




US 20200382560A1

(19) **United States**(12) **Patent Application Publication**
Woolward et al.(10) **Pub. No.: US 2020/0382560 A1**(43) **Pub. Date: Dec. 3, 2020**(54) **VALIDATION OF CLOUD SECURITY
POLICIES**(71) Applicant: **vArmour Networks, Inc.**, Mountain
View, CA (US)(72) Inventors: **Marc Woolward**, Bude (GB); **Meng
Xu**, Los Altos, CA (US); **Hong Xiao**,
San Jose, CA (US); **Keith Stewart**, San
Francisco, CA (US); **Matthew M.
Williamson**, Marblehead, MA (US)(21) Appl. No.: **16/428,849**(22) Filed: **May 31, 2019****Publication Classification**(51) **Int. Cl.**
H04L 29/06 (2006.01)
H04L 29/08 (2006.01)
G06F 16/901 (2006.01)(52) **U.S. Cl.**CPC **H04L 63/205** (2013.01); **G06F 16/9024**
(2019.01); **H04L 67/1097** (2013.01); **H04L**
63/0263 (2013.01)

(57)

ABSTRACT

Methods and systems for validating security policy in a cloud computing environment are provided. An example method includes providing a graph database, the graph database representing workloads of the cloud computing environment as nodes and relationships between the workloads as edges, receiving a security policy, the security policy logically describing rules for the relationships between the workloads, determining, based on the security policy and the graph database, a list of violations, the list of violations including at least one relationship from the relationships between the workloads in the graph database, the at least one relationship being not allowed by at least one of the rules in the security policy, and providing the list of violations to a user.



100 ↘




FIG. 1




FIG. 2




FIG. 3





FIG. 4A



500 ↘




FIG. 5





FIG. 6

**FIG. 7**




FIG. 8

VALIDATION OF CLOUD SECURITY POLICIES

CROSS-REFERENCE TO RELATED APPLICATION

[0001] The present application is related to U.S. patent application Ser. No. _____ (Attorney Docket No. PA9273US) filed _____, titled “Cloud Security Management”. The subject matter of the aforementioned application is incorporated herein by reference for all purposes.

FIELD OF THE INVENTION

[0002] The present technology pertains to communications networks, and more specifically to security in cloud computing environments.

BACKGROUND ART

[0003] The approaches described in this section could be pursued but are not necessarily approaches that have previously been conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.

[0004] Due to the extensive use of computer networks by enterprises, there has been a dramatic rise in network attacks, a proliferation of computer viruses, and a constant distribution of other types of malicious content that attempts to attack, infect, or otherwise infiltrate the computer networks. Attackers breach internal networks and public clouds to steal critical data. For example, attackers target low-profile assets to enter the internal network. Inside the internal network and public clouds, and behind the hardware firewall, attackers move laterally across the internal network, exploiting East-West traffic flows, to critical enterprise assets. Once there, attackers siphon off valuable company and customer data.

SUMMARY OF THE INVENTION

[0005] This summary is provided to introduce a selection of concepts in a simplified form that are further described in the Detailed Description below. This summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

[0006] The present disclosure is related to various methods and systems for validating security policy in a cloud computing environment. Specifically, a method for validating security policy in a cloud computing environment may include providing a graph database. The graph database may represent workloads of the cloud computing environment as nodes and relationships between the workloads as edges. The method may include receiving a security policy. The security policy may logically describe rules for the relationships between the workloads. The method may include determining, based on the security policy and the graph database, a list of violations. The list of violations may include at least one relationship from the relationships between the workloads in the graph database such that the relationship is not allowed by at least one of the rules in the security policy. The method may further include providing the list of violations to a user.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Embodiments are illustrated by way of example, and not by limitation, in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

[0008] FIG. 1 is a simplified block diagram of a cloud computing environment, according to some embodiments.

[0009] FIG. 2 is a simplified block diagram of a system for cloud security management, according to various embodiments.

[0010] FIG. 3 depicts simplified graph of a cloud computing environment, in accordance with some embodiments.

[0011] FIG. 4A shows another graph of a cloud computing environment and FIG. 4B depicts a graph of an application, in accordance with various embodiments.

[0012] FIG. 5 is a simplified flow diagram of a method for cloud security management, according to some embodiments.

[0013] FIG. 6 is a simplified block diagrams showing functionality of a protect in a cloud security management, according to some embodiments.

[0014] FIG. 7 is a simplified flow diagram of a method for validating security policy in a cloud computing environment, according to some embodiments.

[0015] FIG. 8 is a simplified block diagram of a computing system, according to various embodiments.

DETAILED DESCRIPTION

[0016] While this technology is susceptible of embodiment in many different forms, there is shown in the drawings and will herein be described in detail several specific embodiments with the understanding that the present disclosure is to be considered as an exemplification of the principles of the technology and is not intended to limit the technology to the embodiments illustrated. The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the technology. As used herein, the singular forms “a,” “an,” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “comprises,” “comprising,” “includes,” and/or “including,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof. It will be understood that like or analogous elements and/or components, referred to herein, may be identified throughout the drawings with like reference characters. It will be further understood that several of the figures are merely schematic representations of the present technology. As such, some of the components may have been distorted from their actual scale for pictorial clarity.

[0017] FIG. 1 shows cloud computing environment 100 including workloads 110_{1,1}-110_{X,Y}, according to some embodiments. Cloud computing environment 110 provides on-demand availability of computer system resources, such as data storage and computing power. Cloud computing environment 110 can physically reside in one or more data centers and/or be physically distributed over multiple locations. Cloud computing environment 110 can be hosted by more than one cloud service, such as those provided by Amazon, Microsoft, and Google. Cloud computing environ-

ment **110** can be limited to a single organization (referred to as an enterprise cloud), available to many organizations (referred to as a public cloud) or a combination of both (referred to as a hybrid cloud). Examples of public clouds include Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP).

[0018] Each of workloads **110_{1,1}-110_{x,y}** can be a unit of computing resource, such as a physical computing system (also referred to as a bare metal server), virtual machine, container, pod, and combinations thereof. A physical computing system is computer hardware and not a virtual computing system, such as a virtual machine and container. In addition to running operating systems and applications, physical computing systems can be the hardware that virtual computing systems run on.

[0019] A virtual machine provides a substitute for a physical computing system, including functionality to execute entire operating systems. Virtual machines are created and run by a hypervisor or virtual machine monitor (VMM). A hypervisor is computer software or firmware which can run on workloads **110_{1,1}-110_{x,y}**. A hypervisor uses native execution to share and manage hardware, allowing for multiple environments which are isolated from one another, yet exist on the same physical computing system.

[0020] Containers are an operating system-level virtualization method for deploying and running distributed applications without launching an entire virtual machine for each application. Containers can look like physical computing systems from the point of view of programs running in them. Generally, a computer program running on an operating system can see all resources (e.g., connected devices, files and folders, network shares, CPU power, etc.) of that physical computing system. However, programs running inside a container can only see the container's contents and devices assigned to the container. A pod is a group of containers with shared storage and/or network resources, and a shared specification for how to run the containers.

[0021] A container is an instance of an image. An image can be a file, comprised of multiple layers, with information to create a complete and executable version of an application. Containers can be arranged, coordinated, and managed—including means of discovery and communications between containers—by container orchestration (e.g., Docker Swarm®, Kubernetes®, Amazon EC2 Container Service (ECS), Diego, Red Hat OpenShift, and Apache® Mesos™). In contrast to hypervisor-based virtualization, containers may be an abstraction performed at the operating system (OS) level, whereas virtual machines are an abstraction of physical hardware.

[0022] Typically, workloads **110_{1,1}-110_{x,y}** of cloud computing environment **100** individually and/or collectively run applications and/or services. Applications and/or services are programs designed to carry out operations for a specific purpose. By way of non-limiting example, applications can be a database (e.g., Microsoft SQL Server®, MongoDB, Hadoop Distributed File System (HDFS), etc.), email server (e.g., Sendmail®, Postfix, qmail, Microsoft® Exchange Server, etc.), message queue (e.g., Apache® Qpid™, RabbitMQ®, etc.), web server (e.g., Apache® HTTP Server™, Microsoft® Internet Information Services (IIS), Nginx, etc.), Session Initiation Protocol (SIP) server (e.g., Kamailio® SIP Server, Avaya® Aura® Application Server 5300, etc.), other media server (e.g., video and/or audio streaming, live broadcast, etc.), file server (e.g., Linux server, Micro-

soft® Windows Server®, etc.), service-oriented architecture (SOA) and/or microservices process, object-based storage (e.g., Lustre®, EMC® Centera®, Scality® RING®, etc.), directory service (e.g., Microsoft® Active Directory®, Domain Name System (DNS) hosting service, etc.), and the like.

[0023] Physical computing systems and cloud computing environments are described further in relation to FIG. 6.

[0024] FIG. 2 shows system **200** for cloud security management, according to some embodiments. System **200** includes controller **210**. Controller **210** can receive streaming telemetry **275** from network logs **270**, events **285** from cloud control plane **280**, and inventory **295** from configuration management database (CMDB) **290**.

[0025] Network logs **270** can be data sources such as flow logs from cloud services **260₁-260_z** (e.g., Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)), vArmour DSS Distributed Security System, Software Defined Networking (SDN) (e.g., VMware NSX and Cisco Application Centric Infrastructure (ACI)), monitoring agents (e.g., Tanium Asset and Falco) and the like. Generally, streaming telemetry **275** can be low-level data about relationships between applications. Streaming telemetry **275** can include 5-tuple, layer 7 (application layer) process information, management plane logs, and the like. 5-tuple refers to a set of five different values that comprise a Transmission Control Protocol/Internet Protocol (TCP/IP) connection: a source IP address/port number, destination IP address/port number and the protocol in use. Streaming telemetry can alternatively or additionally include a volume of data (i.e., how much data is or how many data packets are) exchanged between workloads (e.g., workloads **110_{1,1}-110_{x,y}** in FIG. 1) in a network, (dates and) times at which communications (e.g., data packets) are exchanged between workloads, and the like.

[0026] Cloud control plane **280** establishes and controls the network and computing resources within a cloud computing environment (e.g., cloud computing environment **100** in FIG. 1). Cloud control plane **280** can include interfaces for managing assets (e.g., launching virtual machines and/or containers, configuring the network, etc.) in a cloud computing environment. For example, cloud control plane **280** can include one or more instances of container orchestration, such as Docker Swarm®, Kubernetes®, Amazon EC2 Container Service (ECS), Diego, and Apache® Mesos™. By way of further non-limiting example, cloud control plane **280** can include VMware vSphere, application programming interfaces (APIs) provided by cloud services **260₁-260_z**, and the like.

[0027] Events **285** can include information about a container (and/or a pod) being created, having a state change, having an error, and the like. For example, when a container is created, information about the workload such as a service name, image deployed, and the like can be received in events **285**. By way of further example, additional information from an image registry corresponding to the deployed image can be gathered by controller **210**.

[0028] Configuration management database (CMDB) **290** can be a database of information about the hardware and software components (also known as assets) used in a cloud computing environment (e.g., cloud computing environment **100** in FIG. 1) and relationships between those components and business functions. CMDB **290** can include information about upstream sources or dependencies of components, and

the downstream targets of components. For example, inventory 295 can be used to associate an application name and other information (e.g., regulatory requirements, business unit ownership, business criticality, and the like) with the workload (e.g., workloads 110_{1,1}-110_{X,Y} in FIG. 1) it is running on.

[0029] Streaming telemetry 275, events 285, and inventory 295 can be ingested by graph 220. Graph 220 normalizes information received in streaming telemetry 275, events 285, and inventory 295 into a standard data format and/or model, graph database 225. Graph database 225 uses a graph data model comprised of nodes (also referred to as vertices), which is an entity such as a workload (e.g., of workloads 110_{1,1}-110_{X,Y} in FIG. 1), and edges, which represent the relationship between two nodes. Edges can be referred to as relationships. An edge can have a start node, end node, type, and direction, and an edge can describe parent-child relationships, actions, ownership, and the like. In contrast to relational databases, relationships are (most) important in graph database 225. In other words, connected data is equally (or more) important than individual data points.

[0030] Conventionally, security management systems stored raw logs of each and every individual communication between workloads. The amount of data scaled linearly and consumed massive amounts of storage. In contrast, streaming telemetry 275, events 285, and inventory 295, graph 220 (FIG. 2) can be used by graph 220 to create and update graph (database) 300. The individual communications are not stored. In this way, graph database 225 is advantageously scalable. For example, graph database 225 for a large cloud computing environments of 30,000-50,000 workloads can be stored in memory of a workload (e.g., of workloads 110_{1,1}-110_{X,Y} in FIG. 1).

[0031] FIG. 3 depicts (simplified) graph (database) 300 of a cloud computing environment, according to various embodiments. Graph 300 is a simplified example, purely for illustrative purposes, of a graph in graph database 225 (FIG. 2). Graph 300 can include three workloads (e.g., of workloads 110_{1,1}-110_{X,Y} in FIG. 1): node 310, node 330, and node 350. As shown in FIG. 3, edge (relationship) 320 is between nodes 310 and 330 have; edge (relationship) 340 is between nodes 330 and 350; edge (relationship) 360 is between nodes 350 and 310.

[0032] Using streaming telemetry 275, events 285, and inventory 295, graph 220 (FIG. 2) can determine information 335 about node 330. By way of non-limiting example, information 335 can include an application name, application function, business organization (e.g., division within a company), realm (e.g., production system, development system, and the like), (geographic) location/zone, and other metadata. Moreover, using layer 7 information (when available), the name of the database can be determined.

[0033] Referring back to FIG. 2, graph 220 can employ various techniques to manage entropy. In a cloud computing environment (e.g., cloud computing environment 100 in FIG. 1), entropy is change to the workloads (e.g., created and removed), communications among workloads (e.g., which workloads communicate with other workloads), applications and services provided in the network, and the like. Typically in a (closed) enterprise cloud, entropy is low. For example, after monitoring an enterprise cloud for one month, another month of monitoring will reveal little that is new.

[0034] On the other hand, a web server connected to the Internet will have high entropy, because the number of relationships (connections) to clients on the Internet (nodes) is huge and continues to grow. To protect the size of graph database 225, graph 220 can recognize when there is high entropy and summarize the nodes. For example, the vast (and growing) number of clients on the Internet is represented by a single “Internet” object with one edge to the web server node.

[0035] According to some embodiments, a new relationship can be created around a particular node in graph database 225, as streaming telemetry 275, events 285, and inventory 295 are processed by graph 220. Graph 220 (FIG. 2) can further re-analyze the edges (relationships) connected to the particular node, to classify what the particular node is. For example, if the node accepts database client connections from systems that are known to be application servers, then graph 220 may classify the node as a database management system (i.e., a certain group). Classification criteria can include heuristic rules. Graph 220 can use machine learning algorithms and measure how close a particular node is to satisfying conditions for membership in a group. Classification is described further in U.S. Pat. No. 10,264,025 issued Apr. 16, 2019, titled “Security Policy Generation for Virtualization, Bare-Metal Server, and Cloud Computing Environments,” which is hereby incorporated by reference for disclosure of classification.

[0036] Visualize 230 can visually present information from graph database 225 to users according to various criteria, such as by application, application type, organization, and the like. FIGS. 4A and 4B show example visual presentations 400A and 400B, respectively, in accordance with some embodiments.

[0037] Visualize 230 can visually organize information from graph database 225. In some embodiments, nodes that behave similarly can be clustered together (i.e., be put in a cluster). For example, when two nodes have similar edges (relationships) and behave in a similar fashion (e.g., run the same application, are associated with the same organization, and the like), the two nodes can be clustered together. Nodes that are clustered together can be visually presented as a shape (e.g., circle, rectangle, and the like) which denotes that there are a certain number of workloads fulfilling the same function, instead of presenting a shape for each workload in the cluster.

[0038] In various embodiments, visualize 230 can detect and present communities. Communities are workloads (e.g., of workloads 110_{1,1}-110_{X,Y} in FIG. 1) that have a close set of edges (relationships). The constituent workloads of a community do not have to be the same—they can each perform different functions, such as web server, database server, application server, and the like—but the workloads are densely connected. In other words, the nodes communicate with each other often and in high volume. Workloads in a community act collectively to perform an application, service, and/or business function. Instead of displaying a shape (e.g., circle, rectangle, and the like) for each of the hundreds or thousands of workloads in a community, the community can be represented by a single shape denoting the application performed, the number of constituent workloads, and the like.

[0039] Protect 240 can use information in the graph database 225 to design security policies. Security policies can implement security controls, for example, to protect an

application wherever it is in a cloud computing environment (e.g., cloud computing environment **100** in FIG. 1). A security policy can specify what is to be protected (“nouns”), for example, applications run for a particular organization. A security policy can further specify a security intent (“verbs”), that is, how to protect. For example, a security intent can be to implement Payment Card Industry Data Security Standard (PCI DSS) network segmentation requirements (a regulatory requirement), implement a security best practices for databases, implement a whitelist architecture, and the like. By way of further example, a security intent can be specified in a template by a user (responsible for system administration, security, and the like).

[0040] Nouns and verbs can be described in a security template. A security template can include logic about how to process information in graph database **225** relating to workloads having a particular label/selection (nouns). Labels can be provided by logs **270** (e.g., layer 7 information), cloud control planes **280** (e.g., container orchestration), and CMDB **290**. Protect **240** uses a security template to extract workloads to be protected (nouns) from graph database **225**. Protect **240** further applies logic in the security template about how to protect the workloads (verbs) to produce a security policy. In various embodiments, security templates are JavaScript Object Notation (JSON) documents, documents in Jinja (or Jinja2), YAML Ain’t Markup Language (YAML) documents, Open Policy Agent (OPA) rules, and the like. Jinja and Jinja2 are a web template engine for the Python programming language. YAML is a human-readable data-serialization language. OPA is an open source, general-purpose policy engine that enables unified, context-aware policy enforcement. Security templates are described further in U.S. patent application Ser. No. _____ (Attorney Docket No. PA9274US) filed _____, titled “Template-Driven Intent-Based Security,” which is hereby incorporated by reference for disclosure of generating a security policy using security templates.

[0041] Protect **240** can produce multiple security policies, each reflecting independent pieces of security logic that can be implemented by protect **240**. In various embodiments, security policies are JavaScript Object Notation (JSON) documents which are described to a user (responsible for system administration, security, and the like) in natural language. A natural language is any language that has evolved naturally in humans through use and repetition without conscious planning or premeditation. Natural language can broadly be defined in contrast to artificial or constructed languages such as computer programming languages. The multiple security policies can be placed in an order of precedence to resolve potential conflicts. Visualize **230** can be used to visualize the security policy (or security policies), showing the workloads protected, permitted relationships, and prohibited relationships. Protect **240** can then be used to edit the security policy. For example, there can be a primary and backup server (e.g., of workloads **110_{1,1}-110_{x,y}** in FIG. 1). The backup server may have never been used and may not have the same edges (relationships) as the primary server in graph database **225**. The security policy can be edited to give the backup server the same permissions as the primary server.

[0042] Protect **240** can validate a security policy. The security policy can be simulated using graph database **225**. For example, a simulation can report which applications are broken (e.g., communications among nodes needed by the

application to operate are prohibited) by the security policy, are unnecessarily exposed by weak policy, and the like. Security policy validation is described further in U.S. patent application Ser. No. _____ (Attorney Docket No. PA9275US) filed _____, titled “Validation of Cloud Security Policies,” which is incorporated by reference herein for disclosure of security policy validation.

[0043] Protect **240** can test a security policy. Protect can use historical data in graph database **225** to determine entropy in the cloud computing environment (e.g., cloud computing environment **100** in FIG. 1). For example, when a cloud computing environment first starts up, there are initially numerous changes as workloads are brought online and communicate with each other, such that entropy is high. Over time, the cloud computing environment becomes relatively stable with few changes, so entropy becomes low. In general, security policies are less reliable when entropy is high. Protect **240** can determine a level of entropy in the cloud computing environment and produce a reliability score and recommendation for the security policy. Security policy testing is described further in U.S. patent application Ser. No. _____ (Attorney Docket No. PA9276US) filed _____, titled “Reliability Prediction for Cloud Security Policies,” which is incorporated by reference herein for disclosure of security policy reliability prediction.

[0044] Protect **240** can deploy a security policy (or security policies). The security policy is deployed as needed in one or more cloud computing environments of cloud services **260₁-260_z** (e.g., Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)), vArmour DSS Distributed Security System, VMware NSX, and the like). Protect **240** can provide the security policy to one or more of cloud drivers **250₁-250_z**. Cloud drivers **250₁-250_z** maintain an inventory and topology (i.e., current state) of the workloads in the cloud computing environments hosted by cloud services **260₁-260_z**, respectively. Cloud drivers **250₁-250_z** can use their respective inventory and topology to apply the security policy to the appropriate workloads, and respond immediately to changes in workload topology and workload placement.

[0045] Cloud drivers **250₁-250_z** can serve as an interface between protect **240** (having a centralized security policy) and cloud services **260₁-260_z**. In other words, cloud drivers **250₁-250_z** implement the security policy using the different facilities (e.g., application programming interfaces (APIs)) and capabilities available from cloud services **260₁-260_z**. For example, each of cloud services **260₁-260_z** can have different syntax and semantics for implementing security controls. Moreover, each of cloud services **260₁-260_z** can have different security capabilities (e.g., communications/connections between workloads can only be expressly permitted and not expressly prohibited), rule capacity (limit on the number of rules), optimization methods, and the like.

[0046] Cloud drivers **250₁-250_z** can maintain the integrity of the security policy in the cloud computing environments hosted by cloud services **260₁-260_z** (referred to as the “cloud”). Cloud drivers **250₁-250_z** can check that the security policy actually deployed in the cloud is as it should be, using the security policy’s JSON source. When the security policy deployed in the cloud does not comport with the centralized security policy—such as when a bad actor logs into one of the cloud services and removes all the security rules—the responsible cloud driver (of cloud drivers **250₁-250_z**) can re-deploy the security policy and/or raise an

operational alert. Where supported, cloud services 260₁-260_z can notify the respective cloud driver (of cloud drivers 250₁-250_z) of changes to the topology and/or configuration. Otherwise, the respective cloud driver (of cloud drivers 250₁-250_z) can poll the cloud service (cloud services 260₁-260_z) to ensure the security rules are in place.

[0047] As described above, a security policy can be pushed down to the cloud computing environments hosted by cloud services 260₁-260_z using cloud drivers 250₁-250_z, respectively. Additionally or alternatively, as new data comes into graph 220 as network logs 270, events 285 from cloud control plane 280, and inventory 295, protect 240 can check the new data against the security policy to detect violations and or drift (e.g., change in the environment and/or configuration).

[0048] Protect 240 can dynamically update a security policy as changes occur in the cloud computing environments hosted by cloud services 260₁-260_z. For example, when a container (or pod) is deployed by container orchestration, it can be given a label, and cloud control plane 290 reports a container is deployed (as event 295). Labels can be predefined to specify identifying attributes of containers (and pods), such the container's application function. When the label corresponds to an attribute covered by an active (deployed) security policy, protect 240 can dynamically add the new container to the active security policy (as a target). For example, when a pod is deployed for a particular organization and there is an active policy for that organization, the new workload is added to the security policy. Similarly, when a container is killed, the workload is removed from the security policy. Dynamically updating security policy is described further in U.S. Pat. No. 9,521, 115 issued Dec. 13, 2016, titled "Security Policy Generation Using Container Metadata," which is hereby incorporated by reference for disclosure of dynamically updating security policy.

[0049] FIG. 5 shows method 500 for managing cloud security, according to some embodiments. Method 500 can be performed by system 200 (FIG. 2), including controller 210. Method 500 can commence at step 510 where data from a cloud computing environment (e.g., cloud computing environment 100 in FIG. 1) can be received. For example, graph 220 (FIG. 2) can receive streaming telemetry 275 from network logs 270, events 285 from cloud control plane 280, and inventory 295 from configuration management database (CMDB) 290.

[0050] At step 520, a graph database can be created or updated using the cloud data. For example, streaming telemetry 275, events 285, and inventory 295 (FIG. 2) can be normalized into a standard data format and stored in graph database 225.

[0051] At step 530, a visual representation of the cloud computing environment as modeled by the graph database can be provided. For example, visualize 230 (FIG. 2) can present a graph using data in graph database 225. In some embodiments, nodes (representing workloads in the cloud computing environment) can be clustered and/or placed in communities for visual clarity.

[0052] At step 540, a security template can be received. A security template can include logic about how to extract information from graph database 225 to identify workloads to be targets of a security policy. In addition, a security template can specify how the workloads are to be protected (e.g., security intent).

[0053] At step 550, a security policy can be created. For example, protect 240 can use the security template to extract information from graph database 225 (FIG. 2) to produce a security policy for the security intent of the security template.

[0054] At step 560, the security policy can be validated. For example, protect 240 (FIG. 2) test the security policy against a historical data set stored in graph database 225. Protect 240 can generate a report around the risks and implications of the security policy being implemented.

[0055] At step 570, the security policy can be tested. For example, protect 240 (FIG. 2) can measure entropy and a rate of change in the data set stored graph database 225 to predict—when the security policy is deployed—the cloud computing environment (e.g., cloud computing environment 100 in FIG. 1) will change such that applications and/or services will break (e.g., be prevented from proper operation by the security policy).

[0056] At step 580, the security policy can be deployed to the cloud computing environment (e.g., cloud computing environment 100 in FIG. 1). For example, cloud drivers 250₁-250_z can produce requests, instructions, commands, and the like which are suitable for and accepted by cloud services 260₁-260_z (respectively) to implement the security policy in the cloud computing environments hosted by cloud services 260₁-260_z (respectively).

[0057] Optionally at step 580, the security policy can be maintained. For example, cloud drivers 250₁-250_z can make sure the security policy remains in force at the cloud computing environment hosted by a respective one of cloud services 260₁-260_z. Optionally at step 580, the security policy can be dynamically updated as workloads subject to the deployed security policy are deployed and/or killed.

[0058] Although steps 510-580 are shown in a particular sequential order, various embodiments can perform steps 510-580 in different orders, perform some of steps 510-580 concurrently, and/or omit some of steps 510-580.

[0059] FIG. 6 is a simplified block diagram 600 showing functionality of protect 240 in cloud security management 200, according to some embodiments. The cloud security management 200 and the protect 240 is described in FIG. 2.

[0060] The protect 240 may receive a security policy 610. The security policy may include one or more rules for relationship between workloads 110_{1,1}-110_{x,r} of cloud computing environment 100. In some embodiments, the security policy may include JSON documents. The multiple security policies can be placed in order to resolve potential conflicts. In example of FIG. 6, the security policy 610 includes a rule allowing an HTTP service between a source workload A1 and a destination workload L1, a rule denying HTTPS service between source workload A2 and destination workload L2. The security policy 610 may end with a rule that denies any service between any source workload labeled with A and any destination workload labeled with L.

[0061] The security policy 630 can be generated based on a security intent of the security template. The security policy 630 can be created by users (operators) who deploy the security policy using the protect 240. The security policy 630 can be generated based on the security template and then reviewed and edited by the users.

[0062] Before deploying the security policy 610 to the cloud computing environment 100, the protect 240 may simulate deployment of the security policy 610 by applying the security policy 610 to the graph database 225. Specifi-

cally, the protect **240** may traverse through the nodes and edges of the graph database **225** and determine edges that represent services between the workloads which are not allowed by the security policy **610**. The protect **240** may determine a type of service or connection existing between two nodes representing two workloads in the graph database. The protect **240** may also determine a type of service or connection between the two nodes that was present in one or more previous time snapshot of the graph database **225**. The protect **240** may further determine whether this type of service for these two nodes (workloads) is permitted according to the security policy **610**. In example of FIG. 6, the graph database includes an HTTPS service between the source workload **A2** and the destination workload **L2** which is not permitted after the security policy **610** is deployed to the cloud computing environment **100**.

[0063] Upon completion of traversing the graph database **225**, the protect **240** may generate list of violations **640**. The list of violations **640** may include a subset of nodes and edges in the graph database. Each entry of the list of violations **640** may represent a label for a first workload and a label of a second workload and a type of service or connection between the first workload and the second workload that is not allowed by the security policy **610** but present in the graph database **225**. The list of violations **640** may be provided to the user. The entries from the list of violations **640** can be displayed using a graphical representation in a form of the nodes representing workloads and the edges representing service and connections between the workloads that are not allowed in the security policy **610**. Additionally, the entries from the list of violations **640** can be displayed in a table and in a JSON document.

[0064] The protect **240** may further receive, from the user, a user input indicating either allowance or disallowance of a type of service or connection for an entry from the list of violations. Upon receiving the user input, the protect **240** may modify the security policy **610** based on the user input. The protect **240** may further deploy the modified security policy to the cloud computing environment.

[0065] FIG. 7 is a simplified flow diagram of a method for validating security policies in a cloud computing environment, according to some embodiments. Method **700** can be performed by system **200** (FIG. 2), including controller **210**.

[0066] Method **700** can commence at step **710** with providing a graph database. The graph database may represent workloads of the cloud computing environment (e.g., cloud computing environment **100** in FIG. 1) as nodes and relationships between the workloads as edges. The graph database can be created and updated based on data concerning the cloud computing environment. The data may include streaming telemetry from network logs, events from a cloud control plane, and inventory from a configuration management database.

[0067] At step **720**, a security policy can be received. The security policy can logically describe rules for the relationships between the workloads of the cloud computing environment. The security policy may include a JSON document. The security policy can be generated based on a security template. The security template may include protected workloads in the cloud computing environment. The security policy can be created by a user.

[0068] At step **730**, a list of violations can be generated based on the security policy and the graph database. The list of violations may include at least one relationship from the

relationships between the workloads in the graph database such that this relationship is not allowed by at least one of the rules of the security policy. The determination of the list of violations may include simulating the security policy using the graph database by traversing the nodes and edges in the graph database. While traversing the graph database, the determination of the list of violations may include determining a type of connection between at least two of the nodes in the graph database and determining that the type of connection between these two nodes is not allowed by the rules of the security policy.

[0069] At step **740**, the list of violations can be provided to the user. For example, a visual representation of the nodes and edges of the graph database can be displayed to the user. In the visual representation, the edges may represent relationships between workloads not allowed by the rules of the security policy but present in the graph database.

[0070] At step **750**, a user input can be received. The user input may indicate allowance or disallowance of a relationship that is not allowed by the rules of the security policy but present the graph database.

[0071] At step **760**, the security policy can be modified based on the user input. The security policy can be then deployed to cloud computing environment.

[0072] FIG. 8 illustrates an exemplary computer system **800** that may be used to implement some embodiments of the present invention. The computer system **800** in FIG. 8 may be implemented in the contexts of the likes of computing systems, networks, servers, or combinations thereof. The computer system **800** in FIG. 8 includes one or more processor unit(s) **810** and main memory **820**. Main memory **820** stores, in part, instructions and data for execution by processor unit(s) **810**. Main memory **820** stores the executable code when in operation, in this example. The computer system **800** in FIG. 8 further includes a mass data storage **830**, portable storage device **840**, output devices **850**, user input devices **860**, a graphics display system **870**, and peripheral device(s) **880**.

[0073] The components shown in FIG. 8 are depicted as being connected via a single bus **890**. The components may be connected through one or more data transport means. Processor unit(s) **810** and main memory **820** are connected via a local microprocessor bus, and the mass data storage **830**, peripheral device(s) **880**, portable storage device **840**, and graphics display system **870** are connected via one or more input/output (I/O) buses.

[0074] Mass data storage **830**, which can be implemented with a magnetic disk drive, solid state drive, or an optical disk drive, is a non-volatile storage device for storing data and instructions for use by processor unit(s) **810**. Mass data storage **830** stores the system software for implementing embodiments of the present disclosure for purposes of loading that software into main memory **820**.

[0075] Portable storage device **840** operates in conjunction with a portable non-volatile storage medium, such as a flash drive, floppy disk, compact disk, digital video disc, or Universal Serial Bus (USB) storage device, to input and output data and code to and from the computer system **800** in FIG. 8. The system software for implementing embodiments of the present disclosure is stored on such a portable medium and input to the computer system **800** via the portable storage device **840**.

[0076] User input devices **860** can provide a portion of a user interface. User input devices **860** may include one or

more microphones, an alphanumeric keypad, such as a keyboard, for inputting alphanumeric and other information, or a pointing device, such as a mouse, a trackball, stylus, or cursor direction keys. User input devices **860** can also include a touchscreen. Additionally, the computer system **800** as shown in FIG. **8** includes output devices **850**. Suitable output devices **850** include speakers, printers, network interfaces, and monitors.

[0077] Graphics display system **870** include a liquid crystal display (LCD) or other suitable display device. Graphics display system **870** is configurable to receive textual and graphical information and processes the information for output to the display device.

[0078] Peripheral device(s) **880** may include any type of computer support device to add additional functionality to the computer system.

[0079] Some of the components provided in the computer system **800** in FIG. **8** can be those typically found in computer systems that may be suitable for use with embodiments of the present disclosure and are intended to represent a broad category of such computer components. Thus, the computer system **800** in FIG. **8** can be a personal computer (PC), hand held computer system, telephone, mobile computer system, workstation, tablet, phablet, mobile phone, server, minicomputer, mainframe computer, wearable, or any other computer system. The computer may also include different bus configurations, networked platforms, multi-processor platforms, and the like. Various operating systems may be used including UNIX, LINUX, WINDOWS, MAC OS, PALM OS, QNX ANDROID, IOS, CHROME, and other suitable operating systems.

[0080] Some of the above-described functions may be composed of instructions that are stored on storage media (e.g., computer-readable medium). The instructions may be retrieved and executed by the processor. Some examples of storage media are memory devices, tapes, disks, and the like. The instructions are operational when executed by the processor to direct the processor to operate in accord with the technology. Those skilled in the art are familiar with instructions, processor(s), and storage media.

[0081] In some embodiments, the computing system **800** may be implemented as a cloud-based computing environment, such as a virtual machine operating within a computing cloud. In other embodiments, the computing system **800** may itself include a cloud-based computing environment, where the functionalities of the computing system **800** are executed in a distributed fashion. Thus, the computing system **800**, when configured as a computing cloud, may include pluralities of computing devices in various forms, as will be described in greater detail below.

[0082] In general, a cloud-based computing environment is a resource that typically combines the computational power of a large grouping of processors (such as within web servers) and/or that combines the storage capacity of a large grouping of computer memories or storage devices. Systems that provide cloud-based resources may be utilized exclusively by their owners or such systems may be accessible to outside users who deploy applications within the computing infrastructure to obtain the benefit of large computational or storage resources.

[0083] The cloud is formed, for example, by a network of web servers that comprise a plurality of computing devices, such as the computing system **800**, with each server (or at least a plurality thereof) providing processor and/or storage

resources. These servers manage workloads provided by multiple users (e.g., cloud resource customers or other users). Typically, each user places workload demands upon the cloud that vary in real-time, sometimes dramatically. The nature and extent of these variations typically depends on the type of business associated with the user.

[0084] It is noteworthy that any hardware platform suitable for performing the processing described herein is suitable for use with the technology. The terms “computer-readable storage medium” and “computer-readable storage media” as used herein refer to any medium or media that participate in providing instructions to a CPU for execution. Such media can take many forms, including, but not limited to, non-volatile media, volatile media and transmission media. Non-volatile media include, for example, optical, magnetic, and solid-state disks, such as a fixed disk. Volatile media include dynamic memory, such as system random-access memory (RAM). Transmission media include coaxial cables, copper wire and fiber optics, among others, including the wires that comprise one embodiment of a bus. Transmission media can also take the form of acoustic or light waves, such as those generated during radio frequency (RF) and infrared (IR) data communications. Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, a hard disk, magnetic tape, any other magnetic medium, a CD-ROM disk, digital video disk (DVD), any other optical medium, any other physical medium with patterns of marks or holes, a RAM, a programmable read-only memory (PROM), an erasable programmable read-only memory (EPROM), an electrically erasable programmable read-only memory (EEPROM), a Flash memory, any other memory chip or data exchange adapter, a carrier wave, or any other medium from which a computer can read.

[0085] Various forms of computer-readable media may be involved in carrying one or more sequences of one or more instructions to a CPU for execution. A bus carries the data to system RAM, from which a CPU retrieves and executes the instructions. The instructions received by system RAM can optionally be stored on a fixed disk either before or after execution by a CPU.

[0086] Computer program code for carrying out operations for aspects of the present technology may be written in any combination of one or more programming languages, including an object oriented programming language such as JAVA, SMALLTALK, C++ or the like and conventional procedural programming languages, such as the “C” programming language or similar programming languages. The program code may execute entirely on the user’s computer, partly on the user’s computer, as a stand-alone software package, partly on the user’s computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user’s computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

[0087] The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present technology has been presented for purposes of

illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. Exemplary embodiments were chosen and described in order to best explain the principles of the present technology and its practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

[0088] Aspects of the present technology are described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0089] These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

[0090] The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0091] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present technology. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based

systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[0092] The description of the present technology has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. Exemplary embodiments were chosen and described in order to best explain the principles of the present technology and its practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

What is claimed is:

1. A computer-implemented method for validating security policy in a cloud computing environment, the method comprising:

providing a graph database, the graph database representing workloads of the cloud computing environment as nodes and relationships between the workloads as edges;

receiving a security policy, the security policy logically describing rules for the relationships between the workloads;

determining, based on the security policy and the graph database, a list of violations, the list of violations including at least one relationship from the relationships between the workloads in the graph database, the at least one relationship being not allowed by at least one of the rules in the security policy; and

providing the list of violations to a user.

2. The computer-implemented method of claim 1, wherein the graph database is created and updated based on a data concerning the cloud computing environment, the data including at least one of streaming telemetry from network logs, events from a cloud control plane, and inventory from a configuration management database.

3. The computer-implemented method of claim 1, wherein the security policy is a programmatically readable document including one or more of the following: a JavaScript Object Notation document, a Language (YAML) document, and an Open Policy Agent (OPA) rules document.

4. The computer-implemented method of claim 1, wherein the security policy is generated based on a security template, the security template including protected workloads in the cloud computing environment.

5. The computer-implemented method of claim 1, wherein the security policy is created by the user based upon manually declared requirements.

6. The computer-implemented method of claim 1, wherein the determining the list of violations includes:

traversing the nodes and the edges of the graph database; and

while traversing:

determining a type of connection between at least two nodes; and

determining that the type of connection between the at least two nodes is not permitted by the rules of the security policy.

7. The computer-implemented method of claim 1, wherein the providing the list of violations to the user includes:

displaying a visual representation of the nodes and the edges of the graph database, wherein the edges correspond to connections not allowed by the at least one of the rules in the security policy; or
a tabular representation using JSON.

8. The computer-implemented method of claim 7, further comprising:

receiving, from the user, a user input indicative of allowing or disallowing the at the least one relationship not allowed by the at least one of the rules of the security policy; and
modifying the security policy based on the user input.

9. The computer-implemented method of claim 8, further comprising:

deploying the security policy to the cloud computing environment.

10. The computer-implemented method of claim 1, wherein the cloud computing environment is hosted by a plurality of different cloud services.

11. A system for managing security in a cloud computing environment, the system comprising:

a processor; and

a memory communicatively coupled to the processor, the memory storing instructions executable by the processor to perform a method comprising:

providing a graph database, the graph database representing workloads of the cloud computing environment as nodes and relationships between the workloads as edges;

receiving a security policy, the security policy logically describing rules for the relationships between the workloads;

determining, based on the security policy and the graph database, a list of violations, the list of violations including at least one relationship from the relationships between the workloads in the graph database, the at least one relationship being not allowed by at least one of the rules in the security policy; and

providing the list of violations to a user.

12. The system of claim 11, wherein the graph database is created and updated based on a data concerning the cloud computing environment, the data including at least one of streaming telemetry from network logs, events from a cloud control plane, and inventory from a configuration management database.

13. The system of claim 11, wherein the security policy is a JavaScript Object Notation document.

14. The system of claim 11, wherein the security policy is generated based on a security template, the security template including protected workloads in the cloud computing environment.

15. The system of claim 11, wherein the security policy is created by the user.

16. The system of claim 11, wherein the determining the list of violations includes:

traversing the nodes and the edges of the graph database; and

while traversing:

determining a type of connection between at least two nodes; and

determining that the type of connection between the at least two nodes is not permitted by the rules of the security policy.

17. The system of claim 11, wherein the providing the list of violations to the user includes:

displaying a visual representation of the nodes and the edges of the graph database, wherein the edges correspond to connections not allowed by the at least one of the rules in the security policy.

18. The system of claim 17, wherein the processor performs the method further comprising:

receiving, from the user, a user input indicative of allowing or disallowing the at the least one relationship not allowed by the at least one of the rules of the security policy; and

modifying the security policy based on the user input.

19. The system of claim 8, wherein the processor performs the method further comprising:

deploying the security policy to the cloud computing environment.

20. A non-transitory processor-readable medium having embodied thereon a program being executable by at least one processor to perform a method for validating security policy in a cloud computing environment, the method comprising:

providing a graph database, the graph database representing workloads of the cloud computing environment as nodes and relationships between the workloads as edges;

receiving a security policy, the security policy logically describing rules for the relationships between the workloads;

determining, based on the security policy and the graph database, a list of violations, the list of violations including at least one relationship from the relationships between the workloads in the graph database, the at least one relationship being not allowed by at least one of the rules in the security policy; and

providing the list of violations to a user.

* * * * *