integrity of the network subsystem.

[0080] FIG. 9 is a block diagram of an illustrative computing device 900 such as may be used to implement various components described above. Generally, one or more computing devices 900 may be used to implement any of the components and/or functionality of the preceding FIGS. 1-8.

25 [0081] In various embodiments, the computing device 900 may include at least one processing unit 902 and system memory 904. Depending on the exact configuration and type of computing device, the system memory 904 may be volatile (such as RAM), non-volatile (such as ROM, flash memory, etc.) or some combination of the two. The system memory 904 may include an operating system 906, one or more program
30 modules 908, and may include program data 910.

[0082] The computing device 900 may also include additional data storage devices (removable and/or non-removable) such as, for example, magnetic disks, optical disks, or tape. Such additional storage devices are illustrated in FIG. 9 as storage 912.

[0083] Non-transitory computer storage media of the computing device 900 may include volatile and nonvolatile, removable and non-removable media, implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. The system memory 904 and storage 912 are all examples of computer-readable storage media. Non-transitory computer-readable storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile discs (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computing device 900. Any such non-transitory computer-readable storage media may be part of the computing device 900.

[0084] In various embodiments, any or all of the system memory 904 and storage 912 may store programming instructions which, when executed, implement some or all of the function functionality described above. Any or all of the system memory 904 and storage 912 may also be used for storage of distributed ledger data, including blockchains.

[0085] The computing device 900 may also have input device(s) 914 such as a keyboard, a mouse, a touch-sensitive display, voice input device, etc. Output device(s) 916 such as a display, speakers, a printer, etc. may also be included. The computing device 900 may also contain communication connections 918 that allow the device to communicate with other computing devices.

[0086] Although features and/or methodological acts are described above, it is to be understood that the appended claims are not necessarily limited to those features or acts. Rather, the features and acts described above are disclosed as example forms of implementing the claims.

CLAIMS

What is claimed is:

1. A method for monitoring integrity of a network slice, comprising:
receiving first state data from components of the network slice, wherein the first
5 state data corresponds to a first time;
calculating a first signature of the first state data;
receiving second state data from the components, wherein the second state data
corresponds to a second time;
calculating a second signature of the second state data;
10 determining that the second signature is different than the first signature; and
in response to detecting that the second signature is different than the first
signature, initiating an integrity threat response.
2. The method of claim 1, wherein the first and second state data comprise
component signatures received from the components of the network slice.
- 15 3. The method of claim 1, wherein the first and second state data comprise
properties of the components.
4. The method of claim 1, further comprising:
calculating a chronological sequence of additional signatures based on further
state data received over time from the components of the network slice; and
20 storing the chronological sequence of additional signatures, wherein the
chronological sequence of additional signatures is available as historical evidence of
the integrity of the network slice.
5. The method of claim 4, further comprising storing the chronological
sequence of additional signatures in a distributed ledger, wherein the distributed ledger
25 uses consensus-based data synchronization between multiple distributed data stores.

6. The method of claim 4, further comprising storing the chronological sequence of additional signatures in a distributed ledger, wherein:

the distributed ledger uses consensus-based data synchronization between multiple distributed data stores; and

5 the additional signatures of the sequence are stored as blocks of a blockchain in the distributed data stores.

7. The method of claim 1, wherein:

calculating the first signature comprises calculating a first cryptographic hash of the first state data; and

10 calculating the second signature comprises calculating a second cryptographic hash of the second state data.

8. The method of claim 1, wherein initiating the threat response comprises generating a notification.

9. A method, comprising:

15 receiving component signatures from components of a network slice, wherein each of the component signatures corresponds to a respective component of the network slice and is generated based on a state of the component at a corresponding time;

calculating a chronological sequence of slice signatures based on based on the component signatures; and

20 storing the chronological sequence of slice signatures as a blockchain, in a distributed blockchain ledger, as a historical record of network slice integrity.

10. The method of claim 9, further comprising:

receiving a customer request for verification of the network slice integrity; and

in response to the customer request, providing the blockchain as verification of

25 the network slice integrity.

11. The method of claim 9, wherein the blockchain comprises multiple linked data blocks respectively containing the slice signatures of the chronological sequence.

12. The method of claim 9, wherein the distributed blockchain ledger uses consensus-based data synchronization between multiple distributed data stores.

13. The method of claim 9, wherein the distributed blockchain ledger comprises a public blockchain ledger that can be accessed by a customer to verify the
5 network slice integrity.

14. The method of claim 9, wherein the components comprise one or more of:

a router;
a switch;
10 a gateway;
an Internet-of-things device;
a virtual network function;
a container;
a software application;
15 a software module;
a virtual machine;
a radio;
a radio access network;
an Internet Multimedia Subsystem (IMSs).

20 15. One or more non-transitory computer-readable media storing computer-executable instructions that, when executed by one or more processors, cause the one or more processors to perform actions comprising:

repeatedly receiving state data from components of a network subsystem, wherein the state data is based at least in part on one or more properties of each
25 component;

calculating a chronological sequence of signatures based on the state data, wherein each of the signatures is based on a set of the state data; and

storing the chronological sequence of signatures as a historical record of network subsystem integrity.

16. The one or more non-transitory computer-readable memory of claim 15,
the actions further comprising:

receiving a customer request for verification of the network subsystem integrity;

and

5 in response to the customer request, providing the chronological sequence of
signatures as verification of the network subsystem integrity.

17. The one or more non-transitory computer-readable memory of claim 15,
further comprising storing the chronological sequence of signatures in a distributed
ledger.

10 18. The one or more non-transitory computer-readable memory of claim 17,
wherein the distributed ledger uses consensus-based data synchronization between
multiple distributed data stores.

19. The one or more non-transitory computer-readable memory of claim 17,
wherein the distributed ledger is a distributed blockchain ledger.

15 20. The one or more non-transitory computer-readable memory of claim 15,
wherein the one or more properties include one or more of:

content of a software image;

content of a firmware image;

content of a file;

20 stored data;

a setting or configuration parameter;

whether an application or other software component is in memory and/or
executing;

a permission;

25 whether a patch or update has been applied;

content of a code image of a patch or update;

whether a file is present;

whether a file is absent; or

whether a resource is available or sufficient.

30




FIG. 1




FIG. 3




FIG. 2




FIG. 5




FIG. 4




FIG. 8




FIG. 7





FIG. 6

**FIG. 9**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2019/044301**A. CLASSIFICATION OF SUBJECT MATTER****H04L 29/06(2006.01)i, H04L 9/32(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; G06F 11/10; G06K 009/00; H04B 1/00; H04B 17/30; H04L 29/08; H04L 9/32Documentation 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: network slice, integrity, state data, time, signature, different**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2015-0349810 A1 (BASTILLE NETWORKS, INC.) 03 December 2015 See paragraphs [0252], [0263]–[0264]; claims 1, 3; and figure 13.	1–20
A	EP 0963637 B1 (KONINKLIJKE PHILIPS ELECTRONICS N.V.) 18 July 2007 See paragraph [0027]; claim 1; and figure 5.	1–20
A	US 2015-0261609 A1 (CLEVERSAFE, INC.) 17 September 2015 See paragraphs [0100]–[0106]; and figure 7.	1–20
A	US 2003-0138135 A1 (KEVIN KWONG-TAI CHUNG et al.) 24 July 2003 See claims 1–8.	1–20
A	VINOD KUMAR CHOYI et al., 'Network slice selection, assignment and routing within 5G Networks', In: 2016 IEEE Conference on Standards for Communications and Networking (CSCN), October 2016 See section III.	1–20



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

"D" document cited by the applicant in the international application

"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

11 November 2019 (11.11.2019)

Date of mailing of the international search report

11 November 2019 (11.11.2019)

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

KIM, Sung Hee

Telephone No. +82-42-481-5659



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2019/044301

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015-0349810 A1	03/12/2015	EP 3149597 A1	05/04/2017
		EP 3149986 A1	05/04/2017
		EP 3213105 A1	06/09/2017
		EP 3276527 A1	31/01/2018
		EP 3285196 A1	21/02/2018
		EP 3287934 A1	28/02/2018
		EP 3296916 A1	21/03/2018
		EP 3296917 A2	21/03/2018
		EP 3327609 A1	30/05/2018
		US 10104098 B2	16/10/2018
		US 10122736 B2	06/11/2018
		US 10338191 B2	02/07/2019
		US 2015-0348341 A1	03/12/2015
		US 2015-0348342 A1	03/12/2015
		US 2015-0350228 A1	03/12/2015
		US 9485266 B2	01/11/2016
		US 9485267 B2	01/11/2016
		US 9551781 B2	24/01/2017
		US 9739868 B2	22/08/2017
		US 9880256 B2	30/01/2018
		US 9945928 B2	17/04/2018
		WO 2015-187726 A1	10/12/2015
		WO 2015-187730 A1	10/12/2015
		WO 2016-070153 A1	06/05/2016
EP 0963637 B1	18/07/2007	CA 2282479 A1	15/07/1999
		CN 1149784 C	12/05/2004
		CN 1254464 A	24/05/2000
		DE 69838094 T2	03/04/2008
		EP 0963637 A2	15/12/1999
		JP 2001-515612 A	18/09/2001
		JP 2010-187419 A	26/08/2010
		KR 10-2000-0075866 A	26/12/2000
		US 6601172 B1	29/07/2003
		WO 99-35785 A2	15/07/1999
		WO 99-35785 A3	16/09/1999
US 2015-0261609 A1	17/09/2015	US 10154034 B2	11/12/2018
		US 2011-0264717 A1	27/10/2011
		US 2011-0264989 A1	27/10/2011
		US 2011-0265143 A1	27/10/2011
		US 2013-0246865 A1	19/09/2013
		US 2013-0262959 A1	03/10/2013
		US 2014-0325259 A1	30/10/2014
		US 2014-0351891 A1	27/11/2014
		US 2014-0359349 A1	04/12/2014
		US 2019-0026044 A1	24/01/2019
		US 8566354 B2	22/10/2013
		US 8769379 B2	01/07/2014

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2019/044301

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2003-0138135 A1	24/07/2003	US 9043689 B2	26/05/2015
		US 9047218 B2	02/06/2015
		US 9063881 B2	23/06/2015
		US 9311186 B2	12/04/2016
		US 9342406 B2	17/05/2016
		US 9606858 B2	28/03/2017
		AU 2001-297675 A1	19/09/2002
		AU 2001-97012 A1	22/04/2002
		CA 2456098 A1	13/02/2003
		CA 2469146 A1	12/09/2002
		EP 1344166 A2	17/09/2003
		EP 1344166 A4	12/05/2004
		EP 1421459 A2	26/05/2004
		EP 1514247 A2	16/03/2005
		EP 1514247 A4	22/07/2009
		EP 1514247 B1	16/01/2013
		HK 1073167 A1	19/09/2013
		US 2002-0044096 A1	18/04/2002
		US 2002-0077886 A1	20/06/2002
		US 2003-0006878 A1	09/01/2003
		US 2003-0026462 A1	06/02/2003
		US 2004-0156537 A1	12/08/2004
		US 2006-0255145 A1	16/11/2006
		US 2007-0136154 A1	14/06/2007
		US 6657543 B1	02/12/2003
		US 6694045 B2	17/02/2004
		US 7036729 B2	02/05/2006
		US 7036730 B2	02/05/2006
		US 7077313 B2	18/07/2006
		US 7431209 B2	07/10/2008
		US 7436989 B2	14/10/2008
		US 7461787 B2	09/12/2008
		US 7501954 B1	10/03/2009
		US 7508308 B2	24/03/2009
		US 7513425 B2	07/04/2009
		US 7894634 B2	22/02/2011
		US 7975920 B2	12/07/2011
		US 7988047 B2	02/08/2011
		WO 02-031629 A3	13/02/2003
		WO 02-070998 A2	12/09/2002
		WO 02-070998 A3	17/07/2003
		WO 02-31629 A2	18/04/2002
		WO 02-33511 A2	25/04/2002
		WO 03-012595 A2	13/02/2003
		WO 03-012595 A3	18/09/2003
		WO 03-062961 A2	31/07/2003