



(51) International Patent Classification:
H04L 29/00 (2006.01)

(21) International Application Number:
PCT/US2019/035704

(22) International Filing Date:
06 June 2019 (06.06.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/681,279 06 June 2018 (06.06.2018) US
62/737,558 27 September 2018 (27.09.2018) US
62/817,943 13 March 2019 (13.03.2019) US

(71) Applicant: RELIAQUEST HOLDINGS, LLC [US/US];
777 S. Harbour Island Blvd., Suite 500, Tampa, Florida
33602 (US).

(72) Inventors: MURPHY, Brian P.; 777 S. Harbour Island
Blvd., Suite 500, Tampa, Florida 33602 (US). PARTLOW,
Joe; 777 S. Harbour Island Blvd., Suite 500, Tampa, Flori-
da 33602 (US). O' CONNOR, Colin; 777 S. Harbour Is-

land Blvd., Suite 500, Tampa, FL 33602 (US). PFEIFFER,
Jason; 777 S. Harbour Island Blvd., Suite 500, Tampa, FL
33602 (US).

(74) Agent: COLANDREO, Brian J. et al.; Holland & Knight
LLP, 10 St. James Avenue, Boston, Massachusetts 02116
(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.

(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,

(54) Title: THREAT MITIGATION SYSTEM AND METHOD

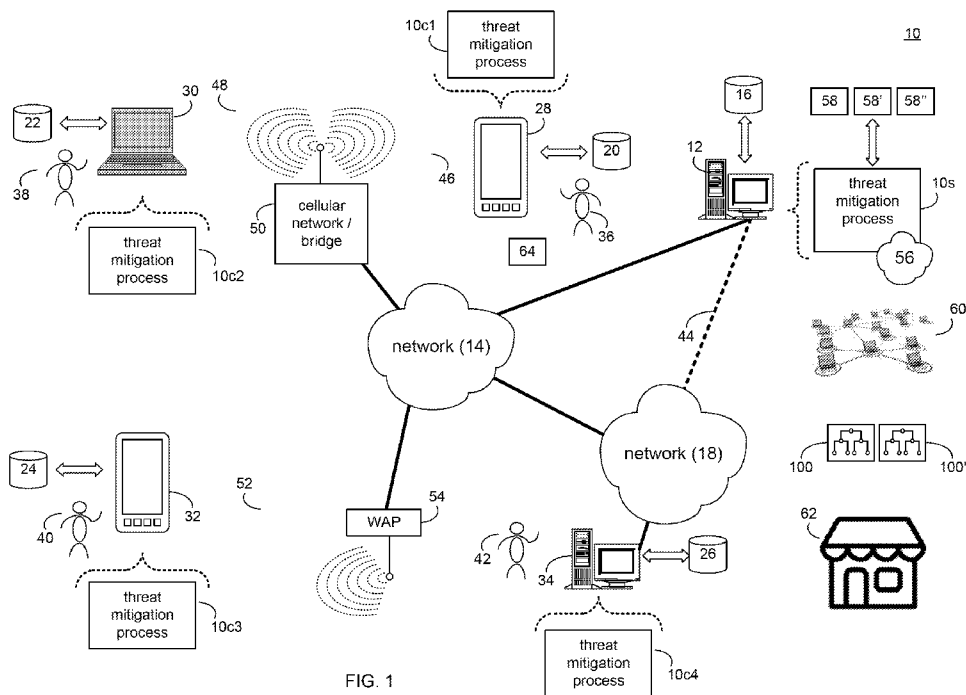


FIG. 1

(57) Abstract: A computer-implemented method, computer program product and computing system for: obtaining system-defined consolidated platform information for a computing platform from an independent information source; obtaining client-defined consolidated platform information for the computing platform from a client information source; and presenting differential consolidated platform information for the computing platform to the third-party.



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))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

Threat Mitigation System and Method

Related Application(s)

[001] This application claims the benefit of the following U.S. Provisional Application Nos.: 62/681,279, filed on 06 June 2018; 62/737,558, filed on 27 September 2018; and 62/817,943 filed on 13 March 2019, their entire contents of which are herein incorporated by reference.

Technical Field

[002] This disclosure relates to threat mitigation systems and, more particularly, to threat mitigation systems that utilize Artificial Intelligence (AI) and Machine Learning (ML).

Background

[003] In the computer world, there is a constant battle occurring between bad actors that want to attack computing platforms and good actors who try to prevent the same. Unfortunately, the complexity of such computer attacks is constantly increasing, so technology needs to be employed that understands the complexity of these attacks and is capable of addressing the same. Additionally, the use of Artificial Intelligence (AI) and Machine Learning (ML) has revolutionized the manner in which large quantities of content may be processed so that information may be extracted that is not readily discernible to a human user. Accordingly and though the use of AI / ML, the good actors may gain the upper hand in this never ending battle.

Summary of Disclosure

Concept 3)

[004] In one implementation, a computer-implemented method is executed on a computing device and includes: obtaining system-defined consolidated platform information for a computing platform from an independent information source; obtaining client-defined consolidated platform information for the computing platform from a client information source; and presenting differential consolidated platform information for the computing platform to the third-party.

[005] One or more of the following features may be included. The system-defined consolidated platform information may be compared to the client-defined consolidated platform information to define the differential consolidated platform information for the computing platform.

The system-defined consolidated platform information may be processed prior to comparing the system-defined consolidated platform information to the client-defined consolidated platform information to define differential consolidated platform information for the computing platform. The system-defined consolidated platform information may be obtained from one or more log files defined for the computing platform. The system-defined consolidated platform information may be obtained from two or more security-relevant subsystems deployed within the computing platform. The client-defined consolidated platform information may be obtained from one or more client-completed questionnaires. The client-defined consolidated platform information may be obtained from one or more client-deployed platform monitors.

[006] In another implementation, a computer program product resides on a computer readable medium and has a plurality of instructions stored on it. When executed by a processor, the instructions cause the processor to perform operations including: obtaining system-defined consolidated platform information for a computing platform from an independent information source; obtaining client-defined consolidated platform information for the computing platform from a client information source; and presenting differential consolidated platform information for the computing platform to the third-party.

[007] One or more of the following features may be included. The system-defined consolidated platform information may be compared to the client-defined consolidated platform information to define the differential consolidated platform information for the computing platform. The system-defined consolidated platform information may be processed prior to comparing the system-defined consolidated platform information to the client-defined consolidated platform information to define differential consolidated platform information for the computing platform. The system-defined consolidated platform information may be obtained from one or more log files defined for the computing platform. The system-defined consolidated platform information may be obtained from two or more security-relevant subsystems deployed within the computing platform. The client-defined consolidated platform information may be obtained from one or more client-completed questionnaires. The client-defined consolidated platform information may be obtained from one or more client-deployed platform monitors.

[008] In another implementation, a computing system includes a processor and memory is configured to perform operations including: obtaining system-defined consolidated platform information for a computing platform from an independent information source; obtaining client-defined consolidated platform information for the computing platform from a client information source; and presenting differential consolidated platform information for the computing platform to the third-party.

[009] One or more of the following features may be included. The system-defined consolidated platform information may be compared to the client-defined consolidated platform information to define the differential consolidated platform information for the computing platform. The system-defined consolidated platform information may be processed prior to comparing the system-defined consolidated platform information to the client-defined consolidated platform information to define differential consolidated platform information for the computing platform. The system-defined consolidated platform information may be obtained from one or more log files defined for the computing platform. The system-defined consolidated platform information may be obtained from two or more security-relevant subsystems deployed within the computing platform. The client-defined consolidated platform information may be obtained from one or more client-completed questionnaires. The client-defined consolidated platform information may be obtained from one or more client-deployed platform monitors.

[0010] The details of one or more implementations are set forth in the accompanying drawings and the description below. Other features and advantages will become apparent from the description, the drawings, and the claims.

Brief Description of the Drawings

[0011] FIG. 1 is a diagrammatic view of a distributed computing network including a computing device that executes a threat mitigation process according to an embodiment of the present disclosure;

[0012] FIG. 2 is a diagrammatic view of an exemplary probabilistic model rendered by a probabilistic process of the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0013] FIG. 3 is a diagrammatic view of the computing platform of FIG. 1 according to an embodiment of the present disclosure;

[0014] FIG. 4 is a flowchart of an implementation of the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0015] FIGS. 5-6 are diagrammatic views of screens rendered by the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0016] FIGS. 7-9 are flowcharts of other implementations of the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0017] FIG. 10 is a diagrammatic view of a screen rendered by the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0018] FIG. 11 is a flowchart of another implementation of the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0019] FIG. 12 is a diagrammatic view of a screen rendered by the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0020] FIG. 13 is a flowchart of another implementation of the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0021] FIG. 14 is a diagrammatic view of a screen rendered by the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0022] FIG. 15 is a flowchart of another implementation of the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0023] FIG. 16 is a diagrammatic view of screens rendered by the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0024] FIGS. 17-23 are flowcharts of other implementations of the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure;

[0025] FIG. 24 is a diagrammatic view of a screen rendered by the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure; and

[0026] FIGS. 25-30 are flowcharts of other implementations of the threat mitigation process of FIG. 1 according to an embodiment of the present disclosure.

[0027] Like reference symbols in the various drawings indicate like elements.

Detailed Description of the Preferred Embodiments

System Overview

[0028] Referring to FIG 1, there is shown threat mitigation process 10. Threat mitigation process 10 may be implemented as a server-side process, a client-side process, or a hybrid server-side / client-side process. For example, threat mitigation process 10 may be implemented as a purely server-side process via threat mitigation process 10s. Alternatively, threat mitigation process 10 may be implemented as a purely client-side process via one or more of threat mitigation process 10c1, threat mitigation process 10c2, threat mitigation process 10c3, and threat mitigation process 10c4. Alternatively still, threat mitigation process 10 may be implemented as a hybrid server-side / client-side process via threat mitigation process 10s in combination with one or more of threat mitigation process 10c1, threat mitigation process 10c2, threat mitigation process 10c3, and threat mitigation process 10c4. Accordingly, threat mitigation process 10 as used in this disclosure may include any combination of threat mitigation process 10s, threat mitigation process 10c1, threat mitigation process 10c2, threat mitigation process, and threat mitigation process 10c4.

[0029] Threat mitigation process 10s may be a server application and may reside on and may be executed by computing device 12, which may be connected to network 14 (e.g., the Internet or a local area network). Examples of computing device 12 may include, but are not limited to: a personal computer, a laptop computer, a personal digital assistant, a data-enabled cellular telephone, a notebook computer, a television with one or more processors embedded therein or coupled thereto, a cable / satellite receiver with one or more processors embedded therein or coupled thereto, a server computer, a series of server computers, a mini computer, a mainframe computer, or a cloud-based computing network.

[0030] The instruction sets and subroutines of threat mitigation process 10s, which may be stored on storage device 16 coupled to computing device 12, may be executed by one or more processors (not shown) and one or more memory architectures (not shown) included within computing device 12. Examples of storage device 16 may include but are not limited to: a hard disk drive; a RAID device; a random access memory (RAM); a read-only memory (ROM); and all forms of flash memory storage devices.

[0031] Network 14 may be connected to one or more secondary networks (e.g., network 18), examples of which may include but are not limited to: a local area network; a wide area network; or an intranet, for example.

[0032] Examples of threat mitigation processes 10c1, 10c2, 10c3, 10c4 may include but are not limited to a client application, a web browser, a game console user interface, or a specialized application (e.g., an application running on e.g., the Android™ platform or the iOS™ platform). The instruction sets and subroutines of threat mitigation processes 10c1, 10c2, 10c3, 10c4, which may be stored on storage devices 20, 22, 24, 26 (respectively) coupled to client electronic devices 28, 30, 32, 34 (respectively), may be executed by one or more processors (not shown) and one or more memory architectures (not shown) incorporated into client electronic devices 28, 30, 32, 34 (respectively). Examples of storage device 16 may include but are not limited to: a hard disk drive; a RAID device; a random access memory (RAM); a read-only memory (ROM); and all forms of flash memory storage devices.

[0033] Examples of client electronic devices 28, 30, 32, 34 may include, but are not limited to, data-enabled, cellular telephone 28, laptop computer 30, personal digital assistant 32, personal computer 34, a notebook computer (not shown), a server computer (not shown), a gaming console (not shown), a smart television (not shown), and a dedicated network device (not shown). Client electronic devices 28, 30, 32, 34 may each execute an operating system, examples of which may include but are not limited to Microsoft Windows™, Android™, WebOS™, iOS™, Redhat Linux™, or a custom operating system.

[0034] Users 36, 38, 40, 42 may access threat mitigation process 10 directly through network 14 or through secondary network 18. Further, threat mitigation process 10 may be connected to network 14 through secondary network 18, as illustrated with link line 44.

[0035] The various client electronic devices (e.g., client electronic devices 28, 30, 32, 34) may be directly or indirectly coupled to network 14 (or network 18). For example, data-enabled, cellular telephone 28 and laptop computer 30 are shown wirelessly coupled to network 14 via wireless communication channels 46, 48 (respectively) established between data-enabled, cellular telephone 28, laptop computer 30 (respectively) and cellular network / bridge 50, which is shown directly coupled to network 14. Further, personal digital assistant 32 is shown wirelessly coupled to

network 14 via wireless communication channel 52 established between personal digital assistant 32 and wireless access point (i.e., WAP) 54, which is shown directly coupled to network 14. Additionally, personal computer 34 is shown directly coupled to network 18 via a hardwired network connection.

[0036] WAP 54 may be, for example, an IEEE 802.11a, 802.11b, 802.11g, 802.11n, Wi-Fi, and/or Bluetooth device that is capable of establishing wireless communication channel 52 between personal digital assistant 32 and WAP 54. As is known in the art, IEEE 802.11x specifications may use Ethernet protocol and carrier sense multiple access with collision avoidance (i.e., CSMA/CA) for path sharing. The various 802.11x specifications may use phase-shift keying (i.e., PSK) modulation or complementary code keying (i.e., CCK) modulation, for example. As is known in the art, Bluetooth is a telecommunications industry specification that allows e.g., mobile phones, computers, and personal digital assistants to be interconnected using a short-range wireless connection.

Artificial Intelligence / Machines Learning Overview:

[0037] Assume for illustrative purposes that threat mitigation process 10 includes probabilistic process 56 (e.g., an artificial intelligence / machine learning process) that is configured to process information (e.g., information 58). As will be discussed below in greater detail, examples of information 58 may include but are not limited to platform information (e.g., structured or unstructured content) being scanned to detect security events (e.g., access auditing; anomalies; authentication; denial of services; exploitation; malware; phishing; spamming; reconnaissance; and/or web attack) within a monitored computing platform (e.g., computing platform 60).

[0038] As is known in the art, structured content may be content that is separated into independent portions (e.g., fields, columns, features) and, therefore, may have a pre-defined data model and/or is organized in a pre-defined manner. For example, if the structured content concerns an employee list: a first field, column or feature may define the first name of the employee; a second field, column or feature may define the last name of the employee; a third field, column or feature may define the home address of the employee; and a fourth field, column or feature may define the hire date of the employee.

[0039] Further and as is known in the art, unstructured content may be content that is not separated into independent portions (e.g., fields, columns, features) and, therefore, may not have a pre-defined data model and/or is not organized in a pre-defined manner. For example, if the unstructured content concerns the same employee list: the first name of the employee, the last name of the employee, the home address of the employee, and the hire date of the employee may all be combined into one field, column or feature.

[0040] For the following illustrative example, assume that information 58 is unstructured content, an example of which may include but is not limited to unstructured user feedback received by a company (e.g., text-based feedback such as text-messages, social media posts, and email messages; and transcribed voice-based feedback such as transcribed voice mail, and transcribed voice messages).

[0041] When processing information 58, probabilistic process 56 may use probabilistic modeling to accomplish such processing, wherein examples of such probabilistic modeling may include but are not limited to discriminative modeling, generative modeling, or combinations thereof.

[0042] As is known in the art, probabilistic modeling may be used within modern artificial intelligence systems (e.g., probabilistic process 56), in that these probabilistic models may provide artificial intelligence systems with the tools required to autonomously analyze vast quantities of data (e.g., information 58).

[0043] Examples of the tasks for which probabilistic modeling may be utilized may include but are not limited to:

- predicting media (music, movies, books) that a user may like or enjoy based upon media that the user has liked or enjoyed in the past;
- transcribing words spoken by a user into editable text;
- grouping genes into gene clusters;
- identifying recurring patterns within vast data sets;
- filtering email that is believed to be spam from a user's inbox;
- generating clean (i.e., non-noisy) data from a noisy data set;
- analyzing (voice-based or text-based) customer feedback; and

- diagnosing various medical conditions and diseases.

[0044] For each of the above-described applications of probabilistic modeling, an initial probabilistic model may be defined, wherein this initial probabilistic model may be subsequently (e.g., iteratively or continuously) modified and revised, thus allowing the probabilistic models and the artificial intelligence systems (e.g., probabilistic process 56) to “learn” so that future probabilistic models may be more precise and may explain more complex data sets.

[0045] Accordingly, probabilistic process 56 may define an initial probabilistic model for accomplishing a defined task (e.g., the analyzing of information 58). For the illustrative example, assume that this defined task is analyzing customer feedback (e.g., information 58) that is received from customers of e.g., store 62 via an automated feedback phone line. For this example, assume that information 58 is initially voice-based content that is processed via e.g., a speech-to-text process that results in unstructured text-based customer feedback (e.g., information 58).

[0046] With respect to probabilistic process 56, a probabilistic model may be utilized to go from initial observations about information 58 (e.g., as represented by the initial branches of a probabilistic model) to conclusions about information 58 (e.g., as represented by the leaves of a probabilistic model).

[0047] As used in this disclosure, the term “branch” may refer to the existence (or non-existence) of a component (e.g., a sub-model) of (or included within) a model. Examples of such a branch may include but are not limited to: an execution branch of a probabilistic program or other generative model, a part (or parts) of a probabilistic graphical model, and/or a component neural network that may (or may not) have been previously trained.

[0048] While the following discussion provides a detailed example of a probabilistic model, this is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, the following discussion may concern any type of model (e.g., be it probabilistic or other) and, therefore, the below-described probabilistic model is merely intended to be one illustrative example of a type of model and is not intended to limit this disclosure to probabilistic models.

[0049] Additionally, while the following discussion concerns word-based routing of messages through a probabilistic model, this is for illustrative purposes only and is not intended to be a

limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. Examples of other types of information that may be used to route messages through a probabilistic model may include: the order of the words within a message; and the punctuation interspersed throughout the message.

[0050] For example and referring also to FIG 2, there is shown one simplified example of a probabilistic model (e.g., probabilistic model 100) that may be utilized to analyze information 58 (e.g., unstructured text-based customer feedback) concerning store 62. The manner in which probabilistic model 100 may be automatically-generated by probabilistic process 56 will be discussed below in detail. In this particular example, probabilistic model 100 may receive information 58 (e.g., unstructured text-based customer feedback) at branching node 102 for processing. Assume that probabilistic model 100 includes four branches off of branching node 102, namely: service branch 104; selection branch 106; location branch 108; and value branch 110 that respectively lead to service node 112, selection node 114, location node 116, and value node 118.

[0051] As stated above, service branch 104 may lead to service node 112, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) feedback concerning the customer service of store 62. For example, service node 112 may define service word list 120 that may include e.g., the word service, as well as synonyms of (and words related to) the word service (e.g., cashier, employee, greeter and manager). Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) includes the word cashier, employee, greeter and/or manager, that portion of information 58 may be considered to be text-based customer feedback concerning the service received at store 62 and (therefore) may be routed to service node 112 of probabilistic model 100 for further processing. Assume for this illustrative example that probabilistic model 100 includes two branches off of service node 112, namely: good service branch 122 and bad service branch 124.

[0052] Good service branch 122 may lead to good service node 126, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) good feedback concerning the customer service of store 62. For example, good service node 126 may define good service word list 128 that may include e.g., the

word good, as well as synonyms of (and words related to) the word good (e.g., courteous, friendly, lovely, happy, and smiling). Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) that was routed to service node 112 includes the word good, courteous, friendly, lovely, happy, and/or smiling, that portion of information 58 may be considered to be text-based customer feedback indicative of good service received at store 62 (and, therefore, may be routed to good service node 126).

[0053] Bad service branch 124 may lead to bad service node 130, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) bad feedback concerning the customer service of store 62. For example, bad service node 130 may define bad service word list 132 that may include e.g., the word bad, as well as synonyms of (and words related to) the word bad (e.g., rude, mean, jerk, miserable, and scowling). Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) that was routed to service node 112 includes the word bad, rude, mean, jerk, miserable, and/or scowling, that portion of information 58 may be considered to be text-based customer feedback indicative of bad service received at store 62 (and, therefore, may be routed to bad service node 130).

[0054] As stated above, selection branch 106 may lead to selection node 114, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) feedback concerning the selection available at store 62. For example, selection node 114 may define selection word list 134 that may include e.g., words indicative of the selection available at store 62. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) includes any of the words defined within selection word list 134, that portion of information 58 may be considered to be text-based customer feedback concerning the selection available at store 62 and (therefore) may be routed to selection node 114 of probabilistic model 100 for further processing. Assume for this illustrative example that probabilistic model 100 includes two branches off of selection node 114, namely: good selection branch 136 and bad selection branch 138.

[0055] Good selection branch 136 may lead to good selection node 140, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer

feedback) that concerns (in whole or in part) good feedback concerning the selection available at store 62. For example, good selection node 140 may define good selection word list 142 that may include words indicative of a good selection at store 62. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) that was routed to selection node 114 includes any of the words defined within good selection word list 142, that portion of information 58 may be considered to be text-based customer feedback indicative of a good selection available at store 62 (and, therefore, may be routed to good selection node 140).

[0056] Bad selection branch 138 may lead to bad selection node 144, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) bad feedback concerning the selection available at store 62. For example, bad selection node 144 may define bad selection word list 146 that may include words indicative of a bad selection at store 62. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) that was routed to selection node 114 includes any of the words defined within bad selection word list 146, that portion of information 58 may be considered to be text-based customer feedback indicative of a bad selection being available at store 62 (and, therefore, may be routed to bad selection node 144).

[0057] As stated above, location branch 108 may lead to location node 116, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) feedback concerning the location of store 62. For example, location node 116 may define location word list 148 that may include e.g., words indicative of the location of store 62. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) includes any of the words defined within location word list 148, that portion of information 58 may be considered to be text-based customer feedback concerning the location of store 62 and (therefore) may be routed to location node 116 of probabilistic model 100 for further processing. Assume for this illustrative example that probabilistic model 100 includes two branches off of location node 116, namely: good location branch 150 and bad location branch 152.

[0058] Good location branch 150 may lead to good location node 154, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer

feedback) that concerns (in whole or in part) good feedback concerning the location of store 62. For example, good location node 154 may define good location word list 156 that may include words indicative of store 62 being in a good location. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) that was routed to location node 116 includes any of the words defined within good location word list 156, that portion of information 58 may be considered to be text-based customer feedback indicative of store 62 being in a good location (and, therefore, may be routed to good location node 154).

[0059] Bad location branch 152 may lead to bad location node 158, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) bad feedback concerning the location of store 62. For example, bad location node 158 may define bad location word list 160 that may include words indicative of store 62 being in a bad location. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) that was routed to location node 116 includes any of the words defined within bad location word list 160, that portion of information 58 may be considered to be text-based customer feedback indicative of store 62 being in a bad location (and, therefore, may be routed to bad location node 158).

[0060] As stated above, value branch 110 may lead to value node 118, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) feedback concerning the value received at store 62. For example, value node 118 may define value word list 162 that may include e.g., words indicative of the value received at store 62. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) includes any of the words defined within value word list 162, that portion of information 58 may be considered to be text-based customer feedback concerning the value received at store 62 and (therefore) may be routed to value node 118 of probabilistic model 100 for further processing. Assume for this illustrative example that probabilistic model 100 includes two branches off of value node 118, namely: good value branch 164 and bad value branch 166.

[0061] Good value branch 164 may lead to good value node 168, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that

concerns (in whole or in part) good value being received at store 62. For example, good value node 168 may define good value word list 170 that may include words indicative of receiving good value at store 62. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) that was routed to value node 118 includes any of the words defined within good value word list 170, that portion of information 58 may be considered to be text-based customer feedback indicative of good value being received at store 62 (and, therefore, may be routed to good value node 168).

[0062] Bad value branch 166 may lead to bad value node 172, which may be configured to process the portion of information 58 (e.g., unstructured text-based customer feedback) that concerns (in whole or in part) bad value being received at store 62. For example, bad value node 172 may define bad value word list 174 that may include words indicative of receiving bad value at store 62. Accordingly and in the event that a portion of information 58 (e.g., a text-based customer feedback message) that was routed to value node 118 includes any of the words defined within bad value word list 174, that portion of information 58 may be considered to be text-based customer feedback indicative of bad value being received at store 62 (and, therefore, may be routed to bad value node 172).

[0063] Once it is established that good or bad customer feedback was received concerning store 62 (i.e., with respect to the service, the selection, the location or the value), representatives and/or agents of store 62 may address the provider of such good or bad feedback via e.g., social media postings, text-messages and/or personal contact.

[0064] Assume for illustrative purposes that user 36 uses data-enabled, cellular telephone 28 to provide feedback 64 (e.g., a portion of information 58) to an automated feedback phone line concerning store 62. Upon receiving feedback 64 for analysis, probabilistic process 56 may identify any pertinent content that is included within feedback 64.

[0065] For illustrative purposes, assume that user 36 was not happy with their experience at store 62 and that feedback 64 provided by user 36 was “my cashier was rude and the weather was rainy”. Accordingly and for this example, probabilistic process 56 may identify the pertinent content (included within feedback 64) as the phrase “my cashier was rude” and may ignore / remove the irrelevant content “the weather was rainy”. As (in this example) feedback 64 includes

the word “cashier”, probabilistic process 56 may route feedback 64 to service node 112 via service branch 104. Further, as feedback 64 also includes the word “rude”, probabilistic process 56 may route feedback 64 to bad service node 130 via bad service branch 124 and may consider feedback 64 to be text-based customer feedback indicative of bad service being received at store 62.

[0066] For further illustrative purposes, assume that user 36 was happy with their experience at store 62 and that feedback 64 provided by user 36 was “the clothing I purchased was classy but my cab got stuck in traffic”. Accordingly and for this example, probabilistic process 56 may identify the pertinent content (included within feedback 64) as the phrase “the clothing I purchased was classy” and may ignore / remove the irrelevant content “my cab got stuck in traffic”. As (in this example) feedback 64 includes the word “clothing”, probabilistic process 56 may route feedback 64 to selection node 114 via selection branch 106. Further, as feedback 64 also includes the word “classy”, probabilistic process 56 may route feedback 64 to good selection node 140 via good selection branch 136 and may consider feedback 64 to be text-based customer feedback indicative of a good selection being available at store 62.

Model Generation Overview:

[0067] While the following discussion concerns the automated generation of a probabilistic model, this is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, the following discussion of automated generation may be utilized on any type of model. For example, the following discussion may be applicable to any other form of probabilistic model or any form of generic model (such as Dempster Shaffer theory or fuzzy logic).

[0068] As discussed above, probabilistic model 100 may be utilized to categorize information 58, thus allowing the various messages included within information 58 to be routed to (in this simplified example) one of eight nodes (e.g., good service node 126, bad service node 130, good selection node 140, bad selection node 144, good location node 154, bad location node 158, good value node 168, and bad value node 172). For the following example, assume that store 62 is a long-standing and well established shopping establishment. Further, assume that information 58 is a very large quantity of voice mail messages (>10,000 messages) that were left by customers of store 62 on a voice-based customer feedback line. Additionally, assume that this very large quantity

of voice mail messages (>10,000) have been transcribed into a very large quantity of text-based messages (>10,000).

[0069] Probabilistic process 56 may be configured to automatically define probabilistic model 100 based upon information 58. Accordingly, probabilistic process 56 may receive content (e.g., a very large quantity of text-based messages) and may be configured to define one or more probabilistic model variables for probabilistic model 100. For example, probabilistic process 56 may be configured to allow a user to specify such probabilistic model variables. Another example of such variables may include but is not limited to values and/or ranges of values for a data flow variable. For the following discussion and for this disclosure, examples of a “variable” may include but are not limited to variables, parameters, ranges, branches and nodes.

[0070] Specifically and for this example, assume that probabilistic process 56 defines the initial number of branches (i.e., the number of branches off of branching node 102) within probabilistic model 100 as four (i.e., service branch 104, selection branch 106, location branch 108 and value branch 110). The defining of the initial number of branches (i.e., the number of branches off of branching node 102) within probabilistic model 100 as four may be effectuated in various ways (e.g., manually or algorithmically). Further and when defining probabilistic model 100 based, at least in part, upon information 58 and the one or more model variables (i.e., defining the number of branches off of branching node 102 as four), probabilistic process 56 may process information 58 to identify the pertinent content included within information 58. As discussed above, probabilistic process 56 may identify the pertinent content (included within information 58) and may ignore / remove the irrelevant content.

[0071] This type of processing of information 58 may continue for all of the very large quantity of text-based messages (>10,000) included within information 58. And using the probabilistic modeling technique described above, probabilistic process 56 may define a first version of the probabilistic model (e.g., probabilistic model 100) based, at least in part, upon pertinent content found within information 58. Accordingly, a first text-based message included within information 58 may be processed to extract pertinent information from that first message, wherein this pertinent information may be grouped in a manner to correspond (at least temporarily) with the requirement that four branches originate from branching node 102 (as defined above).

[0072] As probabilistic process 56 continues to process information 58 to identify pertinent content included within information 58, probabilistic process 56 may identify patterns within these text-based message included within information 58. For example, the messages may all concern one or more of the service, the selection, the location and/or the value of store 62. Further and e.g., using the probabilistic modeling technique described above, probabilistic process 56 may process information 58 to e.g.: a) sort text-based messages concerning the service into positive or negative service messages; b) sort text-based messages concerning the selection into positive or negative selection messages; c) sort text-based messages concerning the location into positive or negative location messages; and/or d) sort text-based messages concerning the value into positive or negative service messages. For example, probabilistic process 56 may define various lists (e.g., lists 128, 132, 142, 146, 156, 160, 170, 174) by starting with a root word (e.g., good or bad) and may then determine synonyms for these words and use those words and synonyms to populate lists 128, 132, 142, 146, 156, 160, 170, 174.

[0073] Continuing with the above-stated example, once information 58 (or a portion thereof) is processed by probabilistic process 56, probabilistic process 56 may define a first version of the probabilistic model (e.g., probabilistic model 100) based, at least in part, upon pertinent content found within information 58. Probabilistic process 56 may compare the first version of the probabilistic model (e.g., probabilistic model 100) to information 58 to determine if the first version of the probabilistic model (e.g., probabilistic model 100) is a good explanation of the content.

[0074] When determining if the first version of the probabilistic model (e.g., probabilistic model 100) is a good explanation of the content, probabilistic process 56 may use an ML algorithm to fit the first version of the probabilistic model (e.g., probabilistic model 100) to the content, wherein examples of such an ML algorithm may include but are not limited to one or more of: an inferencing algorithm, a learning algorithm, an optimization algorithm, and a statistical algorithm.

[0075] For example and as is known in the art, probabilistic model 100 may be used to generate messages (in addition to analyzing them). For example and when defining a first version of the probabilistic model (e.g., probabilistic model 100) based, at least in part, upon pertinent content found within information 58, probabilistic process 56 may define a weight for each branch within probabilistic model 100 based upon information 58. For example, threat mitigation process 10 may

equally weight each of branches 104, 106, 108, 110 at 25%. Alternatively, if e.g., a larger percentage of information 58 concerned the service received at store 62, threat mitigation process 10 may equally weight each of branches 106, 108, 110 at 20%, while more heavily weighting branch 104 at 40%.

[0076] Accordingly and when probabilistic process 56 compares the first version of the probabilistic model (e.g., probabilistic model 100) to information 58 to determine if the first version of the probabilistic model (e.g., probabilistic model 100) is a good explanation of the content, probabilistic process 56 may generate a very large quantity of messages e.g., by auto-generating messages using the above-described probabilities, the above-described nodes & node types, and the words defined in the above-described lists (e.g., lists 128, 132, 142, 146, 156, 160, 170, 174), thus resulting in generated information 58'. Generated information 58' may then be compared to information 58 to determine if the first version of the probabilistic model (e.g., probabilistic model 100) is a good explanation of the content. For example, if generated information 58' exceeds a threshold level of similarity to information 58, the first version of the probabilistic model (e.g., probabilistic model 100) may be deemed a good explanation of the content. Conversely, if generated information 58' does not exceed a threshold level of similarity to information 58, the first version of the probabilistic model (e.g., probabilistic model 100) may be deemed not a good explanation of the content.

[0077] If the first version of the probabilistic model (e.g., probabilistic model 100) is not a good explanation of the content, probabilistic process 56 may define a revised version of the probabilistic model (e.g., revised probabilistic model 100'). When defining revised probabilistic model 100', probabilistic process 56 may e.g., adjust weighting, adjust probabilities, adjust node counts, adjust node types, and/or adjust branch counts to define the revised version of the probabilistic model (e.g., revised probabilistic model 100'). Once defined, the above-described process of auto-generating messages (this time using revised probabilistic model 100') may be repeated and this newly-generated content (e.g., generated information 58'') may be compared to information 58 to determine if e.g., revised probabilistic model 100' is a good explanation of the content. If revised probabilistic model 100' is not a good explanation of the content, the above-described process may be repeated until a proper probabilistic model is defined.

The Threat Mitigation Process

[0078] As discussed above, threat mitigation process 10 may include probabilistic process 56 (e.g., an artificial intelligence / machine learning process) that may be configured to process information (e.g., information 58), wherein examples of information 58 may include but are not limited to platform information (e.g., structured or unstructured content) that may be scanned to detect security events (e.g., access auditing; anomalies; authentication; denial of services; exploitation; malware; phishing; spamming; reconnaissance; and/or web attack) within a monitored computing platform (e.g., computing platform 60).

[0079] Referring also to FIG 3, the monitored computing platform (e.g., computing platform 60) utilized by business today may be a highly complex, multi-location computing system / network that may span multiple buildings / locations / countries. For this illustrative example, the monitored computing platform (e.g., computing platform 60) is shown to include many discrete computing devices, examples of which may include but are not limited to: server computers (e.g., server computers 200, 202), desktop computers (e.g., desktop computer 204), and laptop computers (e.g., laptop computer 206), all of which may be coupled together via a network (e.g., network 208), such as an Ethernet network. Computing platform 60 may be coupled to an external network (e.g., Internet 210) through WAF (i.e., Web Application Firewall) 212. A wireless access point (e.g., WAP 214) may be configured to allow wireless devices (e.g., smartphone 216) to access computing platform 60. Computing platform 60 may include various connectivity devices that enable the coupling of devices within computing platform 60, examples of which may include but are not limited to: switch 216, router 218 and gateway 220. Computing platform 60 may also include various storage devices (e.g., NAS 222), as well as functionality (e.g., API Gateway 224) that allows software applications to gain access to one or more resources within computing platform 60.

[0080] In addition to the devices and functionality discussed above, other technology (e.g., security-relevant subsystems 226) may be deployed within computing platform 60 to monitor the operation of (and the activity within) computing platform 60. Examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management)

systems; DNS (i.e., Domain Name Server) systems, antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform.

[0081] Each of security-relevant subsystems 226 may monitor and log their activity with respect to computing platform 60, resulting in the generation of platform information 228. For example, platform information 228 associated with a client-defined MDM (i.e., Mobile Device Management) system may monitor and log the mobile devices that were allowed access to computing platform 60.

[0082] Further, SEIM (i.e., Security Information and Event Management) system 230 may be deployed within computing platform 60. As is known in the art, SIEM system 230 is an approach to security management that combines SIM (security information management) functionality and SEM (security event management) functionality into one security management system. The underlying principles of a SIEM system is to aggregate relevant data from multiple sources, identify deviations from the norm and take appropriate action. For example, when a security event is detected, SIEM system 230 might log additional information, generate an alert and instruct other security controls to mitigate the security event. Accordingly, SIEM system 230 may be configured to monitor and log the activity of security-relevant subsystems 226 (e.g., CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform).

Computing Platform Analysis & Reporting

[0083] As will be discussed below in greater detail, threat mitigation process 10 may be configured to e.g., analyze computing platform 60 and provide reports to third-parties concerning the same.

Concept 1)

[0084] Referring also to FIGS. 4-6, threat mitigation process 10 may be configured to obtain and combine information from multiple security-relevant subsystem to generate a security profile

for computing platform 60. For example, threat mitigation process 10 may obtain 300 first system-defined platform information (e.g., system-defined platform information 232) concerning a first security-relevant subsystem (e.g., the number of operating systems deployed) within computing platform 60 and may obtain 302 at least a second system-defined platform information (e.g., system-defined platform information 234) concerning at least a second security-relevant subsystem (e.g., the number of antivirus systems deployed) within computing platform 60.

[0085] The first system-defined platform information (e.g., system-defined platform information 232) and the at least a second system-defined platform information (e.g., system-defined platform information 234) may be obtained from one or more log files defined for computing platform 60.

[0086] Specifically, system-defined platform information 232 and/or system-defined platform information 234 may be obtained from SIEM system 230, wherein (and as discussed above) SIEM system 230 may be configured to monitor and log the activity of security-relevant subsystems 226 (e.g., CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform).

[0087] Alternatively, the first system-defined platform information (e.g., system-defined platform information 232) and the at least a second system-defined platform information (e.g., system-defined platform information 234) may be obtained from the first security-relevant subsystem (e.g., the operating systems themselves) and the at least a second security-relevant subsystem (e.g., the antivirus systems themselves). Specifically, system-defined platform information 232 and/or system-defined platform information 234 may be obtained directly from the security-relevant subsystems (e.g., the operating systems and/or the antivirus systems), which (as discussed above) may be configured to self-document their activity.

[0088] Threat mitigation process 10 may combine 308 the first system-defined platform information (e.g., system-defined platform information 232) and the at least a second system-

defined platform information (e.g., system-defined platform information 234) to form system-defined consolidated platform information 236. Accordingly and in this example, system-defined consolidated platform information 236 may independently define the security-relevant subsystems (e.g., security-relevant subsystems 226) present on computing platform 60.

[0089] Threat mitigation process 10 may generate 310 a security profile (e.g., security profile 350) based, at least in part, upon system-defined consolidated platform information 236. Through the use of security profile (e.g., security profile 350), the user / owner / operator of computing platform 60 may be able to see that e.g., they have a security score of 605 out of a possible score of 1,000, wherein the average customer has a security score of 237. While security profile 350 in shown in the example to include several indicators that may enable a user to compare (in this example) computing platform 60 to other computing platforms, this is for illustrative purposes only and is not intended to be a limitation of this disclosure, as it is understood that other configurations are possible and are considered to be within the scope of this disclosure.

[0090] Naturally, the format, appearance and content of security profile 350 may be varied greatly depending upon the design criteria and anticipated performance / use of threat mitigation process 10. Accordingly, the appearance, format, completeness and content of security profile 350 is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, content may be added to security profile 350, removed from security profile 350, and/or reformatted within security profile 350.

[0091] Additionally, threat mitigation process 10 may obtain 312 client-defined consolidated platform information 238 for computing platform 60 from a client information source, examples of which may include but are not limited to one or more client-completed questionnaires (e.g., questionnaires 240) and/or one or more client-deployed platform monitors (e.g., client-deployed platform monitor 242, which may be configured to effectuate SIEM functionality). Accordingly and in this example, client-defined consolidated platform information 238 may define the security-relevant subsystems (e.g., security-relevant subsystems 226) that the client believes are present on computing platform 60.

[0092] When generating 310 a security profile (e.g., security profile 350) based, at least in part, upon system-defined consolidated platform information 236, threat mitigation process 10 may compare 314 the system-defined consolidated platform information (e.g., system-defined consolidated platform information 236) to the client-defined consolidated platform information (e.g., client-defined consolidated platform information 238) to define differential consolidated platform information 352 for computing platform 60.

[0093] Differential consolidated platform information 352 may include comparison table 354 that e.g., compares computing platform 60 to other computing platforms. For example and in this particular implementation of differential consolidated platform information 352, comparison table 354 is shown to include three columns, namely: security-relevant subsystem column 356 (that identifies the security-relevant subsystems in question); system-defined consolidated platform information column 358 (that is based upon system-defined consolidated platform information 236 and independently defines what security-relevant subsystems are present on computing platform 60); and client-defined consolidated platform column 360 (that is based upon client-defined platform information 238 and defines what security-relevant subsystems the client believes are present on computing platform 60). As shown within comparison table 354, there are considerable differences between that is actually present on computing platform 60 and what is believed to be present on computing platform 60 (e.g., 1 IAM system vs. 10 IAM systems; 4,000 operating systems vs. 10,000 operating systems, 6 DNS systems vs. 10 DNS systems; 0 antivirus systems vs. 1 antivirus system, and 90 firewalls vs. 150 firewalls).

[0094] Naturally, the format, appearance and content of differential consolidated platform information 352 may be varied greatly depending upon the design criteria and anticipated performance / use of threat mitigation process 10. Accordingly, the appearance, format, completeness and content of differential consolidated platform information 352 is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, content may be added to differential consolidated platform information 352, removed from differential consolidated platform information 352, and/or reformatted within differential consolidated platform information 352.

Concept 2)

[0095] Referring also to FIG 7, threat mitigation process 10 may be configured to compare what security relevant subsystems are actually included within computing platform 60 versus what security relevant subsystems were believed to be included within computing platform 60. As discussed above, threat mitigation process 10 may combine 308 the first system-defined platform information (e.g., system-defined platform information 232) and the at least a second system-defined platform information (e.g., system-defined platform information 234) to form system-defined consolidated platform information 236.

[0096] Threat mitigation process 10 may obtain 400 system-defined consolidated platform information 236 for computing platform 60 from an independent information source, examples of which may include but are not limited to: one or more log files defined for computing platform 60 (e.g., such as those maintained by SIEM system 230); and two or more security-relevant subsystems (e.g., directly from the operating system security-relevant subsystem and the antivirus security-relevant subsystem) deployed within computing platform 60.

[0097] Further and as discussed above, threat mitigation process 10 may obtain 312 client-defined consolidated platform information 238 for computing platform 60 from a client information source, examples of which may include but are not limited to one or more client-completed questionnaires (e.g., questionnaires 240) and/or one or more client-deployed platform monitors (e.g., client-deployed platform monitor 242, which may be configured to effectuate SIEM functionality).

[0098] Additionally and as discussed above, threat mitigation process 10 may compare 402 system-defined consolidated platform information 236 to client-defined consolidated platform information 238 to define differential consolidated platform information 352 for computing platform 60, wherein differential consolidated platform information 352 may include comparison table 354 that e.g., compares computing platform 60 to other computing platforms..

[0099] Threat mitigation process 10 may process 404 system-defined consolidated platform information 236 prior to comparing 402 system-defined consolidated platform information 236 to client-defined consolidated platform information 238 to define differential consolidated platform information 352 for computing platform 60. Specifically, threat mitigation process 10 may process

404 system-defined consolidated platform information 236 so that it is comparable to client-defined consolidated platform information 238.

[00100] For example and when processing 404 system-defined consolidated platform information 236, threat mitigation process 10 may homogenize 406 system-defined consolidated platform information 236 prior to comparing 402 system-defined consolidated platform information 236 to client-defined consolidated platform information 238 to define differential consolidated platform information 352 for computing platform 60. Such homogenization 406 may result in system-defined consolidated platform information 236 and client-defined consolidated platform information 238 being comparable to each other (e.g., to accommodate for differing data nomenclatures / headers).

[00101] Further and when processing 404 system-defined consolidated platform information 236, threat mitigation process 10 may normalize 408 system-defined consolidated platform information 236 prior to comparing 402 system-defined consolidated platform information 236 to client-defined consolidated platform information 238 to define differential consolidated platform information 352 for computing platform 60 (e.g., to accommodate for data differing scales / ranges).

Concept 3)

[00102] Referring also to FIG 8, threat mitigation process 10 may be configured to compare what security relevant subsystems are actually included within computing platform 60 versus what security relevant subsystems were believed to be included within computing platform 60.

[00103] As discussed above, threat mitigation process 10 may obtain 400 system-defined consolidated platform information 236 for computing platform 60 from an independent information source, examples of which may include but are not limited to: one or more log files defined for computing platform 60 (e.g., such as those maintained by SIEM system 230); and two or more security-relevant subsystems (e.g., directly from the operating system security-relevant subsystem and the antivirus security-relevant subsystem) deployed within computing platform 60

[00104] Further and as discussed above, threat mitigation process 10 may obtain 312 client-defined consolidated platform information 238 for computing platform 60 from a client information source, examples of which may include but are not limited to one or more client-completed questionnaires (e.g., questionnaires 240) and/or one or more client-deployed platform monitors

(e.g., client-deployed platform monitor 242, which may be configured to effectuate SIEM functionality).

[00105] Threat mitigation process 10 may present 450 differential consolidated platform information 352 for computing platform 60 to a third-party, examples of which may include but are not limited to the user / owner / operator of computing platform 60.

[00106] Additionally and as discussed above, threat mitigation process 10 may compare 402 system-defined consolidated platform information 236 to client-defined consolidated platform information 238 to define differential consolidated platform information 352 for computing platform 60, wherein differential consolidated platform information 352 may include comparison table 354 that e.g., compares computing platform 60 to other computing platforms, wherein (and as discussed above) threat mitigation process 10 may process 404 (e.g., via homogenizing 406 and/or normalizing 408) system-defined consolidated platform information 236 prior to comparing 402 system-defined consolidated platform information 236 to client-defined consolidated platform information 236 to define differential consolidated platform information 352 for computing platform 60.

Computing Platform Analysis & Recommendation

[00107] As will be discussed below in greater detail, threat mitigation process 10 may be configured to e.g., analyze & display the vulnerabilities of computing platform 60.

Concept 4)

[00108] Referring also to FIG 9, threat mitigation process 10 may be configured to make recommendations concerning security relevant subsystems that are missing from computing platform 60. As discussed above, threat mitigation process 10 may obtain 500 consolidated platform information for computing platform 60 to identify one or more deployed security-relevant subsystems 226 (e.g., CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform). This consolidated platform information may be obtained from an independent

information source (e.g., such as SIEM system 230 that may provide system-defined consolidated platform information 236) and/or may be obtained from a client information source (e.g., such as questionnaires 240 that may provide client-defined consolidated platform information 238).

[00109] Referring also to FIG 10, threat mitigation process 10 may process 506 the consolidated platform information (e.g., system-defined consolidated platform information 236 and/or client-defined consolidated platform information 238) to identify one or more non-deployed security-relevant subsystems (within computing platform 60) and may then generate 508 a list of ranked & recommended security-relevant subsystems (e.g., non-deployed security-relevant subsystem list 550) that ranks the one or more non-deployed security-relevant subsystems.

[00110] For this particular illustrative example, non-deployed security-relevant subsystem list 550 is shown to include column 552 that identifies six non-deployed security-relevant subsystems, namely: a CDN subsystem, a WAF subsystem, a DAM subsystem; a UBA subsystem; a API subsystem, and an MDM subsystem.

[00111] When generating 508 a list of ranked & recommended security-relevant subsystems (e.g., non-deployed security-relevant subsystem list 550) that ranks the one or more non-deployed security-relevant subsystems, threat mitigation process 10 may rank 510 the one or more non-deployed security-relevant subsystems (e.g., a CDN subsystem, a WAF subsystem, a DAM subsystem; a UBA subsystem; a API subsystem, and an MDM subsystem) based upon the anticipated use of the one or more non-deployed security-relevant subsystems within computing platform 60. This ranking 510 of the non-deployed security-relevant subsystems (e.g., a CDN subsystem, a WAF subsystem, a DAM subsystem; a UBA subsystem; a API subsystem, and an MDM subsystem) may be agnostic in nature and may be based on the functionality / effectiveness of the non-deployed security-relevant subsystems and the anticipated manner in which their implementation may impact the functionality / security of computing platform 60.

[00112] Threat mitigation process 10 may provide 512 the list of ranked & recommended security-relevant subsystems (e.g., non-deployed security-relevant subsystem list 550) to a third-party, examples of which may include but are not limited to a user / owner / operator of computing platform 60.

[00113] Additionally, threat mitigation process 10 may identify 514 a comparative for at least one of the non-deployed security-relevant subsystems (e.g., a CDN subsystem, a WAF subsystem, a DAM subsystem; a UBA subsystem; a API subsystem, and an MDM subsystem) defined within the list of ranked & recommended security-relevant subsystems (e.g., non-deployed security-relevant subsystem list 550). This comparative may include vendor customers in a specific industry comparative and/or vendor customers in any industry comparative.

[00114] For example and in addition to column 552, non-deployed security-relevant subsystem list 550 may include columns 554, 556 for defining the comparatives for the six non-deployed security-relevant subsystems, namely: a CDN subsystem, a WAF subsystem, a DAM subsystem; a UBA subsystem; a API subsystem, and an MDM subsystem. Specifically, column 554 is shown to define comparatives concerning vendor customers that own the non-deployed security-relevant subsystems in a specific industry (i.e., the same industry as the user / owner / operator of computing platform 60). Additionally, column 556 is shown to define comparatives concerning vendor customers that own the non-deployed security-relevant subsystems in any industry (i.e., not necessarily the same industry as the user / owner / operator of computing platform 60). For example and concerning the comparatives of the WAF subsystem: 33% of the vendor customers in the same industry as the user / owner / operator of computing platform 60 deploy a WAF subsystem; while 71% of the vendor customers in any industry deploy a WAF subsystem.

[00115] Naturally, the format, appearance and content of non-deployed security-relevant subsystem list 550 may be varied greatly depending upon the design criteria and anticipated performance / use of threat mitigation process 10. Accordingly, the appearance, format, completeness and content of non-deployed security-relevant subsystem list 550 is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, content may be added to non-deployed security-relevant subsystem list 550, removed from non-deployed security-relevant subsystem list 550, and/or reformatted within non-deployed security-relevant subsystem list 550.

Concept 5)

[00116] Referring also to FIG 11, threat mitigation process 10 may be configured to compare the current capabilities to the possible capabilities of computing platform 60. As discussed above, threat mitigation process 10 may obtain 600 consolidated platform information to identify current security-relevant capabilities for computing platform 60. This consolidated platform information may be obtained from an independent information source (e.g., such as SIEM system 230 that may provide system-defined consolidated platform information 236) and/or may be obtained from a client information source (e.g., such as questionnaires 240 that may provide client-defined consolidated platform information 238). Threat mitigation process 10 may then determine 606 possible security-relevant capabilities for computing platform 60 (i.e., the difference between the current security-relevant capabilities of computing platform 60 and the possible security-relevant capabilities of computing platform 60). For example, the possible security-relevant capabilities may concern the possible security-relevant capabilities of computing platform 60 using the currently-deployed security-relevant subsystems. Additionally / alternatively, the possible security-relevant capabilities may concern the possible security-relevant capabilities of computing platform 60 using one or more supplemental security-relevant subsystems.

[00117] Referring also to FIG 12 and as will be explained below, threat mitigation process 10 may generate 608 comparison information 650 that compares the current security-relevant capabilities of computing platform 60 to the possible security-relevant capabilities of computing platform 60 to identify security-relevant deficiencies. Comparison information 650 may include graphical comparison information, such as multi-axial graphical comparison information that simultaneously illustrates a plurality of security-relevant deficiencies.

[00118] For example, comparison information 650 may define (in this particular illustrative example) graphical comparison information that include five axes (e.g. axes 652, 654, 656, 658, 660) that correspond to five particular types of computer threats. Comparison information 650 includes origin 662, the point at which computing platform 60 has no protection with respect to any of the five types of computer threats that correspond to axes 652, 654, 656, 658, 660. Accordingly, as the capabilities of computing platform 60 are increased to counter a particular type of computer threat, the data point along the corresponding axis is proportionately displaced from origin 652.

[00119] As discussed above, threat mitigation process 10 may obtain 600 consolidated platform information to identify current security-relevant capabilities for computing platform 60. Concerning such current security-relevant capabilities for computing platform 60, these current security-relevant capabilities are defined by data points 664, 666, 668, 670, 672, the combination of which define bounded area 674. Bounded area 674 (in this example) defines the current security-relevant capabilities of computing platform 60.

[00120] Further and as discussed above, threat mitigation process 10 may determine 606 possible security-relevant capabilities for computing platform 60 (i.e., the difference between the current security-relevant capabilities of computing platform 60 and the possible security-relevant capabilities of computing platform 60).

[00121] As discussed above, the possible security-relevant capabilities may concern the possible security-relevant capabilities of computing platform 60 using the currently-deployed security-relevant subsystems. For example, assume that the currently-deployed security relevant subsystems are not currently being utilized to their full potential. Accordingly, certain currently-deployed security relevant subsystems may have certain features that are available but are not utilized and/or disabled. Further, certain currently-deployed security relevant subsystems may have expanded features available if additional licensing fees are paid. Therefore and concerning such possible security-relevant capabilities of computing platform 60 using the currently-deployed security-relevant subsystems, data points 676, 678, 680, 682, 684 may define bounded area 686 (which represents the full capabilities of the currently-deployed security-relevant subsystems within computing platform 60).

[00122] Further and as discussed above, the possible security-relevant capabilities may concern the possible security-relevant capabilities of computing platform 60 using one or more supplemental security-relevant subsystems. For example, assume that supplemental security-relevant subsystems are available for the deployment within computing platform 60. Therefore and concerning such possible security-relevant capabilities of computing platform 60 using such supplemental security-relevant subsystems, data points 688, 690, 692, 694, 696 may define bounded area 698 (which represents the total capabilities of computing platform 60 when utilizing the full

capabilities of the currently-deployed security-relevant subsystems and any supplemental security-relevant subsystems).

[00123] Naturally, the format, appearance and content of comparison information 650 may be varied greatly depending upon the design criteria and anticipated performance / use of threat mitigation process 10. Accordingly, the appearance, format, completeness and content of comparison information 650 is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, content may be added to comparison information 650, removed from comparison information 650, and/or reformatted within comparison information 650.

Concept 6)

[00124] Referring also to FIG 13, threat mitigation process 10 may be configured to generate a threat context score for computing platform 60. As discussed above, threat mitigation process 10 may obtain 600 consolidated platform information to identify current security-relevant capabilities for computing platform 60. This consolidated platform information may be obtained from an independent information source (e.g., such as SIEM system 230 that may provide system-defined consolidated platform information 236) and/or may be obtained from a client information source (e.g., such as questionnaires 240 that may provide client-defined consolidated platform information 238. As will be discussed below in greater detail, threat mitigation process 10 may determine 700 comparative platform information that identifies security-relevant capabilities for a comparative platform, wherein this comparative platform information may concern vendor customers in a specific industry (i.e., the same industry as the user / owner / operator of computing platform 60) and/or vendor customers in any industry (i.e., not necessarily the same industry as the user / owner / operator of computing platform 60).

[00125] Referring also to FIG 14 and as will be discussed below, threat mitigation process 10 may generate 702 comparison information 750 that compares the current security-relevant capabilities of computing platform 60 to the comparative platform information determined 700 for the comparative platform to identify a threat context indicator for computing platform 60, wherein comparison information 750 may include graphical comparison information 752.

[00126] Graphical comparison information 752 (which in this particular example is a bar chart) may identify one or more of: a current threat context score 754 for a client (e.g., the user / owner / operator of computing platform 60); a maximum possible threat context score 756 for the client (e.g., the user / owner / operator of computing platform 60); a threat context score 758 for one or more vendor customers in a specific industry (i.e., the same industry as the user / owner / operator of computing platform 60); and a threat context score 760 for one or more vendor customers in any industry (i.e., not necessarily the same industry as the user / owner / operator of computing platform 60).

[00127] Naturally, the format, appearance and content of comparison information 750 may be varied greatly depending upon the design criteria and anticipated performance / use of threat mitigation process 10. Accordingly, the appearance, format, completeness and content of comparison information 750 is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, content may be added to comparison information 750, removed from comparison information 750, and/or reformatted within comparison information 750.

Computing Platform Monitoring & Mitigation

[00128] As will be discussed below in greater detail, threat mitigation process 10 may be configured to e.g., monitor the operation and performance of computing platform 60.

Concept 7)

[00129] Referring also to FIG. 15, threat mitigation process 10 may be configured to monitor the health of computing platform 60 and provide feedback to a third-party concerning the same. Threat mitigation process 10 may obtain 800 hardware performance information 244 concerning hardware (e.g., server computers, desktop computers, laptop computers, switches, firewalls, routers, gateways, WAPs, and NASs), deployed within computing platform 60. Hardware performance information 244 may concern the operation and/or functionality of one or more hardware systems (e.g., server computers, desktop computers, laptop computers, switches, firewalls, routers, gateways, WAPs, and NASs) deployed within computing platform 60.

[00130] Threat mitigation process 10 may obtain 802 platform performance information 246 concerning the operation of computing platform 60. Platform performance information 246 may concern the operation and/or functionality of computing platform 60.

[00131] When obtaining 802 platform performance information concerning the operation of computing platform 60, threat mitigation process 10 may (as discussed above): obtain 400 system-defined consolidated platform information 236 for computing platform 60 from an independent information source (e.g., SIEM system 230); obtain 312 client-defined consolidated platform information 238 for computing platform 60 from a client information (e.g., questionnaires 240); and present 450 differential consolidated platform information 352 for computing platform 60 to a third-party, examples of which may include but are not limited to the user / owner / operator of computing platform 60.

[00132] When obtaining 802 platform performance information concerning the operation of computing platform 60, threat mitigation process 10 may (as discussed above): obtain 500 consolidated platform information for computing platform 60 to identify one or more deployed security-relevant subsystems 226 (e.g., CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform); process 506 the consolidated platform information (e.g., system-defined consolidated platform information 236 and/or client-defined consolidated platform information 238) to identify one or more non-deployed security-relevant subsystems (within computing platform 60); generate 508 a list of ranked & recommended security-relevant subsystems (e.g., non-deployed security-relevant subsystem list 550) that ranks the one or more non-deployed security-relevant subsystems; and provide 514 the list of ranked & recommended security-relevant subsystems (e.g., non-deployed security-relevant subsystem list 550) to a third-party, examples of which may include but are not limited to a user / owner / operator of computing platform 60.

[00133] When obtaining 802 platform performance information concerning the operation of computing platform 60, threat mitigation process 10 may (as discussed above): obtain 600 consolidated platform information to identify current security-relevant capabilities for the computing platform; determine 606 possible security-relevant capabilities for computing platform 60; and generate 608 comparison information 650 that compares the current security-relevant capabilities of computing platform 60 to the possible security-relevant capabilities of computing platform 60 to identify security-relevant deficiencies.

[00134] When obtaining 802 platform performance information concerning the operation of computing platform 60, threat mitigation process 10 may (as discussed above): obtain 600 consolidated platform information to identify current security-relevant capabilities for computing platform 60; determine 700 comparative platform information that identifies security-relevant capabilities for a comparative platform; and generate 702 comparison information 750 that compares the current security-relevant capabilities of computing platform 60 to the comparative platform information determined 700 for the comparative platform to identify a threat context indicator for computing platform 60.

[00135] Threat mitigation process 10 may obtain 804 application performance information 248 concerning one or more applications (e.g., operating systems, user applications, security application, and utility application) deployed within computing platform 60. Application performance information 248 may concern the operation and/or functionality of one or more software applications (e.g., operating systems, user applications, security application, and utility application) deployed within computing platform 60.

[00136] Referring also to FIG 16, threat mitigation process 10 may generate 806 holistic platform report (e.g., holistic platform reports 850, 852) concerning computing platform 60 based, at least in part, upon hardware performance information 244, platform performance information 246 and application performance information 248. Threat mitigation process 10 may be configured to receive e.g., hardware performance information 244, platform performance information 246 and application performance information 248 at regular intervals (e.g., continuously, every minute, every ten minutes, etc.).

[00137] As illustrated, holistic platform reports 850, 852 may include various pieces of content such as e.g., thought clouds that identify topics / issues with respect to computing platform 60, system logs that memorialize identified issues within computing platform 60, data sources providing information to computing system 60, and so on. The holistic platform report (e.g., holistic platform reports 850, 852) may identify one or more known conditions concerning the computing platform; and threat mitigation process 10 may effectuate 808 one or more remedial operations concerning the one or more known conditions.

[00138] For example, assume that the holistic platform report (e.g., holistic platform reports 850, 852) identifies that computing platform 60 is under a DoS (i.e., Denial of Services) attack. In computing, a denial-of-service attack (DoS attack) is a cyber-attack in which the perpetrator seeks to make a machine or network resource unavailable to its intended users by temporarily or indefinitely disrupting services of a host connected to the Internet. Denial of service is typically accomplished by flooding the targeted machine or resource with superfluous requests in an attempt to overload systems and prevent some or all legitimate requests from being fulfilled.

[00139] In response to detecting such a DoS attack, threat mitigation process 10 may effectuate 808 one or more remedial operations. For example and with respect to such a DoS attack, threat mitigation process 10 may effectuate 808 e.g., a remedial operation that instructs WAF (i.e., Web Application Firewall) 212 to deny all incoming traffic from the identified attacker based upon e.g., protocols, ports or the originating IP addresses.

[00140] Threat mitigation process 10 may also provide 810 the holistic report (e.g., holistic platform reports 850, 852) to a third-party, examples of which may include but are not limited to a user / owner / operator of computing platform 60.

[00141] Naturally, the format, appearance and content of the holistic platform report (e.g., holistic platform reports 850, 852) may be varied greatly depending upon the design criteria and anticipated performance / use of threat mitigation process 10. Accordingly, the appearance, format, completeness and content of the holistic platform report (e.g., holistic platform reports 850, 852) is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, content may be added to the holistic platform report (e.g., holistic platform reports 850,

852), removed from the holistic platform report (e.g., holistic platform reports 850, 852), and/or reformatted within the holistic platform report (e.g., holistic platform reports 850, 852).

Concept 8)

[00142] Referring also to FIG. 17, threat mitigation process 10 may be configured to monitor computing platform 60 for the occurrence of a security event and (in the event of such an occurrence) gather artifacts concerning the same. For example, threat mitigation process 10 may detect 900 a security event within computing platform 60 based upon identified suspect activity. Examples of such security events may include but are not limited to: DDoS events, DoS events, phishing events, spamming events, malware events, web attacks, and exploitation events.

[00143] When detecting 900 a security event (e.g., DDoS events, DoS events, phishing events, spamming events, malware events, web attacks, and exploitation events) within computing platform 60 based upon identified suspect activity, threat mitigation process 10 may monitor 902 a plurality of sources to identify suspect activity within computing platform 60.

[00144] For example, assume that threat mitigation process 10 detects 900 a security event within computing platform 60. Specifically, assume that threat mitigation process 10 is monitoring 902 a plurality of sources (e.g., the various log files maintained by SIEM system 230). And by monitoring 902 such sources, assume that threat mitigation process 10 detects 900 the receipt of inbound content (via an API) from a device having an IP address located in Uzbekistan; the subsequent opening of a port within WAF (i.e., Web Application Firewall) 212; and the streaming of content from a computing device within computing platform 60 through that recently-opened port in WAF (i.e., Web Application Firewall) 212 and to a device having an IP address located in Moldova.

[00145] Upon detecting 900 such a security event within computing platform 60, threat mitigation process 10 may gather 904 artifacts (e.g., artifacts 250) concerning the above-described security event. When gathering 904 artifacts (e.g., artifacts 250) concerning the above-described security event, threat mitigation process 10 may gather 906 artifacts concerning the security event from a plurality of sources associated with the computing platform, wherein examples of such plurality of sources may include but are not limited to the various log files maintained by SIEM system 230, and the various log files directly maintained by the security-relevant subsystems.

[00146] Once the appropriate artifacts (e.g., artifacts 250) are gathered 904, threat mitigation process 10 may assign 908 a threat level to the above-described security event based, at least in part, upon the artifacts (e.g., artifacts 250) gathered 904.

[00147] When assigning 908 a threat level to the above-described security event, threat mitigation process 10 may assign 910 a threat level using artificial intelligence / machine learning. As discussed above and with respect to artificial intelligence / machine learning being utilized to process data sets, an initial probabilistic model may be defined, wherein this initial probabilistic model may be subsequently (e.g., iteratively or continuously) modified and revised, thus allowing the probabilistic models and the artificial intelligence systems (e.g., probabilistic process 56) to “learn” so that future probabilistic models may be more precise and may explain more complex data sets. As further discussed above, probabilistic process 56 may define an initial probabilistic model for accomplishing a defined task (e.g., the analyzing of information 58), wherein the probabilistic model may be utilized to go from initial observations about information 58 (e.g., as represented by the initial branches of a probabilistic model) to conclusions about information 58 (e.g., as represented by the leaves of a probabilistic model). Accordingly and through the use of probabilistic process 56, massive data sets concerning security events may be processed so that a probabilistic model may be defined (and subsequently revised) to assign 910 a threat level to the above-described security event.

[00148] Once assigned 910 a threat level, threat mitigation process 10 may execute 912 a remedial action plan (e.g., remedial action plan 252) based, at least in part, upon the assigned threat level.

[00149] For example and when executing 912 a remedial action plan, threat mitigation process 10 may allow 914 the above-described suspect activity to continue when e.g., threat mitigation process 10 assigns 908 a “low” threat level to the above-described security event (e.g., assuming that it is determined that the user of the local computing device is streaming video of his daughter’s graduation to his parents in Moldova).

[00150] Further and when executing 912 a remedial action plan, threat mitigation process 10 may generate 916 a security event report (e.g., security event report 254) based, at least in part, upon the artifacts (e.g., artifacts 250) gathered 904; and provide 918 the security event report (e.g.,

security event report 254) to an analyst (e.g., analyst 256) for further review when e.g., threat mitigation process 10 assigns 908 a “moderate” threat level to the above-described security event (e.g., assuming that it is determined that while the streaming of the content is concerning, the content is low value and the recipient is not a known bad actor).

[00151] Further and when executing 912 a remedial action plan, threat mitigation process 10 may autonomously execute 920 a threat mitigation plan (shutting down the stream and closing the port) when e.g., threat mitigation process 10 assigns 908 a “severe” threat level to the above-described security event (e.g., assuming that it is determined that the streaming of the content is very concerning, as the content is high value and the recipient is a known bad actor).

[00152] Additionally, threat mitigation process 10 may allow 922 a third-party (e.g., the user / owner / operator of computing platform 60) to manually search for artifacts within computing platform 60. For example, the third-party (e.g., the user / owner / operator of computing platform 60) may be able to search the various information resources include within computing platform 60, examples of which may include but are not limited to the various log files maintained by SIEM system 230, and the various log files directly maintained by the security-relevant subsystems within computing platform 60.

Computing Platform Aggregation & Searching

[00153] As will be discussed below in greater detail, threat mitigation process 10 may be configured to e.g., aggregate data sets and allow for unified search of those data sets.

Concept 9)

[00154] Referring also to FIG. 18, threat mitigation process 10 may be configured to consolidate multiple separate and discrete data sets to form a single, aggregated data set. For example, threat mitigation process 10 may establish 950 connectivity with a plurality of security-relevant subsystems (e.g., security-relevant subsystems 226) within computing platform 60. As discussed above, examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems,

Antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform.

[00155] When establishing 950 connectivity with a plurality of security-relevant subsystems, threat mitigation process 10 may utilize 952 at least one application program interface (e.g., API Gateway 224) to access at least one of the plurality of security-relevant subsystems. For example, a 1st API gateway may be utilized to access CDN (i.e., Content Delivery Network) system; a 2nd API gateway may be utilized to access DAM (i.e., Database Activity Monitoring) system; a 3rd API gateway may be utilized to access UBA (i.e., User Behavior Analytics) system; a 4th API gateway may be utilized to access MDM (i.e., Mobile Device Management) system; a 5th API gateway may be utilized to access IAM (i.e., Identity and Access Management) system; and a 6th API gateway may be utilized to access DNS (i.e., Domain Name Server) system.

[00156] Threat mitigation process 10 may obtain 954 at least one security-relevant information set (e.g., a log file) from each of the plurality of security-relevant subsystems (e.g., CDN system; DAM system; UBA system; MDM system; IAM system; and DNS system), thus defining plurality of security-relevant information sets 258. As would be expected, plurality of security-relevant information sets 258 may utilize a plurality of different formats and/or a plurality of different nomenclatures. Accordingly, threat mitigation process 10 may combine 956 plurality of security-relevant information sets 258 to form an aggregated security-relevant information set 260 for computing platform 60.

[00157] When combining 956 plurality of security-relevant information sets 258 to form aggregated security-relevant information set 260, threat mitigation process 10 may homogenize 958 plurality of security-relevant information sets 258 to form aggregated security-relevant information set 260. For example, threat mitigation process 10 may process one or more of security-relevant information sets 258 so that they all have a common format, a common nomenclature, and/or a common structure.

[00158] Once threat mitigation process 10 combines 956 plurality of security-relevant information sets 258 to form an aggregated security-relevant information set 260 for computing platform 60, threat mitigation process 10 may enable 960 a third-party (e.g., the user / owner / operator of computing platform 60) to access aggregated security-relevant information set 260

and/or enable 962 a third-party (e.g., the user / owner / operator of computing platform 60) to search aggregated security-relevant information set 260.

Concept 10)

[00159] Referring also to FIG 19, threat mitigation process 10 may be configured to enable the searching of multiple separate and discrete data sets using a single search operation. For example and as discussed above, threat mitigation process 10 may establish 950 connectivity with a plurality of security-relevant subsystems (e., security-relevant subsystems 226) within computing platform 60. As discussed above, examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, Antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform.

[00160] When establishing 950 connectivity with a plurality of security-relevant subsystems, threat mitigation process 10 may utilize 952 at least one application program interface (e.g., API Gateway 224) to access at least one of the plurality of security-relevant subsystems. For example, a 1st API gateway may be utilized to access CDN (i.e., Content Delivery Network) system; a 2nd API gateway may be utilized to access DAM (i.e., Database Activity Monitoring) system; a 3rd API gateway may be utilized to access UBA (i.e., User Behavior Analytics) system; a 4th API gateway may be utilized to access MDM (i.e., Mobile Device Management) system; a 5th API gateway may be utilized to access IAM (i.e., Identity and Access Management) system; and a 6th API gateway may be utilized to access DNS (i.e., Domain Name Server) system.

[00161] Threat mitigation process 10 may receive 1000 unified query 262 from a third-party (e.g., the user / owner / operator of computing platform 60) concerning the plurality of security-relevant subsystems. As discussed above, examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e.,

Domain Name Server) systems, Antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform.

[00162] Threat mitigation process 10 may distribute 1002 at least a portion of unified query 262 to the plurality of security-relevant subsystems, resulting in the distribution of plurality of queries 264 to the plurality of security-relevant subsystems. For example, assume that a third-party (e.g., the user / owner / operator of computing platform 60) wishes to execute a search concerning the activity of a specific employee. Accordingly, the third-party (e.g., the user / owner / operator of computing platform 60) may formulate the appropriate unified query (e.g., unified query 262) that defines the employee name, the computing device(s) of the employee, and the date range of interest. Unified query 262 may then be parsed to form plurality of queries 264, wherein a specific query (within plurality of queries 264) may be defined for each of the plurality of security-relevant subsystems and provided to the appropriate security-relevant subsystems. For example, a 1st query may be included within plurality of queries 264 and provided to CDN (i.e., Content Delivery Network) system; a 2nd query may be included within plurality of queries 264 and provided to DAM (i.e., Database Activity Monitoring) system; a 3rd query may be included within plurality of queries 264 and provided to UBA (i.e., User Behavior Analytics) system; a 4th query may be included within plurality of queries 264 and provided to MDM (i.e., Mobile Device Management) system; a 5th query may be included within plurality of queries 264 and provided to IAM (i.e., Identity and Access Management) system; and a 6th query may be included within plurality of queries 264 and provided to DNS (i.e., Domain Name Server) system.

[00163] Threat mitigation process 10 may effectuate 1004 at least a portion of unified query 262 on each of the plurality of security-relevant subsystems to generate plurality of result sets 266. For example, the 1st query may be executed on CDN (i.e., Content Delivery Network) system to produce a 1st result set; the 2nd query may be executed on DAM (i.e., Database Activity Monitoring) system to produce a 2nd result set; the 3rd query may be executed on UBA (i.e., User Behavior Analytics) system to produce a 3rd result set; the 4th query may be executed on MDM (i.e., Mobile Device Management) system to produce a 4th result set; the 5th query may be executed on IAM (i.e.,

Identity and Access Management) system to produce a 5th result set; and the 6th query may be executed on DNS (i.e., Domain Name Server) system to produce a 6th result set.

[00164] Threat mitigation process 10 may receive 1006 plurality of result sets 266 from the plurality of security-relevant subsystems. Threat mitigation process 10 may then combine 1008 plurality of result sets 266 to form unified query result 268. When combining 1008 plurality of result sets 266 to form unified query result 268, threat mitigation process 10 may homogenize 1010 plurality of result sets 266 to form unified query result 268. For example, threat mitigation process 10 may process one or more discrete result sets included within plurality of result sets 266 so that the discrete result sets within plurality of result sets 266 all have a common format, a common nomenclature, and/or a common structure. Threat mitigation process 10 may then provide 1012 unified query result 268 to the third-party (e.g., the user / owner / operator of computing platform 60).

Concept 11)

[00165] Referring also to FIG 20, threat mitigation process 10 may be configured to utilize artificial intelligence / machine learning to automatically consolidate multiple separate and discrete data sets to form a single, aggregated data set. For example and as discussed above, threat mitigation process 10 may establish 950 connectivity with a plurality of security-relevant subsystems (e.g., security-relevant subsystems 226) within computing platform 60. As discussed above, examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, Antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform.

[00166] As discussed above and when establishing 950 connectivity with a plurality of security-relevant subsystems, threat mitigation process 10 may utilize 952 at least one application program interface (e.g., API Gateway 224) to access at least one of the plurality of security-relevant subsystems. For example, a 1st API gateway may be utilized to access CDN (i.e., Content Delivery Network) system; a 2nd API gateway may be utilized to access DAM (i.e., Database Activity

Monitoring) system; a 3rd API gateway may be utilized to access UBA (i.e., User Behavior Analytics) system; a 4th API gateway may be utilized to access MDM (i.e., Mobile Device Management) system; a 5th API gateway may be utilized to access IAM (i.e., Identity and Access Management) system; and a 6th API gateway may be utilized to access DNS (i.e., Domain Name Server) system.

[00167] As discussed above, threat mitigation process 10 may obtain 954 at least one security-relevant information set (e.g., a log file) from each of the plurality of security-relevant subsystems (e.g., CDN system; DAM system; UBA system; MDM system; IAM system; and DNS system), thus defining plurality of security-relevant information sets 258. As would be expected, plurality of security-relevant information sets 258 may utilize a plurality of different formats and/or a plurality of different nomenclatures.

[00168] Threat mitigation process 10 may process 1050 plurality of security-relevant information sets 258 using artificial learning / machine learning to identify one or more commonalities amongst plurality of security-relevant information sets 258. As discussed above and with respect to artificial intelligence / machine learning being utilized to process data sets, an initial probabilistic model may be defined, wherein this initial probabilistic model may be subsequently (e.g., iteratively or continuously) modified and revised, thus allowing the probabilistic models and the artificial intelligence systems (e.g., probabilistic process 56) to “learn” so that future probabilistic models may be more precise and may explain more complex data sets. As further discussed above, probabilistic process 56 may define an initial probabilistic model for accomplishing a defined task (e.g., the analyzing of information 58), wherein the probabilistic model may be utilized to go from initial observations about information 58 (e.g., as represented by the initial branches of a probabilistic model) to conclusions about information 58 (e.g., as represented by the leaves of a probabilistic model). Accordingly and through the use of probabilistic process 56, plurality of security-relevant information sets 258 may be processed so that a probabilistic model may be defined (and subsequently revised) to identify one or more commonalities (e.g., common headers, common nomenclatures, common data ranges, common data types, common formats, etc.) amongst plurality of security-relevant information sets 258. When processing 1050 plurality of security-relevant information sets 258 using artificial learning /

machine learning to identify one or more commonalities amongst plurality of security-relevant information sets 258, threat mitigation process 10 may utilize 1052 a decision tree (e.g., probabilistic model 100) based, at least in part, upon one or more previously-acquired security-relevant information sets.

[00169] Threat mitigation process 10 may combine 1054 plurality of security-relevant information sets 258 to form aggregated security-relevant information set 260 for computing platform 60 based, at least in part, upon the one or more commonalities identified.

[00170] When combining 1054 plurality of security-relevant information sets 258 to form aggregated security-relevant information set 260 for computing platform 60 based, at least in part, upon the one or more commonalities identified, threat mitigation process 10 may homogenize 1056 plurality of security-relevant information sets 258 to form aggregated security-relevant information set 260. For example, threat mitigation process 10 may process one or more of security-relevant information sets 258 so that they all have a common format, a common nomenclature, and/or a common structure.

[00171] Once threat mitigation process 10 combines 1054 plurality of security-relevant information sets 258 to form an aggregated security-relevant information set 260 for computing platform 60, threat mitigation process 10 may enable 1058 a third-party (e.g., the user / owner / operator of computing platform 60) to access aggregated security-relevant information set 260 and/or enable 1060 a third-party (e.g., the user / owner / operator of computing platform 60) to search aggregated security-relevant information set 260.

Threat Event Information Updating

[00172] As will be discussed below in greater detail, threat mitigation process 10 may be configured to be updated concerning threat event information.

Concept 12)

[00173] Referring also to FIG. 21, threat mitigation process 10 may be configured to receive updated threat event information for security-relevant subsystems 226. For example, threat mitigation process 10 may receive 1100 updated threat event information 270 concerning computing platform 60, wherein updated threat event information 270 may define one or more of: updated threat listings; updated threat definitions; updated threat methodologies; updated threat

sources; and updated threat strategies. Threat mitigation process 10 may enable 1102 updated threat event information 270 for use with one or more security-relevant subsystems 226 within computing platform 60. As discussed above, examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, Antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform.

[00174] When enabling 1102 updated threat event information 270 for use with one or more security-relevant subsystems 226 within computing platform 60, threat mitigation process 10 may install 1104 updated threat event information 270 on one or more security-relevant subsystems 226 within computing platform 60.

[00175] Threat mitigation process 10 may retroactively apply 1106 updated threat event information 270 to previously-generated information associated with one or more security-relevant subsystems 226.

[00176] When retroactively apply 1106 updated threat event information 270 to previously-generated information associated with one or more security-relevant subsystems 226, threat mitigation process 10 may: apply 1108 updated threat event information 270 to one or more previously-generated log files (not shown) associated with one or more security-relevant subsystems 226; apply 1110 updated threat event information 270 to one or more previously-generated data files (not shown) associated with one or more security-relevant subsystems 226; and apply 1112 updated threat event information 270 to one or more previously-generated application files (not shown) associated with one or more security-relevant subsystems 226.

[00177] Additionally, / alternatively, threat mitigation process 10 may proactively apply 1114 updated threat event information 270 to newly-generated information associated with one or more security-relevant subsystems 226.

[00178] When proactively applying 1114 updated threat event information 270 to newly-generated information associated with one or more security-relevant subsystems 226, threat

mitigation process 10 may: apply 1116 updated threat event information 270 to one or more newly-generated log files (not shown) associated with one or more security-relevant subsystems 226; apply 1118 updated threat event information 270 to one or more newly-generated data files (not shown) associated with one or more security-relevant subsystems 226; and apply 1120 updated threat event information 270 to one or more newly-generated application files (not shown) associated with one or more security-relevant subsystems 226.

Concept 13)

[00179] Referring also to FIG 22, threat mitigation process 10 may be configured to receive updated threat event information 270 for security-relevant subsystems 226. For example and as discussed above, threat mitigation process 10 may receive 1100 updated threat event information 270 concerning computing platform 60, wherein updated threat event information 270 may define one or more of: updated threat listings; updated threat definitions; updated threat methodologies; updated threat sources; and updated threat strategies. Further and as discussed above, threat mitigation process 10 may enable 1102 updated threat event information 270 for use with one or more security-relevant subsystems 226 within computing platform 60. As discussed above, examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, Antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform.

[00180] As discussed above and when enabling 1102 updated threat event information 270 for use with one or more security-relevant subsystems 226 within computing platform 60, threat mitigation process 10 may install 1104 updated threat event information 270 on one or more security-relevant subsystems 226 within computing platform 60.

[00181] Sometimes, it may not be convenient and/or efficient to immediately apply updated threat event information 270 to security-relevant subsystems 226. Accordingly, threat mitigation process 10 may schedule 1150 the application of updated threat event information 270 to previously-generated information associated with one or more security-relevant subsystems 226.

[00182] When scheduling 1150 the application of updated threat event information 270 to previously-generated information associated with one or more security-relevant subsystems 226, threat mitigation process 10 may: schedule 1152 the application of updated threat event information 270 to one or more previously-generated log files (not shown) associated with one or more security-relevant subsystems 226; schedule 1154 the application of updated threat event information 270 to one or more previously-generated data files (not shown) associated with one or more security-relevant subsystems 226; and schedule 1156 the application of updated threat event information 270 to one or more previously-generated application files (not shown) associated with one or more security-relevant subsystems 226.

[00183] Additionally, / alternatively, threat mitigation process 10 may schedule 1158 the application of the updated threat event information to newly-generated information associated with the one or more security-relevant subsystems.

[00184] When scheduling 1158 the application of updated threat event information 270 to newly-generated information associated with one or more security-relevant subsystems 226, threat mitigation process 10 may: schedule 1160 the application of updated threat event information 270 to one or more newly-generated log files (not shown) associated with one or more security-relevant subsystems 226; schedule 1162 the application of updated threat event information 270 to one or more newly-generated data files (not shown) associated with one or more security-relevant subsystems 226; and schedule 1164 the application of updated threat event information 270 to one or more newly-generated application files (not shown) associated with one or more security-relevant subsystems 226.

Concept 14)

[00185] Referring also to FIGS. 23-24, threat mitigation process 10 may be configured to initially display analytical data, which may then be manipulated / updated to include automation data. For example, threat mitigation process 10 may display 1200 initial security-relevant information 1250 that includes analytical information (e.g., thought cloud 1252). Examples of such analytical information may include but is not limited to one or more of: investigative information; and hunting information.

[00186] **Investigative Information** (a portion of analytical information): Unified searching and/or automated searching, such as e.g., a security event occurring and searches being performed to gather artifacts concerning that security event.

[00187] **Hunt Information** (a portion of analytical information): Targeted searching / investigations, such as the monitoring and cataloging of the videos that an employee has watched or downloaded over the past 30 days.

[00188] Threat mitigation process 10 may allow 1202 a third-party (e.g., the user / owner / operator of computing platform 60) to manipulate initial security-relevant information 1250 with automation information.

[00189] **Automate Information** (a portion of automation): The execution of a single (and possibly simple) action one time, such as the blocking an IP address from accessing computing platform 60 whenever such an attempt is made.

[00190] **Orchestrate Information** (a portion of automation): The execution of a more complex batch (or series) of tasks, such as sensing an unauthorized download via an API and a) shutting down the API, adding the requesting IP address to a blacklist, and closing any ports opened for the requestor.

[00191] When allowing 1202 a third-party (e.g., the user / owner / operator of computing network 60) to manipulate initial security-relevant information 1250 with automation information, threat mitigation process 10 may allow 1204 a third-party (e.g., the user / owner / operator of computing network 60) to select the automation information to add to initial security-relevant information 1250 to generate revised security-relevant information 1250'. For example and when allowing 1204 a third-party (e.g., the user / owner / operator of computing network 60) to select the automation information to add to initial security-relevant information 1250 to generate revised security-relevant information 1250', threat mitigation process 10 may allow 1206 the third-party (e.g., the user / owner / operator of computing network 60) to choose a specific type of automation information from a plurality of automation information types.

[00192] For example, the third-party (e.g., the user / owner / operator of computing network 60) may choose to add / initiate the automation information to generate revised security-relevant information 1250'. Accordingly, threat mitigation process 10 may render selectable options (e.g.,

selectable buttons 1254, 1256) that the third-party (e.g., the user / owner / operator of computing network 60) may select to manipulate initial security-relevant information 1250 with automation information to generate revised security-relevant information 1250'. For this particular example, the third-party (e.g., the user / owner / operator of computing network 60) may choose two different options to manipulate initial security-relevant information 1250, namely: "block ip" or "search", both of which will result in threat mitigation process 10 generating 1208 revised security-relevant information 1250' (that includes the above-described automation information).

[00193] When generating 1208 revised security-relevant information 1250' (that includes the above-described automation information), threat mitigation process 10 may combine 1210 the automation information (that results from selecting "block IP" or "search") and initial security-relevant information 1250 to generate and render 1212 revised security-relevant information 1250'.

[00194] When rendering 1212 revised security-relevant information 1250', threat mitigation process 10 may render 1214 revised security-relevant information 1250' within interactive report 1258.

Training Routine Generation and Execution

[00195] As will be discussed below in greater detail, threat mitigation process 10 may be configured to allow for the manual or automatic generation of training routines, as well as the execution of the same.

Concept 15)

[00196] Referring also to FIG 25, threat mitigation process 10 may be configured to allow for the manual generation of testing routine 272. For example, threat mitigation process 10 may define 1300 training routine 272 for a specific attack (e.g., a Denial of Services attack) of computing platform 60. Specifically, threat mitigation process 10 may generate 1302 a simulation of the specific attack (e.g., a Denial of Services attack) by executing training routine 272 within a controlled test environment, an example of which may include but is not limited to virtual machine 274 executed on a computing device (e.g., computing device 12).

[00197] When generating 1302 a simulation of the specific attack (e.g., a Denial of Services attack) by executing training routine 272 within the controlled test environment (e.g., virtual

machine 274), threat mitigation process 10 may render 1304 the simulation of the specific attack (e.g., a Denial of Services attack) on the controlled test environment (e.g., virtual machine 274).

[00198] Threat mitigation process 10 may allow 1306 a trainee (e.g., trainee 276) to view the simulation of the specific attack (e.g., a Denial of Services attack) and may allow 1308 the trainee (e.g., trainee 276) to provide a trainee response (e.g., trainee response 278) to the simulation of the specific attack (e.g., a Denial of Services attack). For example, threat mitigation process 10 may execute training routine 272, which trainee 276 may “watch” and provide trainee response 278.

[00199] Threat mitigation process 10 may then determine 1310 the effectiveness of trainee response 278, wherein determining 1310 the effectiveness of the trainee response may include threat mitigation process 10 assigning 1312 a grade (e.g., a letter grade or a number grade) to trainee response 278.

Concept 16)

[00200] Referring also to FIG 26, threat mitigation process 10 may be configured to allow for the automatic generation of testing routine 272. For example, threat mitigation process 10 may utilize 1350 artificial intelligence / machine learning to define training routine 272 for a specific attack (e.g., a Denial of Services attack) of computing platform 60.

[00201] As discussed above and with respect to artificial intelligence / machine learning being utilized to process data sets, an initial probabilistic model may be defined, wherein this initial probabilistic model may be subsequently (e.g., iteratively or continuously) modified and revised, thus allowing the probabilistic models and the artificial intelligence systems (e.g., probabilistic process 56) to “learn” so that future probabilistic models may be more precise and may explain more complex data sets. As further discussed above, probabilistic process 56 may define an initial probabilistic model for accomplishing a defined task (e.g., the analyzing of information 58), wherein the probabilistic model may be utilized to go from initial observations about information 58 (e.g., as represented by the initial branches of a probabilistic model) to conclusions about information 58 (e.g., as represented by the leaves of a probabilistic model). Accordingly and through the use of probabilistic process 56, information may be processed so that a probabilistic model may be defined (and subsequently revised) to define training routine 272 for a specific attack (e.g., a Denial of Services attack) of computing platform 60.

[00202] When using 1350 artificial intelligence / machine learning to define training routine 272 for a specific attack (e.g., a Denial of Services attack) of computing platform 60, threat mitigation process 10 may process 1352 security-relevant information to define training routine 272 for specific attack (e.g., a Denial of Services attack) of computing platform 60. Further and when using 1350 artificial intelligence / machine learning to define training routine 272 for a specific attack (e.g., a Denial of Services attack) of computing platform 60, threat mitigation process 10 may utilize 1354 security-relevant rules to define training routine 272 for a specific attack (e.g., a Denial of Services attack) of computing platform 60. Accordingly, security-relevant information that e.g., defines the symptoms of e.g., a Denial of Services attack and security-relevant rules that define the behavior of e.g., a Denial of Services attack may be utilized by threat mitigation process 10 when defining training routine 272.

[00203] As discussed above, threat mitigation process 10 may generate 1302 a simulation of the specific attack (e.g., a Denial of Services attack) by executing training routine 272 within a controlled test environment, an example of which may include but is not limited to virtual machine 274 executed on a computing device (e.g., computing device 12).

[00204] Further and as discussed above, when generating 1302 a simulation of the specific attack (e.g., a Denial of Services attack) by executing training routine 272 within the controlled test environment (e.g., virtual machine 274), threat mitigation process 10 may render 1304 the simulation of the specific attack (e.g., a Denial of Services attack) on the controlled test environment (e.g., virtual machine 274).

[00205] Threat mitigation process 10 may allow 1306 a trainee (e.g., trainee 276) to view the simulation of the specific attack (e.g., a Denial of Services attack) and may allow 1308 the trainee (e.g., trainee 276) to provide a trainee response (e.g., trainee response 278) to the simulation of the specific attack (e.g., a Denial of Services attack). For example, threat mitigation process 10 may execute training routine 272, which trainee 276 may “watch” and provide trainee response 278.

[00206] Threat mitigation process 10 may utilize 1356 artificial intelligence / machine learning to revise training routine 272 for the specific attack (e.g., a Denial of Services attack) of computing platform 60 based, at least in part, upon trainee response 278.

[00207] As discussed above, threat mitigation process 10 may then determine 1310 the effectiveness of trainee response 278, wherein determining 1310 the effectiveness of the trainee response may include threat mitigation process 10 assigning 1312 a grade (e.g., a letter grade or a number grade) to trainee response 278.

Concept 17)

[00208] Referring also to FIG 27, threat mitigation process 10 may be configured to allow a trainee to choose their training routine. For example mitigation process 10 may allow 1400 a third-party (e.g., the user / owner / operator of computing network 60) to select a training routine for a specific attack (e.g., a Denial of Services attack) of computing platform 60, thus defining a selected training routine. When allowing 1400 a third-party (e.g., the user / owner / operator of computing network 60) to select a training routine for a specific attack (e.g., a Denial of Services attack) of computing platform 60, threat mitigation process 10 may allow 1402 the third-party (e.g., the user / owner / operator of computing network 60) to choose a specific training routine from a plurality of available training routines. For example, the third-party (e.g., the user / owner / operator of computing network 60) may be able to select a specific type of attack (e.g., DDoS events, DoS events, phishing events, spamming events, malware events, web attacks, and exploitation events) and/or select a specific training routine (that may or may not disclose the specific type of attack).

[00209] Once selected, threat mitigation process 10 may analyze 1404 the requirements of the selected training routine (e.g., training routine 272) to determine a quantity of entities required to effectuate the selected training routine (e.g., training routine 272), thus defining one or more required entities. For example, assume that training routine 272 has three required entities (e.g., an attacked device and two attacking devices). According, threat mitigation process 10 may generate 1406 one or more virtual machines (e.g., such as virtual machine 274) to emulate the one or more required entities. In this particular example, threat mitigation process 10 may generate 1406 three virtual machines, a first VM for the attacked device, a second VM for the first attacking device and a third VM for the second attacking device. As is known in the art, a virtual machine (VM) is an virtual emulation of a physical computing system. Virtual machines may be based on computer architectures and may provide the functionality of a physical computer, wherein their implementations may involve specialized hardware, software, or a combination thereof.

[00210] Threat mitigation process 10 may generate 1408 a simulation of the specific attack (e.g., a Denial of Services attack) by executing the selected training routine (e.g., training routine 272). When generating 1408 the simulation of the specific attack (e.g., a Denial of Services attack) by executing the selected training routine (e.g., training routine 272), threat mitigation process 10 may render 1410 the simulation of the specific attack (e.g., a Denial of Services attack) by executing the selected training routine (e.g., training routine 272) within a controlled test environment (e.g., such as virtual machine 274).

[00211] As discussed above, threat mitigation process 10 may allow 1306 a trainee (e.g., trainee 276) to view the simulation of the specific attack (e.g., a Denial of Services attack) and may allow 1308 the trainee (e.g., trainee 276) to provide a trainee response (e.g., trainee response 278) to the simulation of the specific attack (e.g., a Denial of Services attack). For example, threat mitigation process 10 may execute training routine 272, which trainee 276 may “watch” and provide trainee response 278.

[00212] Further and as discussed above, threat mitigation process 10 may then determine 1310 the effectiveness of trainee response 278, wherein determining 1310 the effectiveness of the trainee response may include threat mitigation process 10 assigning 1312 a grade (e.g., a letter grade or a number grade) to trainee response 278.

[00213] When training is complete, threat mitigation process 10 may cease 1412 the simulation of the specific attack (e.g., a Denial of Services attack), wherein ceasing 1412 the simulation of the specific attack (e.g., a Denial of Services attack) may include threat mitigation process 10 shutting down 1414 the one or more virtual machines (e.g., the first VM for the attacked device, the second VM for the first attacking device and the third VM for the second attacking device).

Information Routing

[00214] As will be discussed below in greater detail, threat mitigation process 10 may be configured to route information based upon whether the information is more threat-pertinent or less threat-pertinent.

Concept 18)

[00215] Referring also to FIG 28, threat mitigation process 10 may be configured to route more threat-pertinent content in a specific manner. For example, threat mitigation process 10 may receive 1450 platform information (e.g., log files) from a plurality of security-relevant subsystems (e.g., security-relevant subsystems 226). As discussed above, examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, Antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform.

[00216] Threat mitigation process 10 may process 1452 this platform information (e.g., log files) to generate processed platform information. And when processing 1452 this platform information (e.g., log files) to generate processed platform information, threat mitigation process 10 may: parse 1454 the platform information (e.g., log files) into a plurality of subcomponents (e.g., columns, rows, etc.) to allow for compensation of varying formats and/or nomenclature; enrich 1456 the platform information (e.g., log files) by including supplemental information from external information resources; and/or utilize 1458 artificial intelligence / machine learning (in the manner described above) to identify one or more patterns / trends within the platform information (e.g., log files).

[00217] Threat mitigation process 10 may identify 1460 more threat-pertinent content 280 included within the processed content, wherein identifying 1460 more threat-pertinent content 280 included within the processed content may include processing 1462 the processed content to identify actionable processed content that may be used by a threat analysis engine (e.g., SIEM

system 230) for correlation purposes. Threat mitigation process 10 may route 1464 more threat-pertinent content 280 to this threat analysis engine (e.g., SIEM system 230).

Concept 19)

[00218] Referring also to FIG. 29, threat mitigation process 10 may be configured to route less threat-pertinent content in a specific manner. For example and as discussed above, threat mitigation process 10 may receive 1450 platform information (e.g., log files) from a plurality of security-relevant subsystems (e.g., security-relevant subsystems 226). As discussed above, examples of security-relevant subsystems 226 may include but are not limited to: CDN (i.e., Content Delivery Network) systems; DAM (i.e., Database Activity Monitoring) systems; UBA (i.e., User Behavior Analytics) systems; MDM (i.e., Mobile Device Management) systems; IAM (i.e., Identity and Access Management) systems; DNS (i.e., Domain Name Server) systems, Antivirus systems, operating systems, data lakes; data logs; security-relevant software applications; security-relevant hardware systems; and resources external to the computing platform

[00219] Further and as discussed above, threat mitigation process 10 may process 1452 this platform information (e.g., log files) to generate processed platform information. And when processing 1452 this platform information (e.g., log files) to generate processed platform information, threat mitigation process 10 may: parse 1454 the platform information (e.g., log files) into a plurality of subcomponents (e.g., columns, rows, etc.) to allow for compensation of varying formats and/or nomenclature; enrich 1456 the platform information (e.g., log files) by including supplemental information from external information resources; and/or utilize 1458 artificial intelligence / machine learning (in the manner described above) to identify one or more patterns / trends within the platform information (e.g., log files).

[00220] Threat mitigation process 10 may identify 1500 less threat-pertinent content 282 included within the processed content, wherein identifying 1500 less threat-pertinent content 282 included within the processed content may include processing 1502 the processed content to identify non-actionable processed content that is not usable by a threat analysis engine (e.g., SIEM system 230) for correlation purposes. Threat mitigation process 10 may route 1504 less threat-pertinent content 282 to a long term storage system (e.g., long term storage system 284). Further,

threat mitigation process 10 may be configured to allow 1506 a third-party (e.g., the user / owner / operator of computing network 60) to access and search long term storage system 284.

Automated Analysis

[00221] As will be discussed below in greater detail, threat mitigation process 10 may be configured to automatically analyze a detected security event.

Concept 20)

[00222] Referring also to FIG. 30, threat mitigation process 10 may be configured to automatically classify and investigate a detected security event. As discussed above and in response to a security event being detected, threat mitigation process 10 may obtain 1550 one or more artifacts (e.g., artifacts 250) concerning the detected security event. Examples of such a detected security event may include but are not limited to one or more of: access auditing; anomalies; authentication; denial of services; exploitation; malware; phishing; spamming; reconnaissance; and web attack. These artifacts (e.g., artifacts 250) may be obtained 1550 from a plurality of sources associated with the computing platform, wherein examples of such plurality of sources may include but are not limited to the various log files maintained by SIEM system 230, and the various log files directly maintained by the security-relevant subsystems

[00223] Threat mitigation process 10 may obtain 1552 artifact information (e.g., artifact information 286) concerning the one or more artifacts (e.g., artifacts 250), wherein artifact information 286 may be obtained from information resources include within (or external to) computing platform 60.

[00224] For example and when obtaining 1552 artifact information 286 concerning the one or more artifacts (e.g., artifacts 250), threat mitigation process 10 may obtain 1554 artifact information 286 concerning the one or more artifacts (e.g., artifacts 250) from one or more investigation resources (such as third-party resources that may e.g., provide information on known bad actors).

[00225] Once the investigation is complete, threat mitigation process 10 may generate 1556 a conclusion (e.g., conclusion 288) concerning the detected security event (e.g., a Denial of Services attack) based, at least in part, upon the detected security event (e.g., a Denial of Services attack), the one or more artifacts (e.g., artifacts 250), and artifact information 286. Threat mitigation process 10

may document 1558 the conclusion (e.g., conclusion 288), report 1560 the conclusion (e.g., conclusion 288) to a third-party (e.g., the user / owner / operator of computing network 60). Further, threat mitigation process 10 may obtain 1562 supplemental artifacts and artifact information (if needed to further the investigation).

[00226] While the system is described above as being computer-implemented, this is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example, some or all of the above-described system may be implemented by a human being.

General

[00227] As will be appreciated by one skilled in the art, the present disclosure may be embodied as a method, a system, or a computer program product. Accordingly, the present disclosure may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a “circuit,” “module” or “system.” Furthermore, the present disclosure may take the form of a computer program product on a computer-usable storage medium having computer-usable program code embodied in the medium.

[00228] Any suitable computer usable or computer readable medium may be utilized. The computer-usable or computer-readable medium may be, for example but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, device, or propagation medium. More specific examples (a non-exhaustive list) of the computer-readable medium may include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a transmission media such as those supporting the Internet or an intranet, or a magnetic storage device. The computer-usable or computer-readable medium may also be paper or another suitable medium upon which the program is printed, as the program can be electronically captured, via, for instance, optical scanning of the paper or other medium, then compiled, interpreted, or otherwise processed in a suitable manner, if necessary, and then stored in a computer memory. In the context of this document, a

computer-usable or computer-readable medium may be any medium that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, or device. The computer-usable medium may include a propagated data signal with the computer-usable program code embodied therewith, either in baseband or as part of a carrier wave. The computer usable program code may be transmitted using any appropriate medium, including but not limited to the Internet, wireline, optical fiber cable, RF, etc.

[00229] Computer program code for carrying out operations of the present disclosure may be written in an object oriented programming language such as Java, Smalltalk, C++ or the like. However, the computer program code for carrying out operations of the present disclosure may also be written in 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 a local area network / a wide area network / the Internet (e.g., network 14).

[00230] The present disclosure is described with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the disclosure. 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, may be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer / special purpose computer / other programmable data processing apparatus, 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.

[00231] These computer program instructions may also be stored in a computer-readable memory that may direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable memory produce an

article of manufacture including instruction means which implement the function/act specified in the flowchart and/or block diagram block or blocks.

[00232] The computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer or other programmable apparatus to produce a computer-implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[00233] The flowcharts and block diagrams in the figures may illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present disclosure. 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 illustrations, and combinations of blocks in the block diagrams and/or flowchart illustrations, may be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[00234] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the disclosure. 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" and/or "comprising," 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.

[00235] 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 disclosure has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the disclosure 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 disclosure. The embodiment was chosen and described in order to best explain the principles of the disclosure and the practical application, and to enable others of ordinary skill in the art to understand the disclosure for various embodiments with various modifications as are suited to the particular use contemplated.

[00236] A number of implementations have been described. Having thus described the disclosure of the present application in detail and by reference to embodiments thereof, it will be apparent that modifications and variations are possible without departing from the scope of the disclosure defined in the appended claims.

What Is Claimed Is:**Concept 3)**

1. A computer-implemented method, executed on a computing device, comprising:
 - obtaining system-defined consolidated platform information for a computing platform from an independent information source;
 - obtaining client-defined consolidated platform information for the computing platform from a client information source; and
 - presenting differential consolidated platform information for the computing platform to the third-party.
2. The computer-implemented method of claim 1 further comprising:
 - comparing the system-defined consolidated platform information to the client-defined consolidated platform information to define the differential consolidated platform information for the computing platform.
3. The computer-implemented method of claim 2 further comprising:
 - processing the system-defined consolidated platform information prior to comparing the system-defined consolidated platform information to the client-defined consolidated platform information to define differential consolidated platform information for the computing platform.
4. The computer-implemented method of claim 1 wherein the system-defined consolidated platform information is obtained from one or more log files defined for the computing platform.
5. The computer-implemented method of claim 1 wherein the system-defined consolidated platform information is obtained from two or more security-relevant subsystems deployed within the computing platform.

6. The computer-implemented method of claim 1 wherein the client-defined consolidated platform information is obtained from one or more client-completed questionnaires.
7. The computer-implemented method of claim 1 wherein the client-defined consolidated platform information is obtained from one or more client-deployed platform monitors.
8. A computer program product residing on a computer readable medium having a plurality of instructions stored thereon which, when executed by a processor, cause the processor to perform operations comprising:
 - obtaining system-defined consolidated platform information for a computing platform from an independent information source;
 - obtaining client-defined consolidated platform information for the computing platform from a client information source; and
 - presenting differential consolidated platform information for the computing platform to the third-party.
9. The computer program product of claim 8 further comprising:
 - comparing the system-defined consolidated platform information to the client-defined consolidated platform information to define the differential consolidated platform information for the computing platform.
10. The computer program product of claim 9 further comprising:
 - processing the system-defined consolidated platform information prior to comparing the system-defined consolidated platform information to the client-defined consolidated platform information to define differential consolidated platform information for the computing platform.
11. The computer program product of claim 8 wherein the system-defined consolidated platform information is obtained from one or more log files defined for the computing platform.

12. The computer program product of claim 8 wherein the system-defined consolidated platform information is obtained from two or more security-relevant subsystems deployed within the computing platform.
13. The computer program product of claim 8 wherein the client-defined consolidated platform information is obtained from one or more client-completed questionnaires.
14. The computer program product of claim 8 wherein the client-defined consolidated platform information is obtained from one or more client-deployed platform monitors.
15. A computing system including a processor and memory configured to perform operations comprising:
 - obtaining system-defined consolidated platform information for a computing platform from an independent information source;
 - obtaining client-defined consolidated platform information for the computing platform from a client information source; and
 - presenting differential consolidated platform information for the computing platform to the third-party.
16. The computing system of claim 15 further comprising:
 - comparing the system-defined consolidated platform information to the client-defined consolidated platform information to define the differential consolidated platform information for the computing platform.
17. The computing system of claim 16 further comprising:
 - processing the system-defined consolidated platform information prior to comparing the system-defined consolidated platform information to the client-defined consolidated

platform information to define differential consolidated platform information for the computing platform.

18. The computing system of claim 15 wherein the system-defined consolidated platform information is obtained from one or more log files defined for the computing platform.

19. The computing system of claim 15 wherein the system-defined consolidated platform information is obtained from two or more security-relevant subsystems deployed within the computing platform.

20. The computing system of claim 15 wherein the client-defined consolidated platform information is obtained from one or more client-completed questionnaires.

21. The computing system of claim 15 wherein the client-defined consolidated platform information is obtained from one or more client-deployed platform monitors.

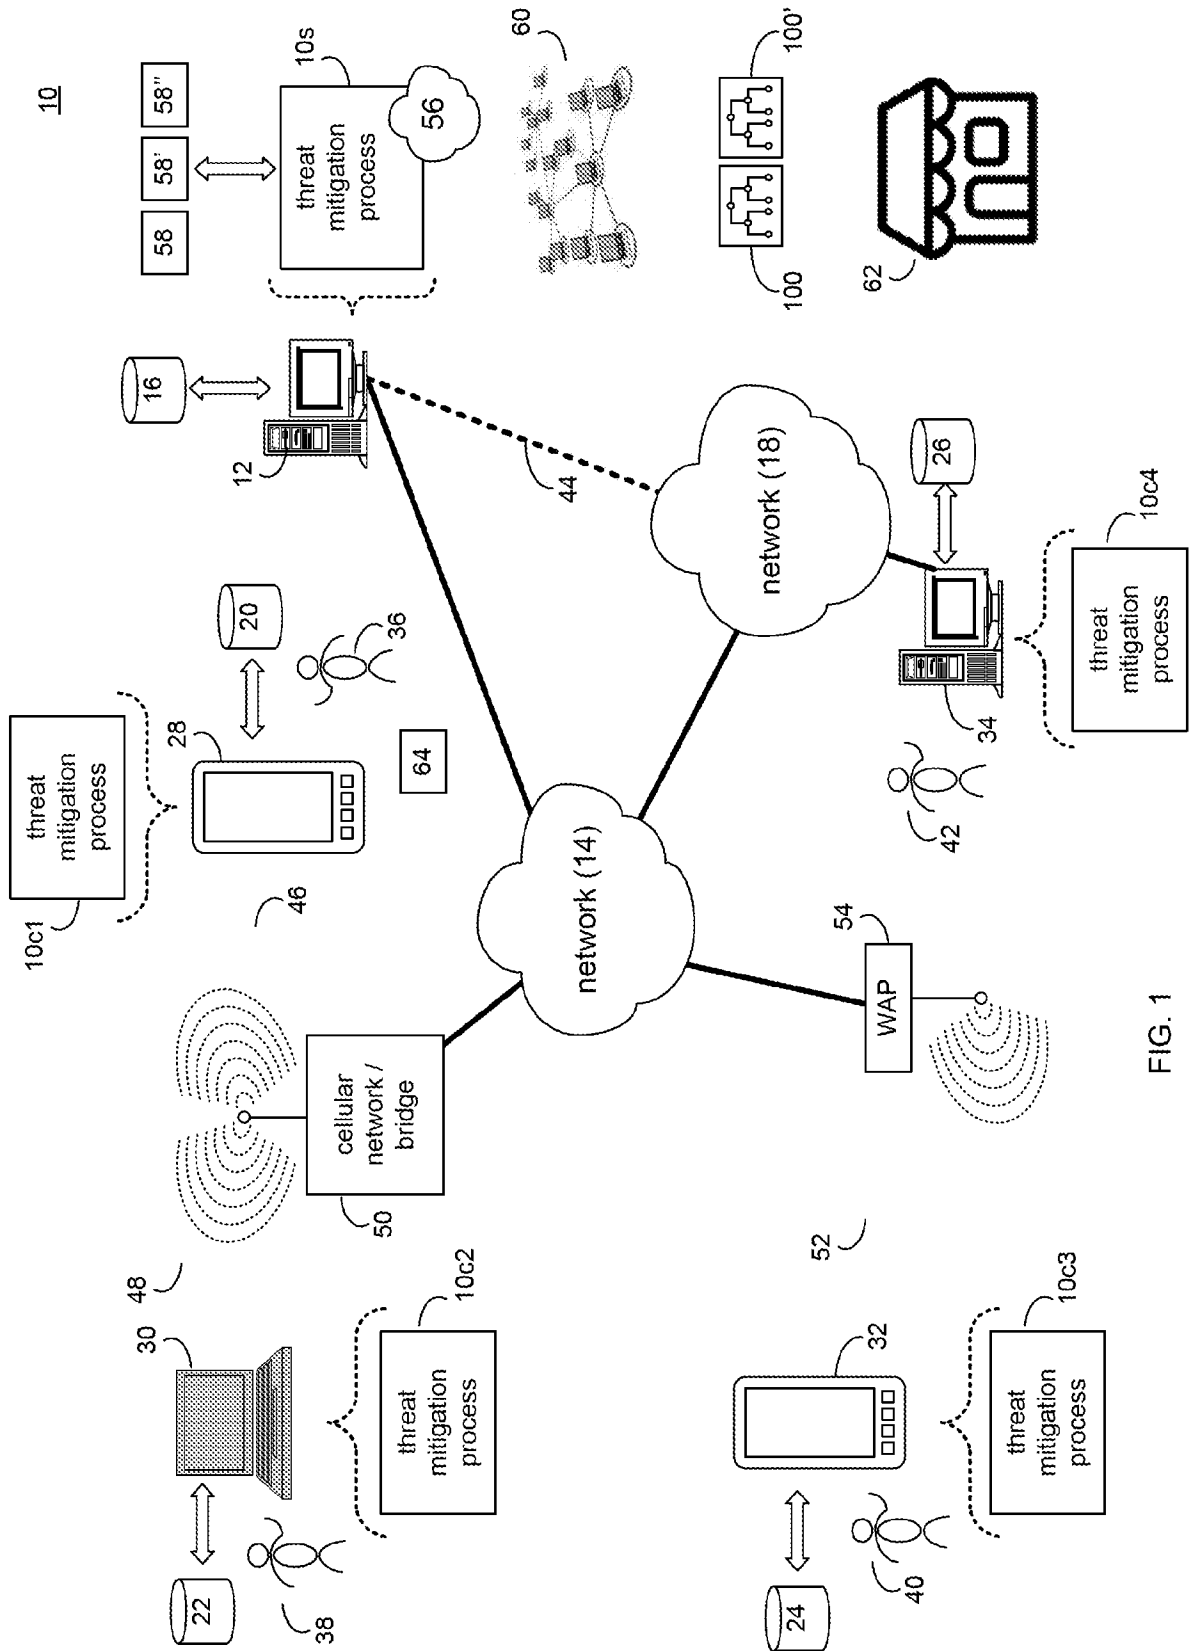


FIG. 1

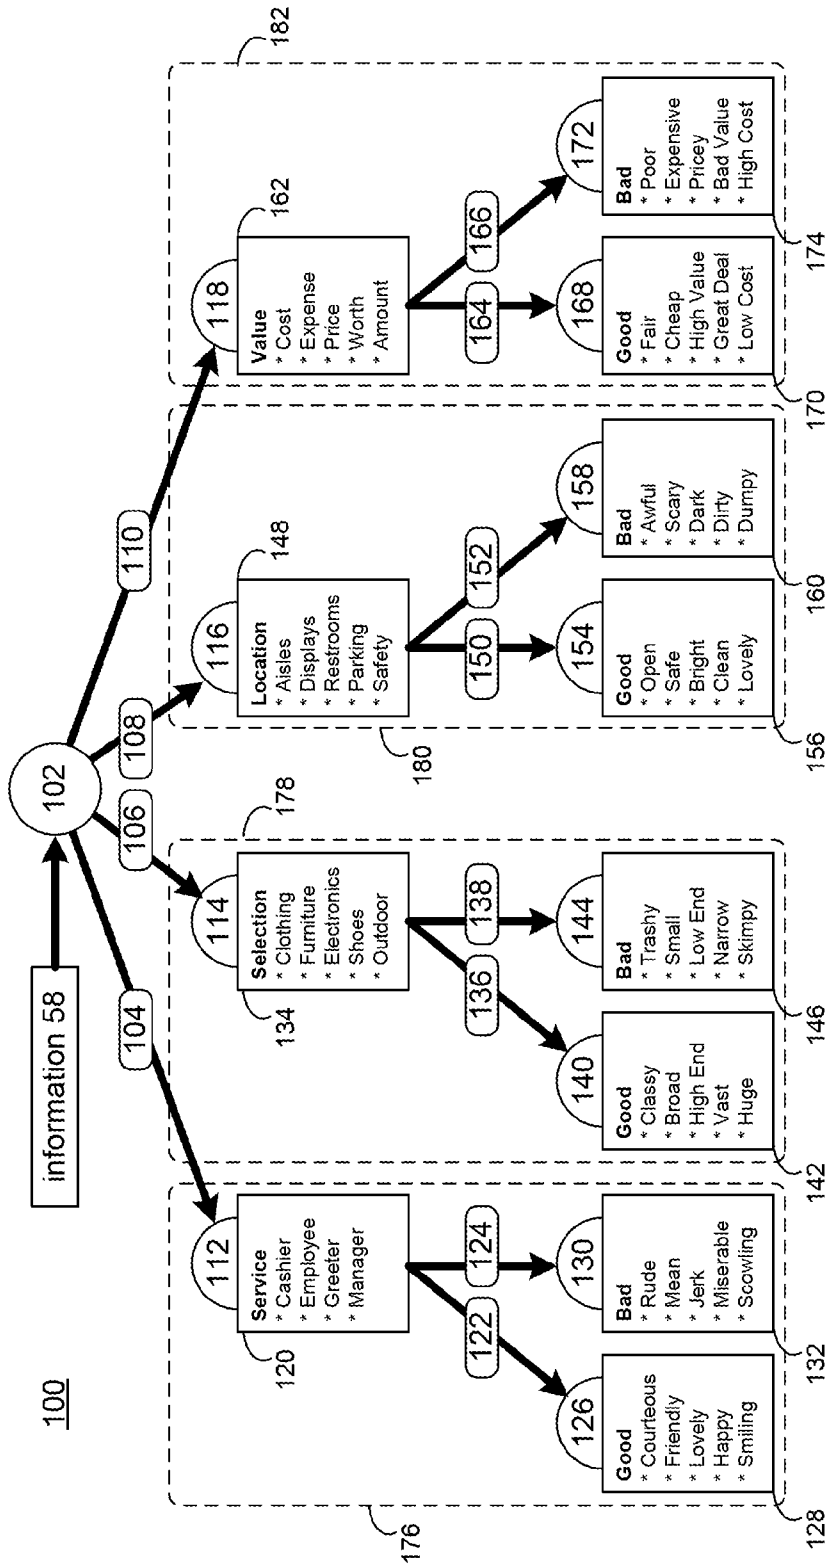


FIG. 2

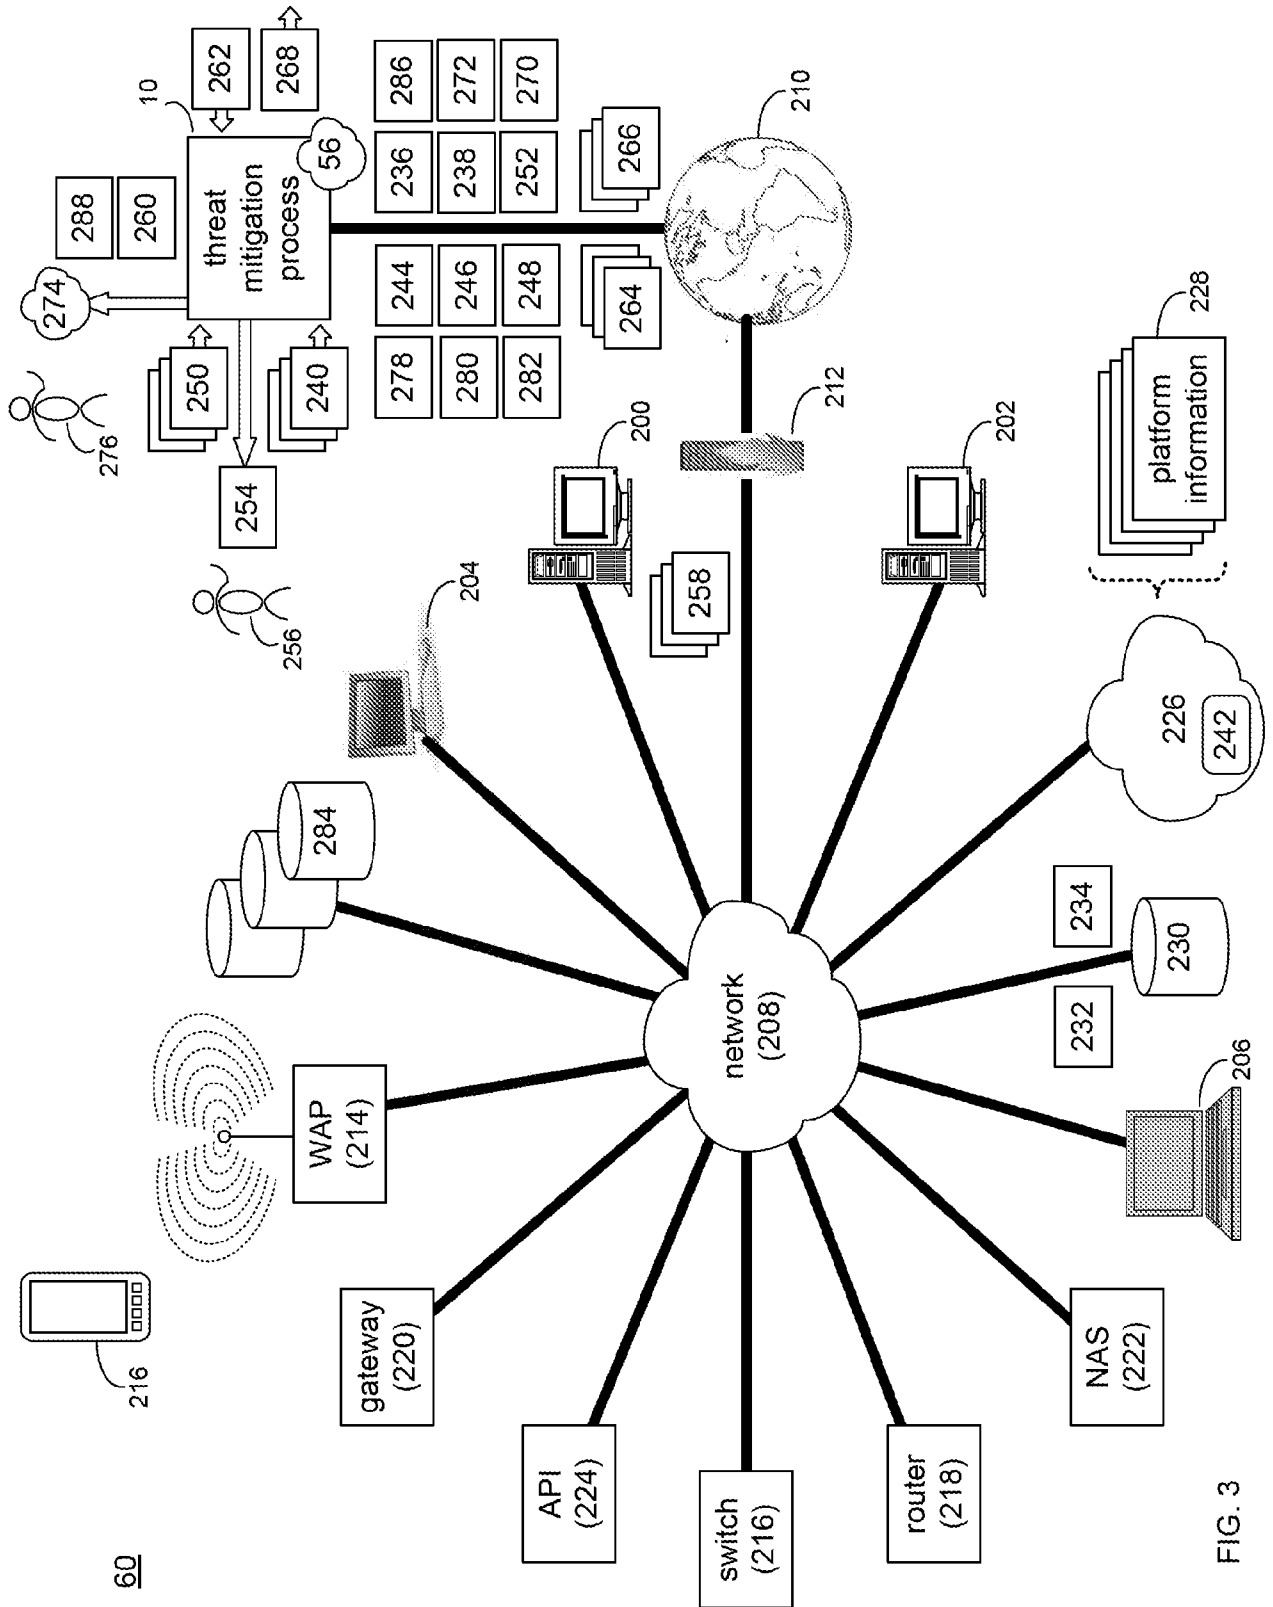


FIG. 3

10

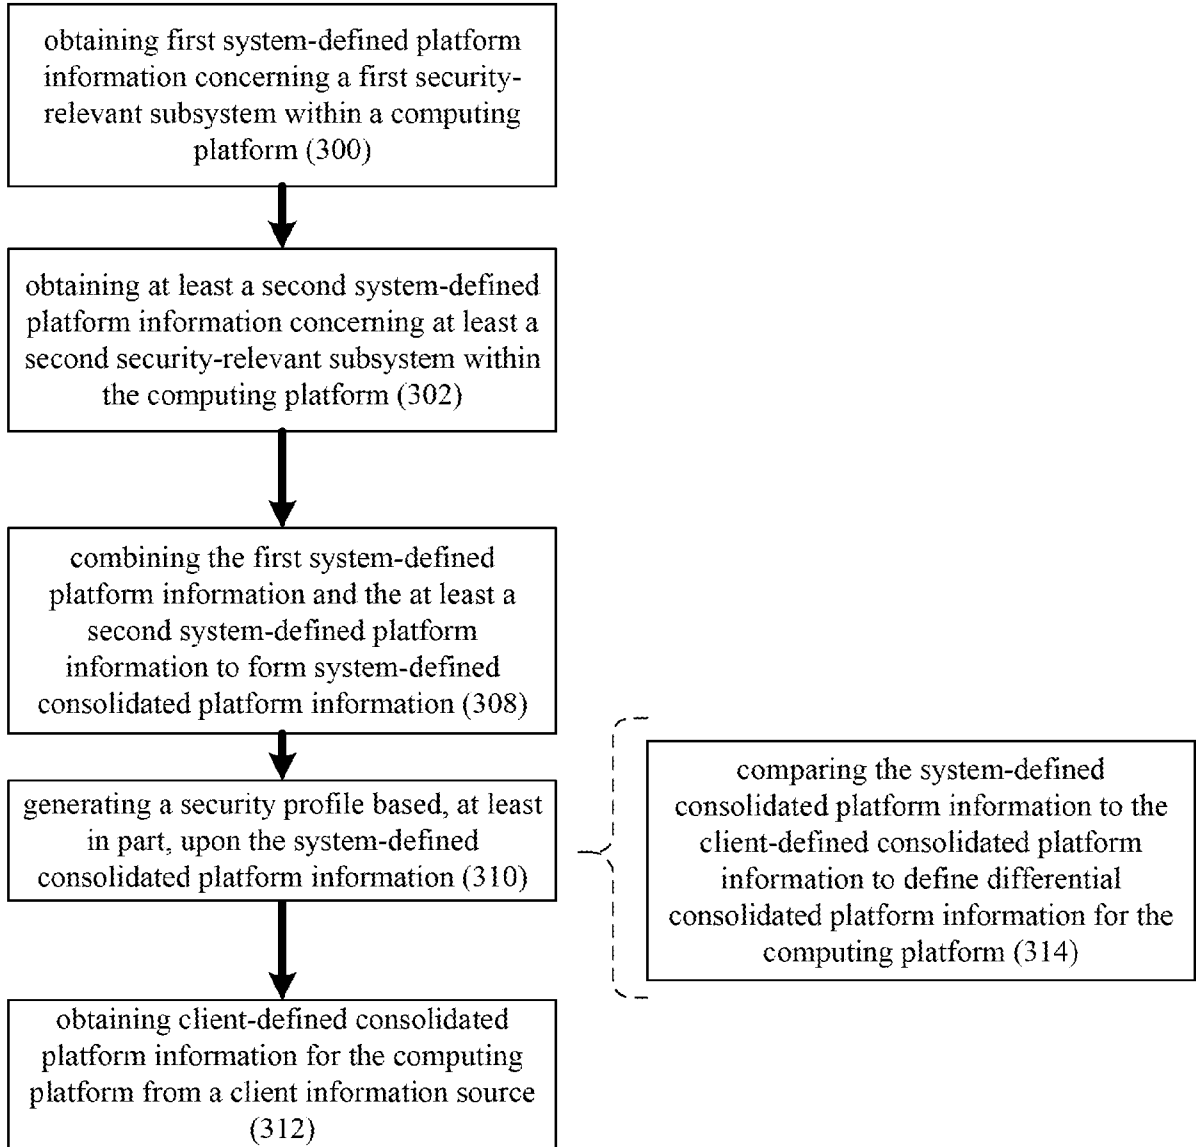


FIG. 4

350

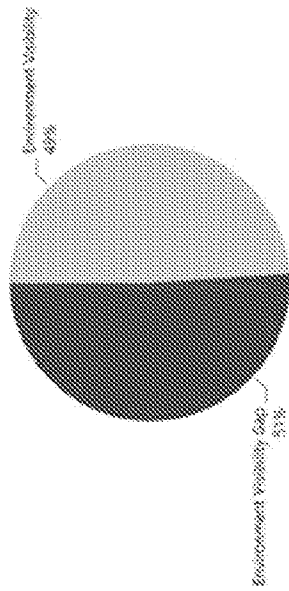
Summary Overview



FIG. 5

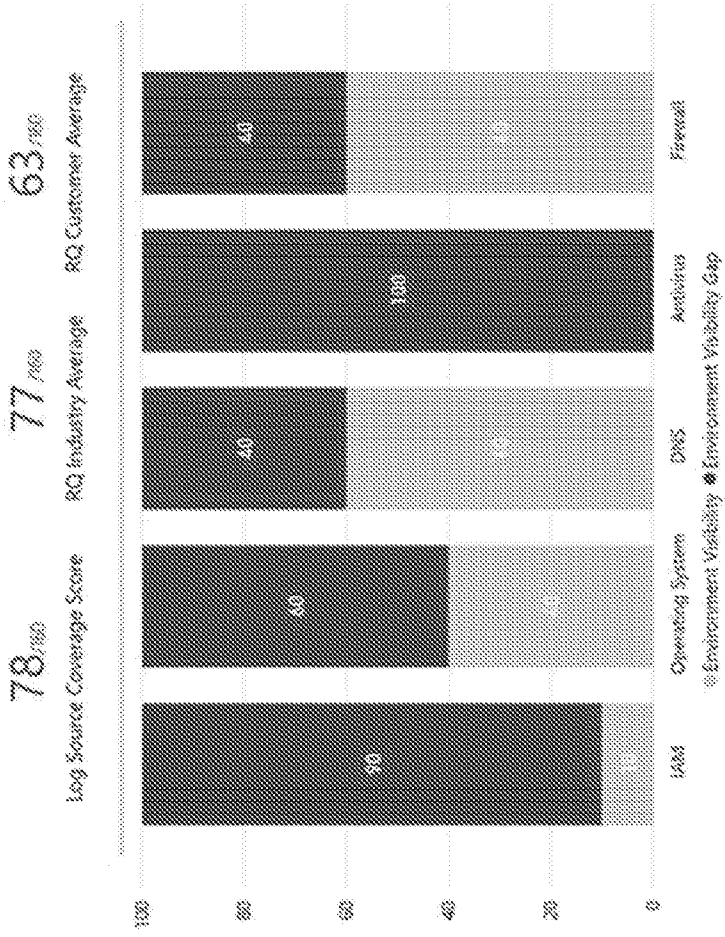
352

Log Source Coverage



354

Priority Function	Visible Log Count	EQ Log Count	(356)	(358)	(360)
1 IAM	1	10			
2 Operating System	4000	10000			
3 DNS	6	10			
4 Antivirus	0	1			
5 Firewall	90	150			



1/7/2019

Reliaquest Model Index

FIG. 6

10

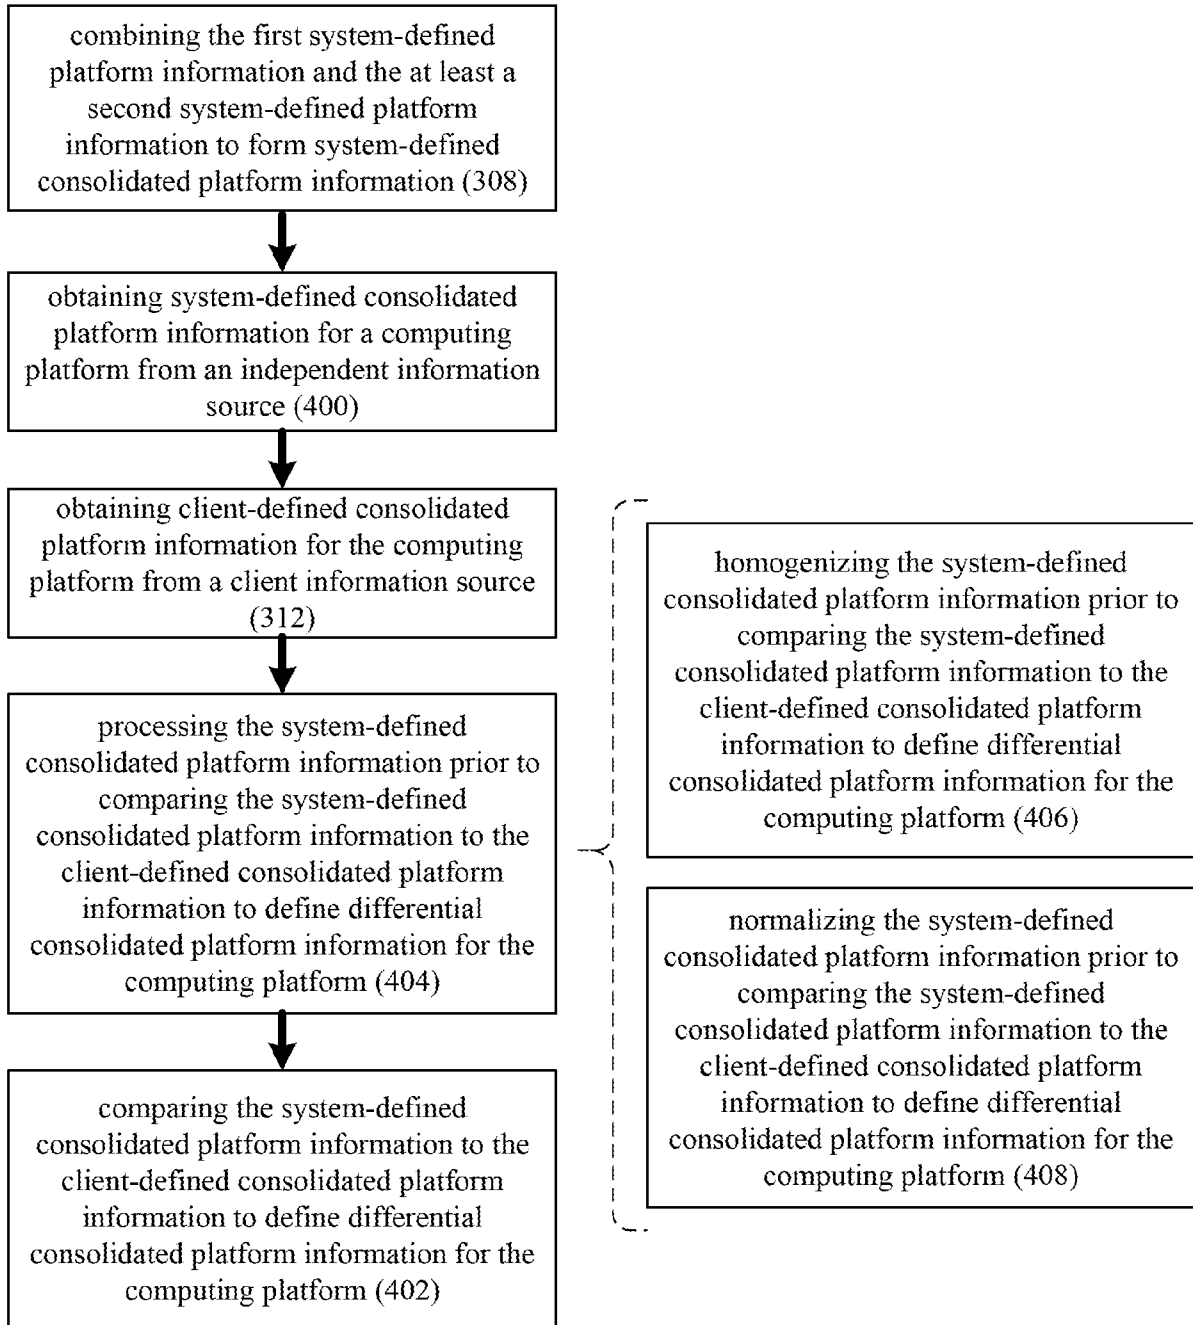


FIG. 7

10

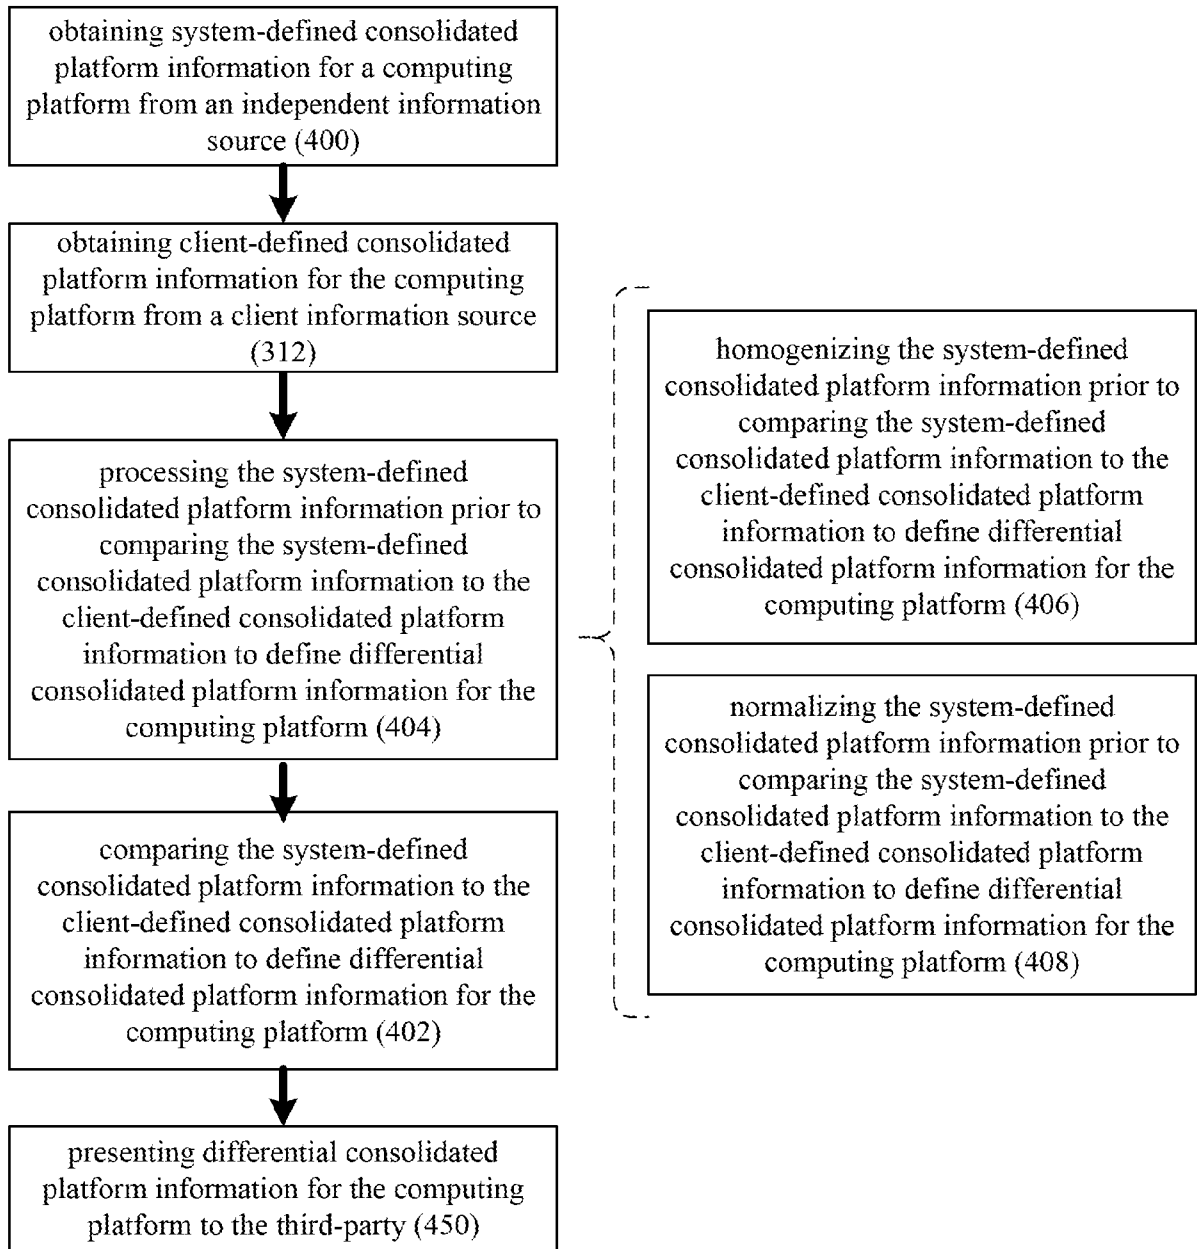


FIG. 8

10

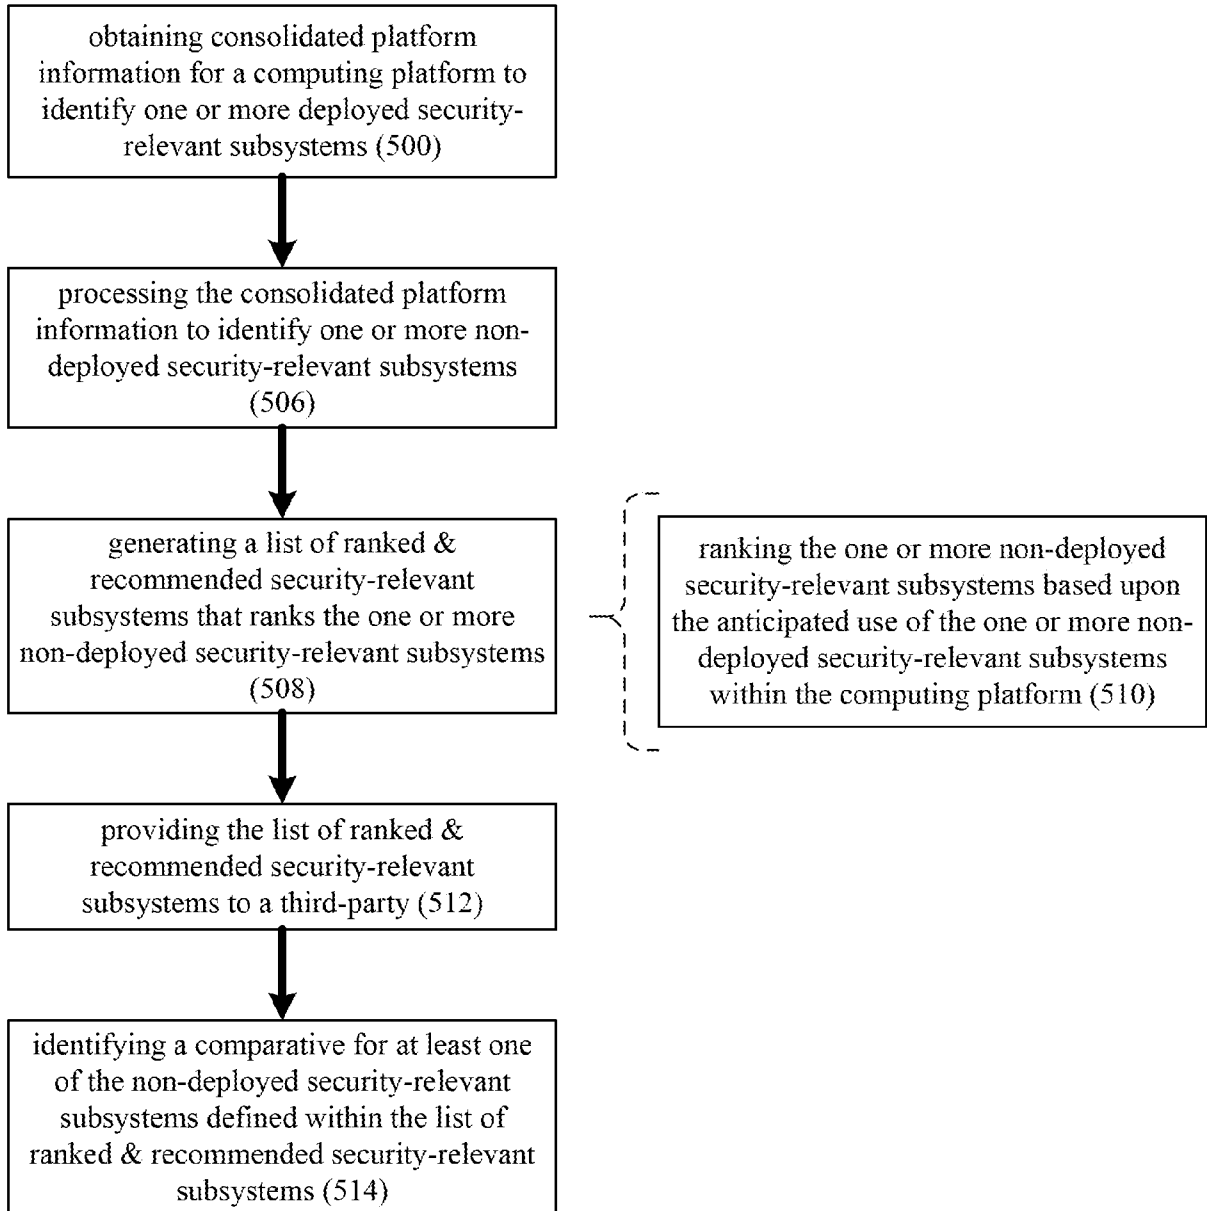
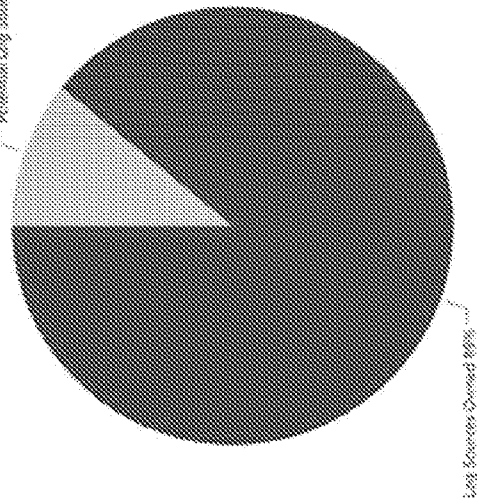


FIG. 9

Log Source Diversity

Potential Log Sources: 12%
 Log Source Diversity Score: **106**
 RQ Industry Average: **47**
 RQ Customer Average: **44**



550

Missing Log Sources by Priority

Priority	Function	RQ Industry Owned	RQ Customer Owned
1	CDN	0%	19%
2	WAF	33%	71%
3	DAM	0%	33%
4	USA	0%	26%
5	API Gateway	0%	22%
6	MDM	0%	19%
		(552)	(554)
			(556)



FIG. 10

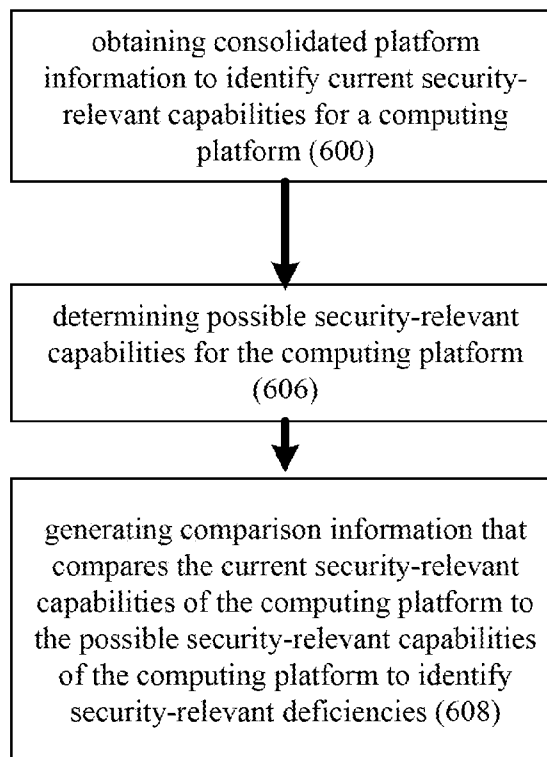
10

FIG. 11




FIG. 12

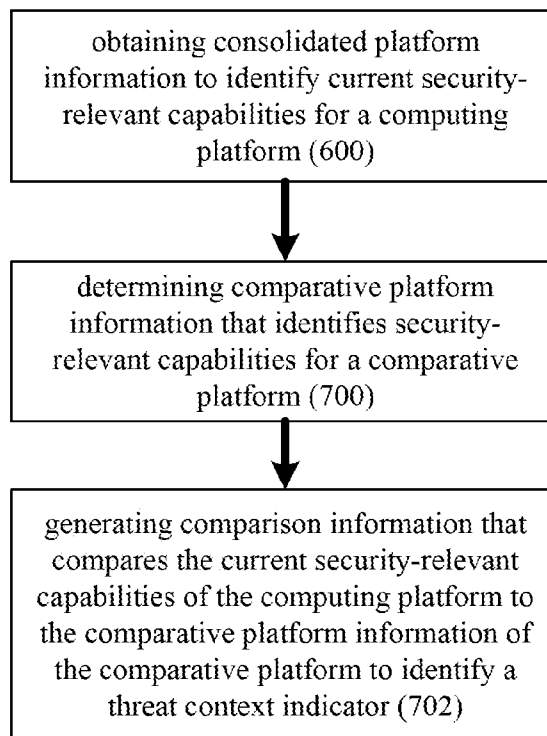
10

FIG. 13

(750)


Threat Context

4	0	3	Automated List Integration SIEM List Update Processing
4	12	3	No Retrospective IOC Hunting
4	3	3	Post Alert Analysis Correlation Sources
4	12	3	Threat Enabled Rules
4	0	3	Open Source Threat Lists

25 /40

Threat Context Score

Trending Threat Context Score



14/30



FIG. 14

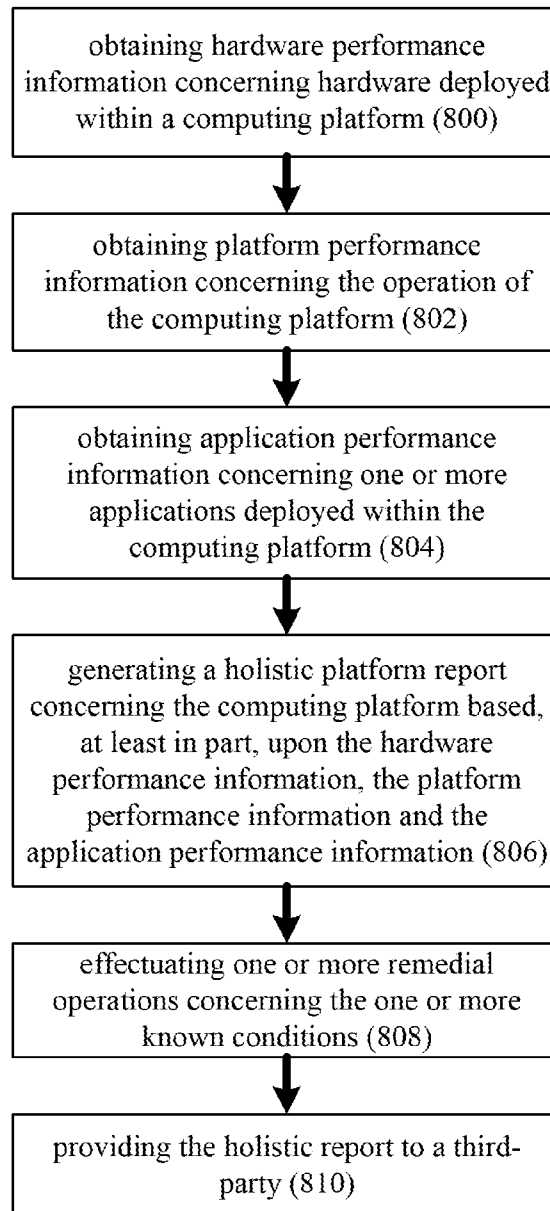
10

FIG. 15

352



352

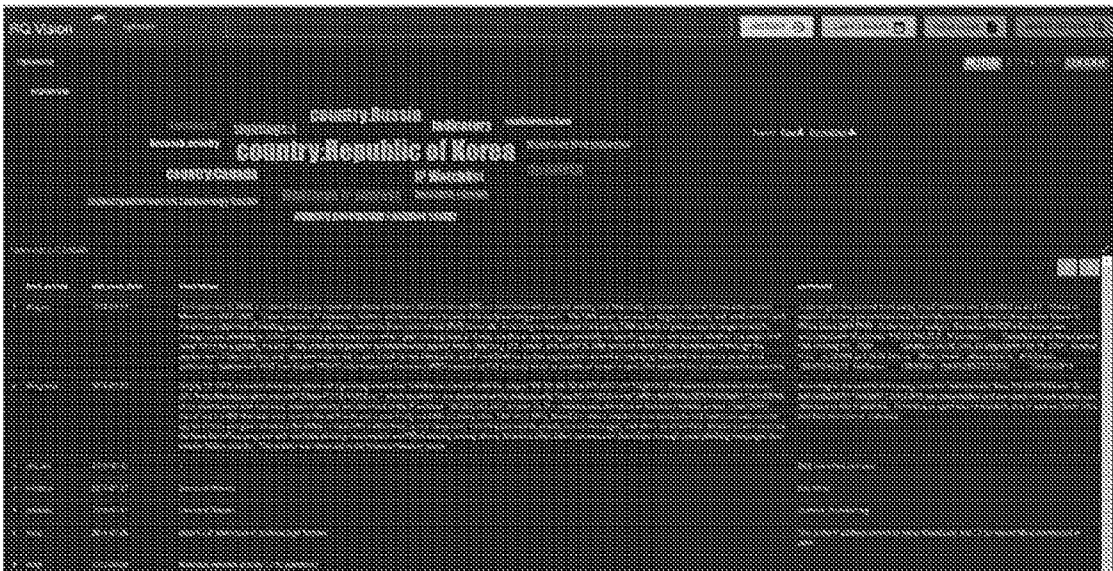


FIG. 16

SUBSTITUTE SHEET (RULE 26)

10

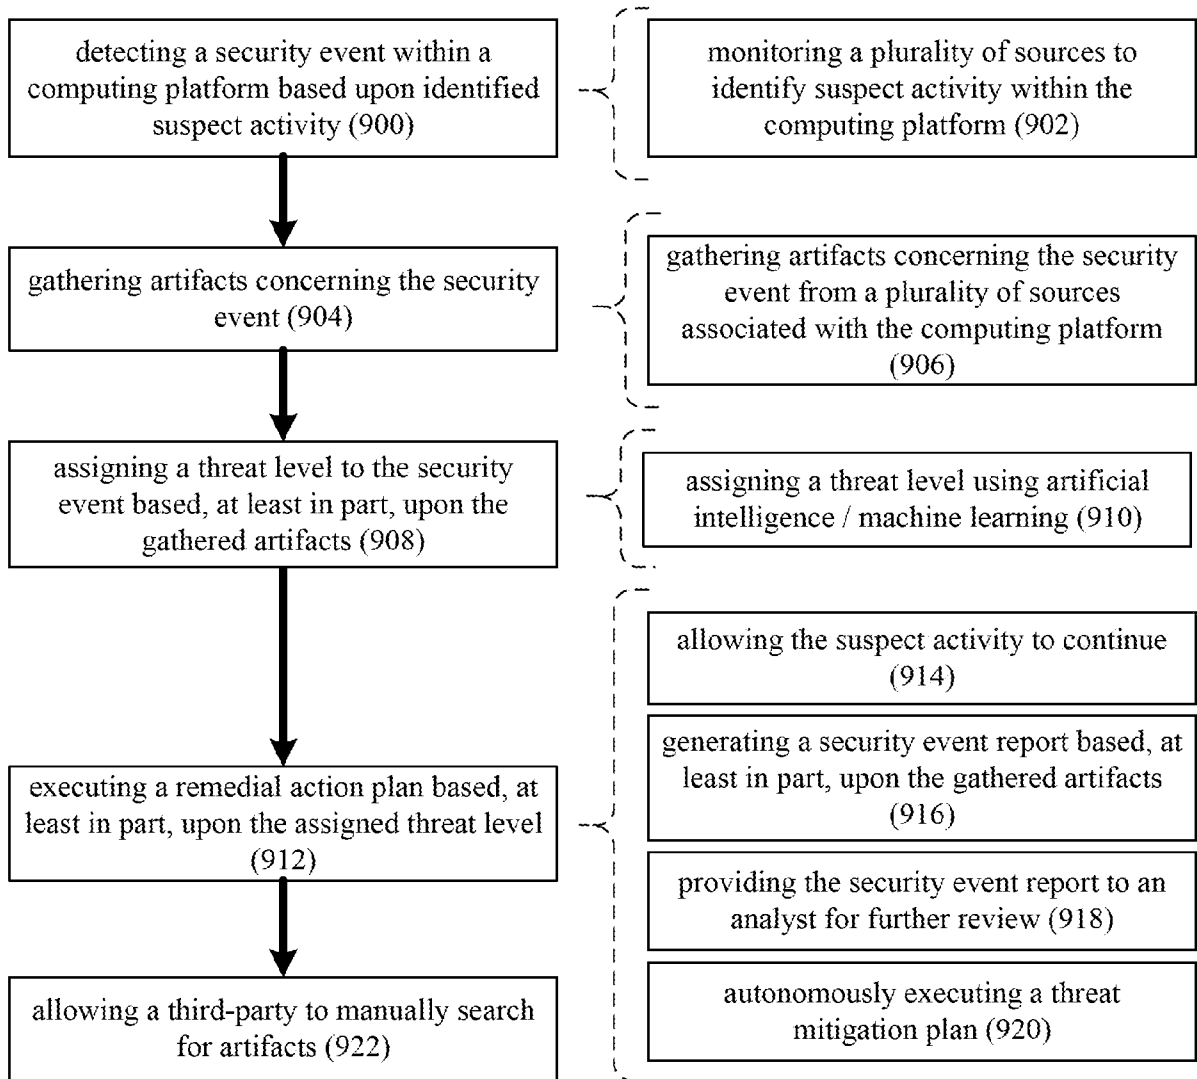


FIG. 17

10

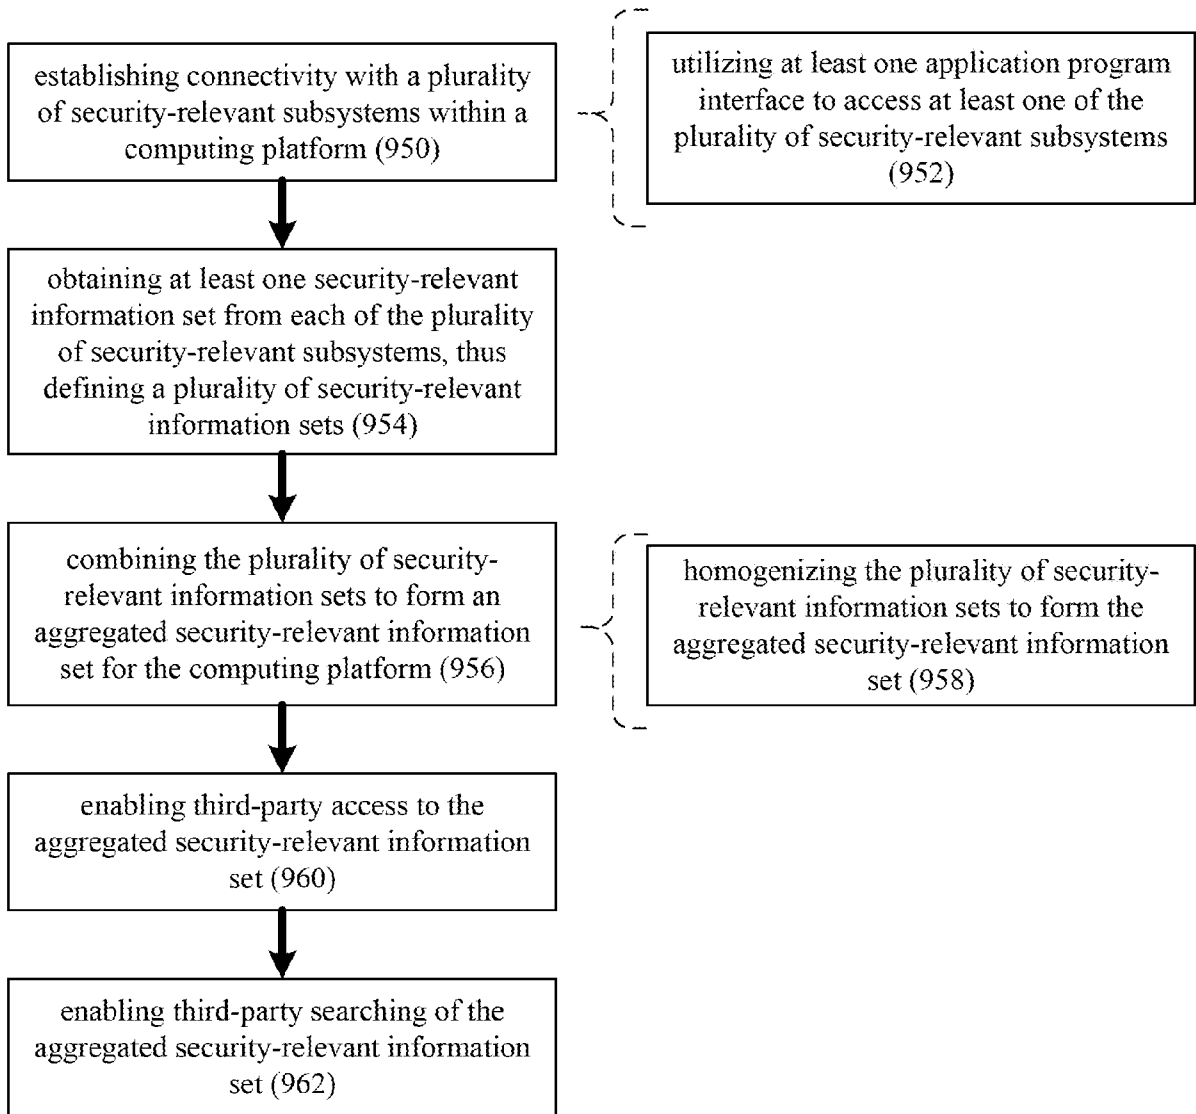


FIG. 18

10

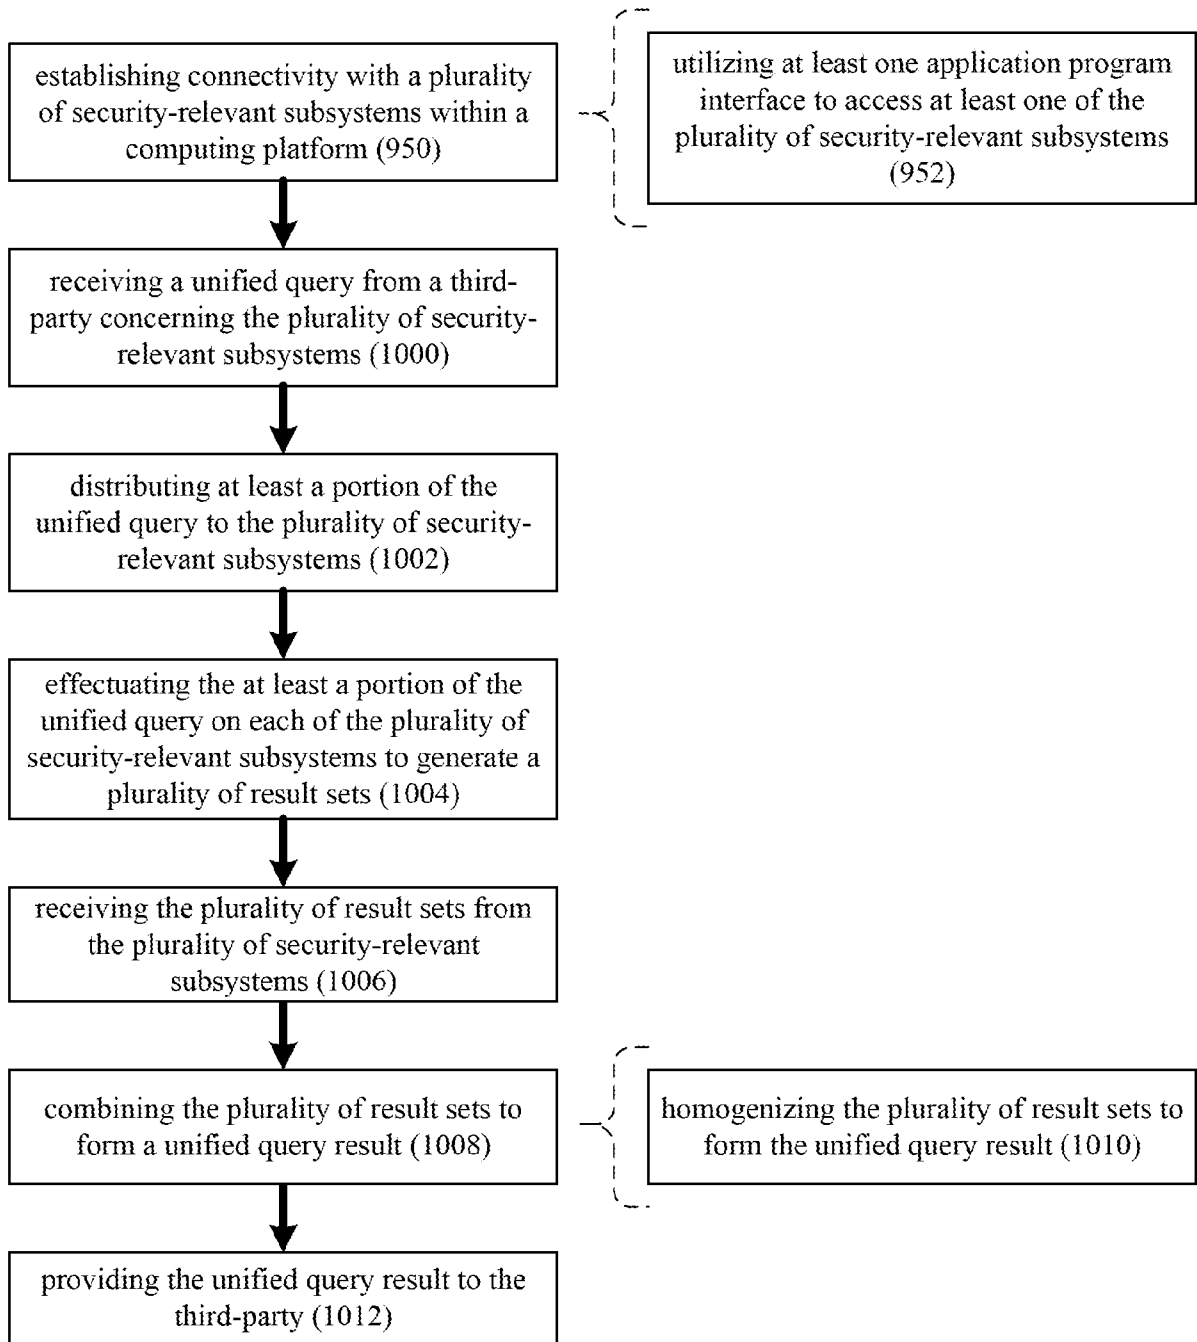


FIG. 19

10

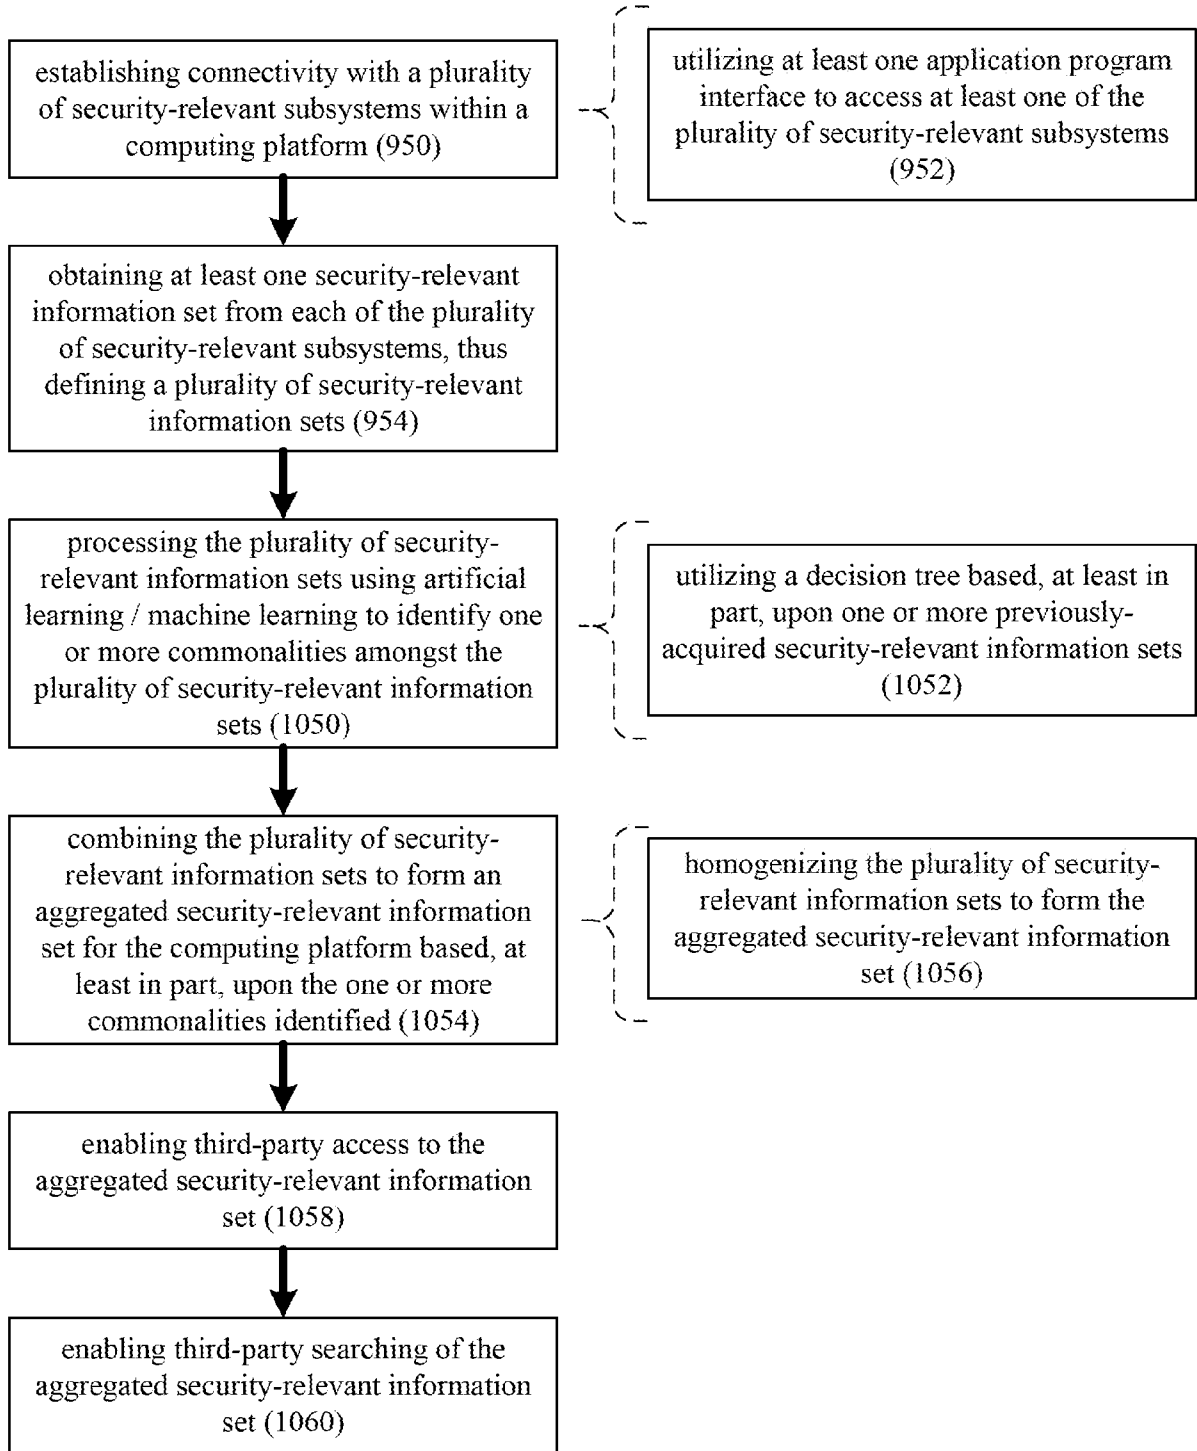


FIG. 20

10

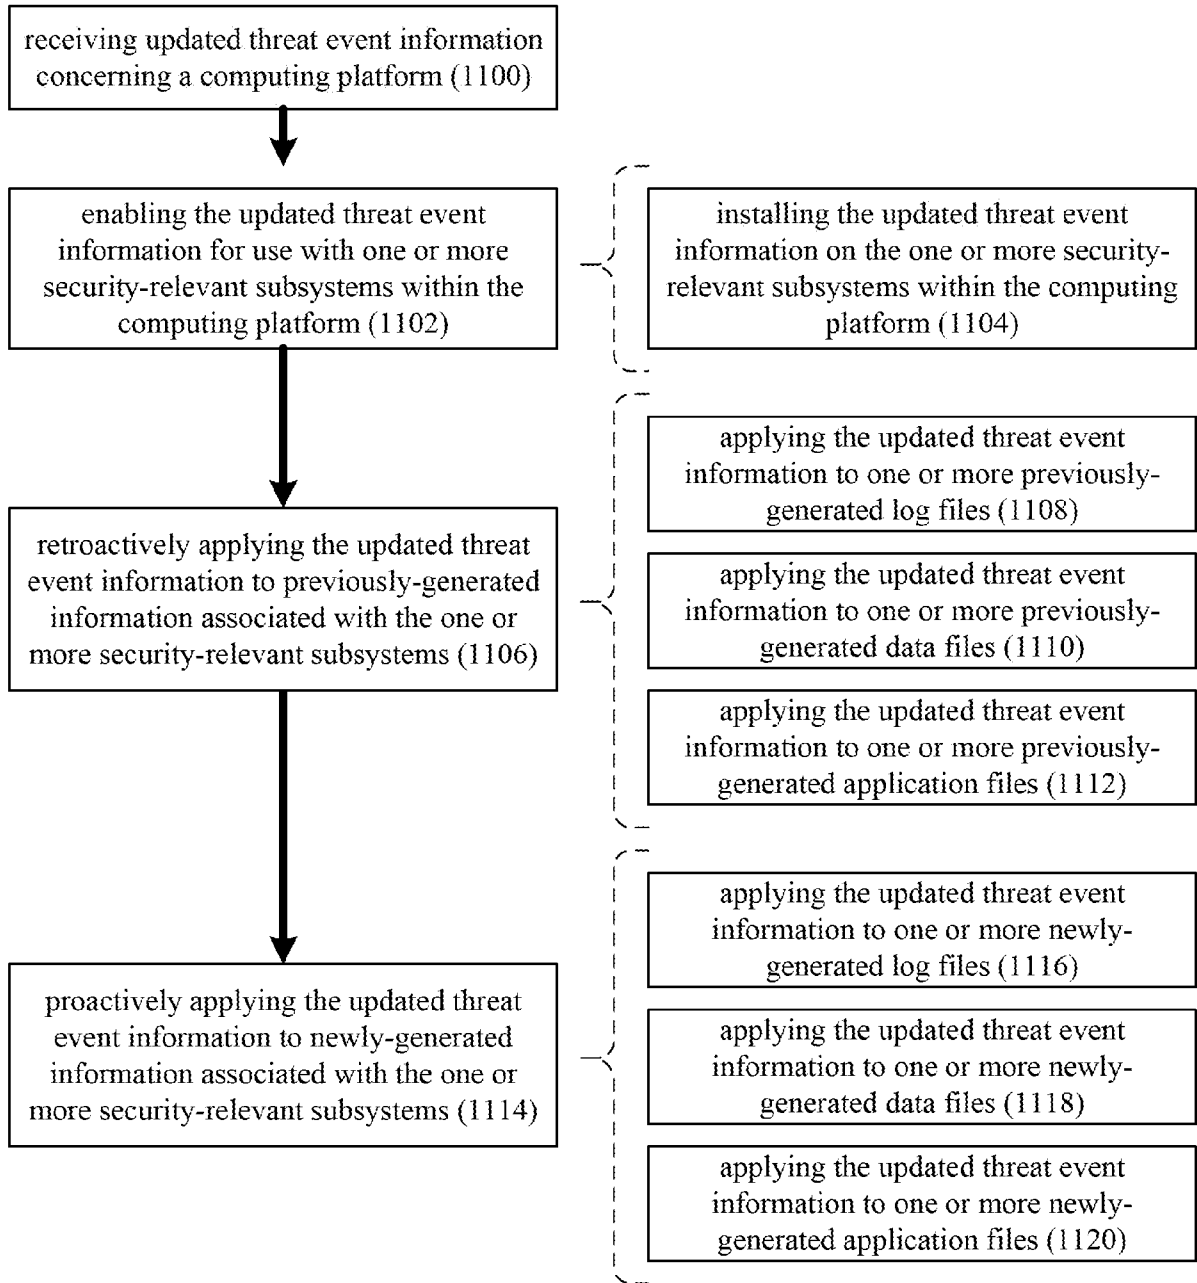


FIG. 21

10

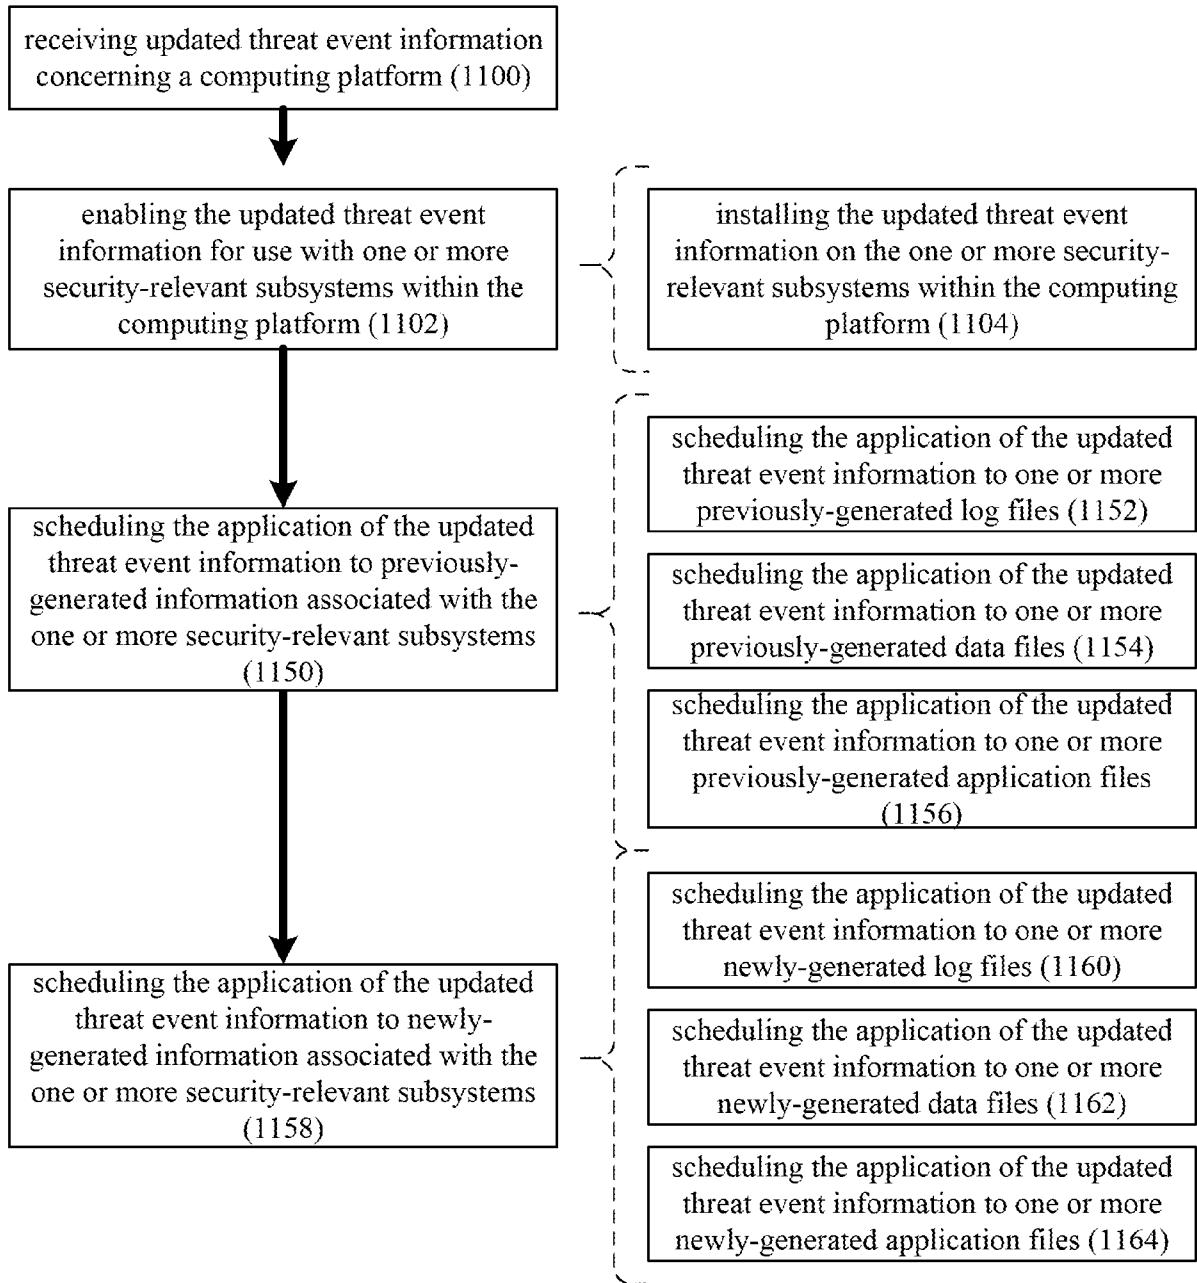


FIG. 22

10

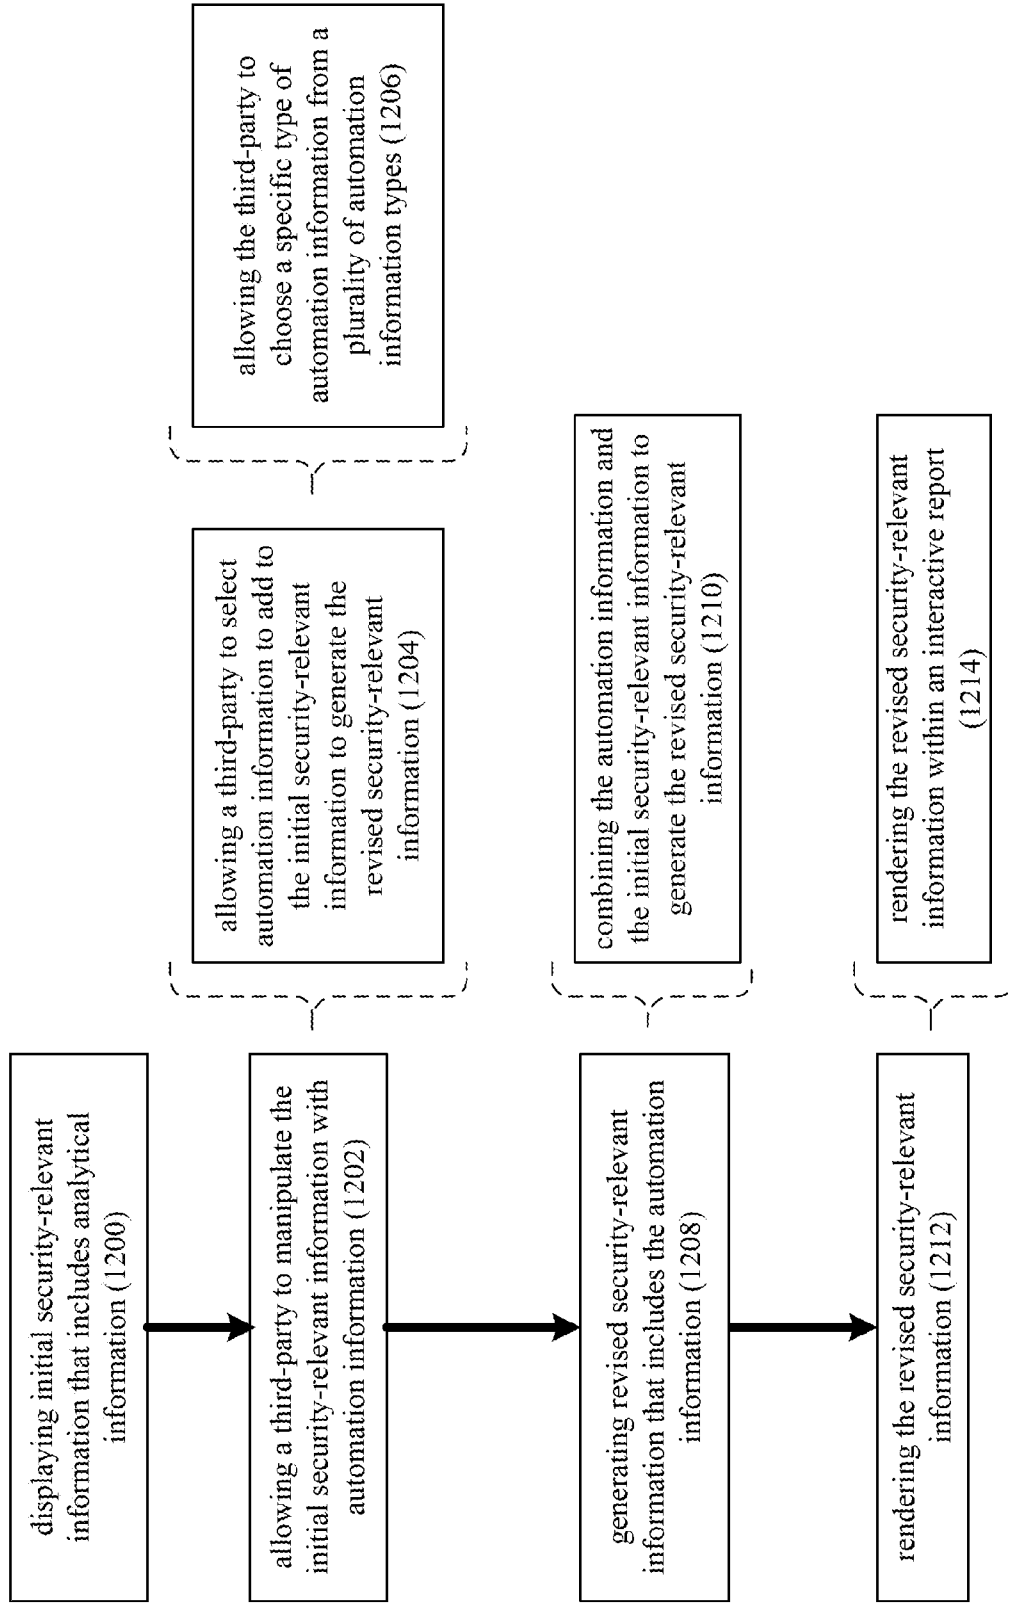


FIG. 23

1258

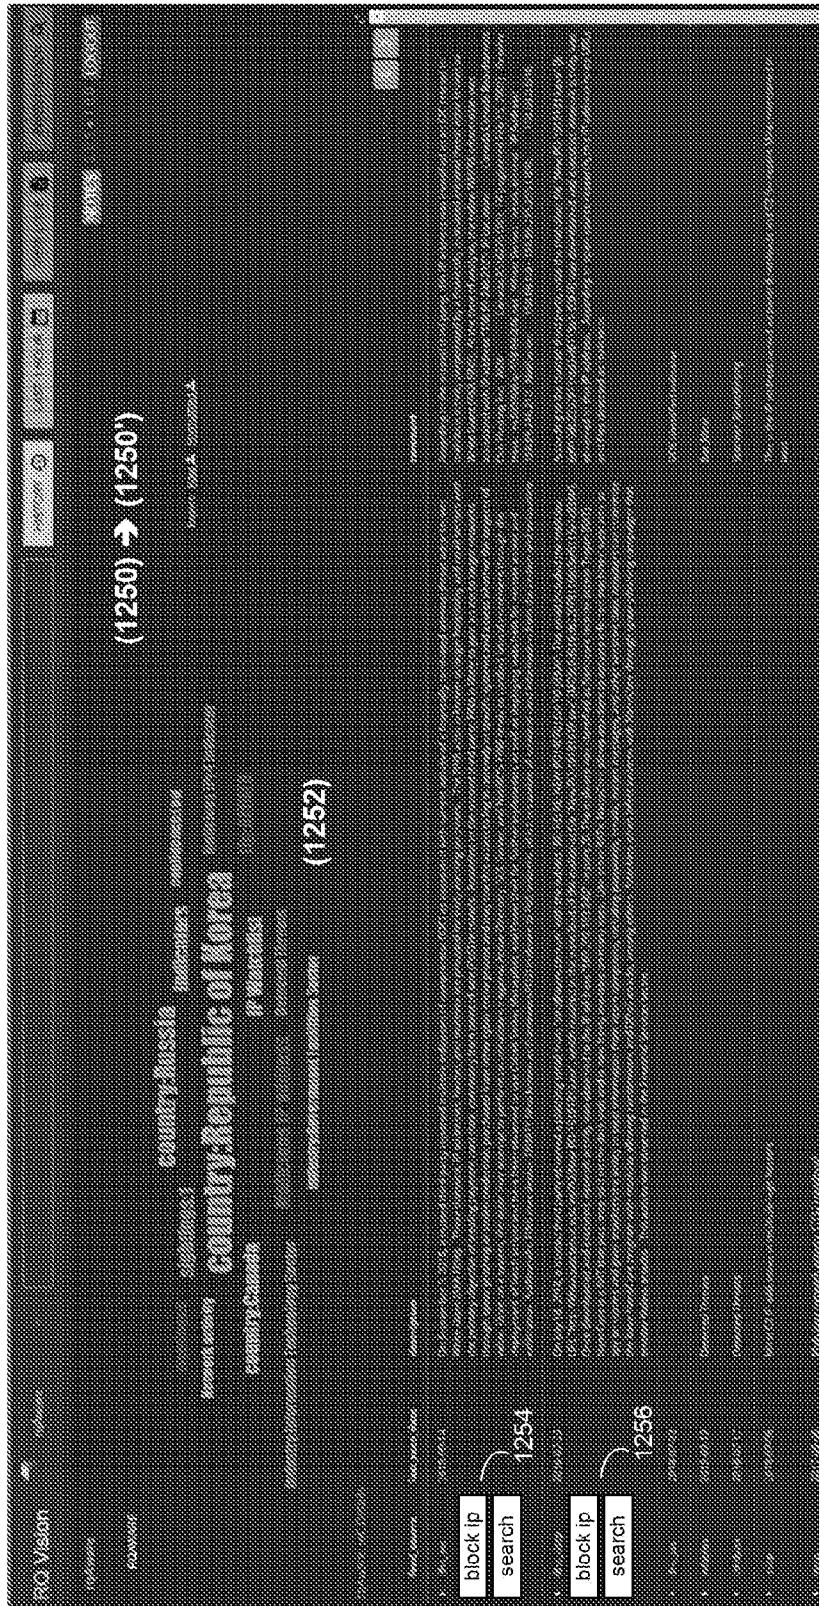


FIG. 24

10

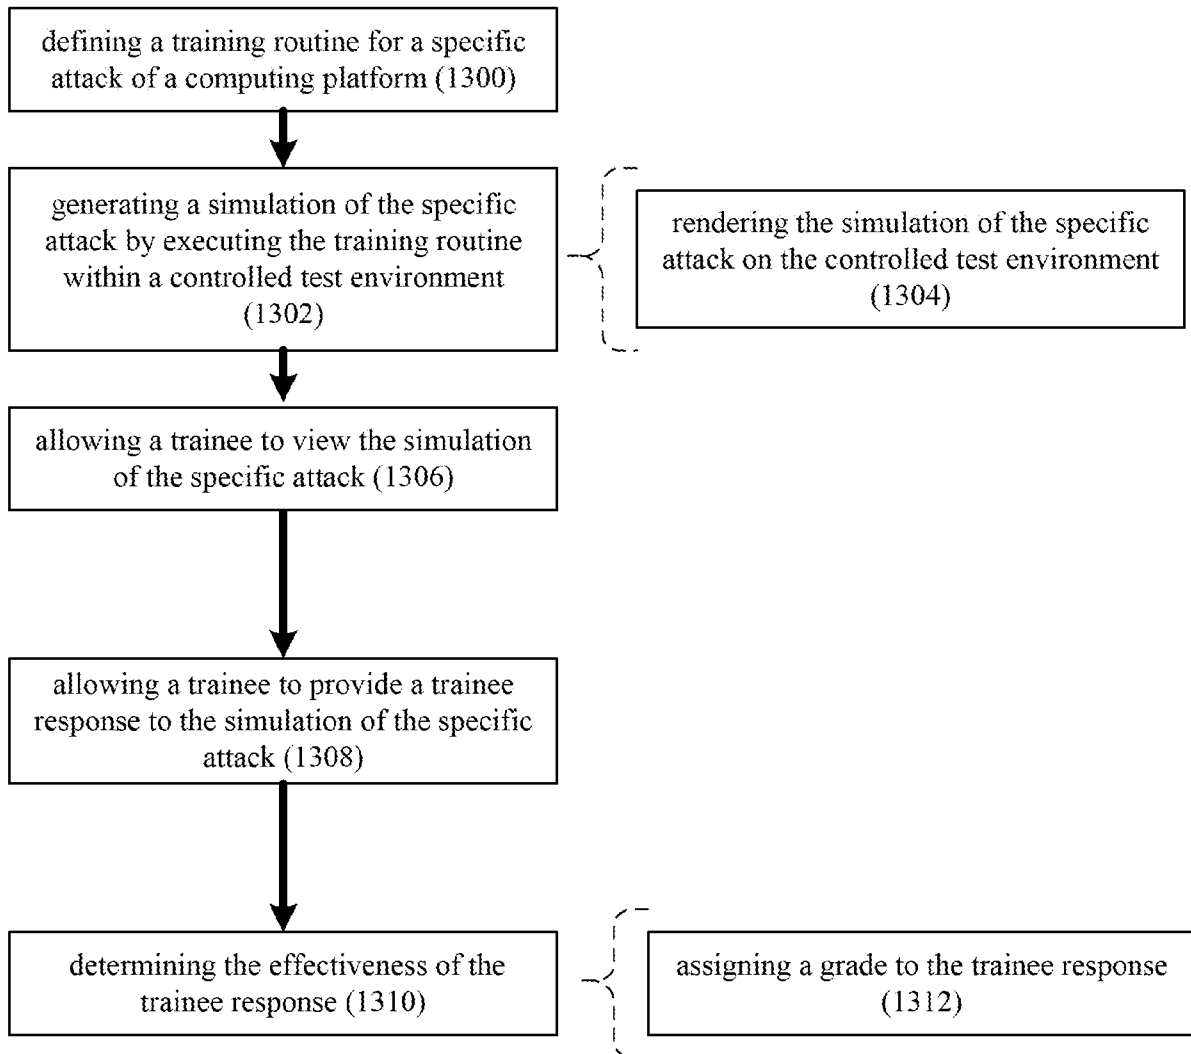


FIG. 25

10

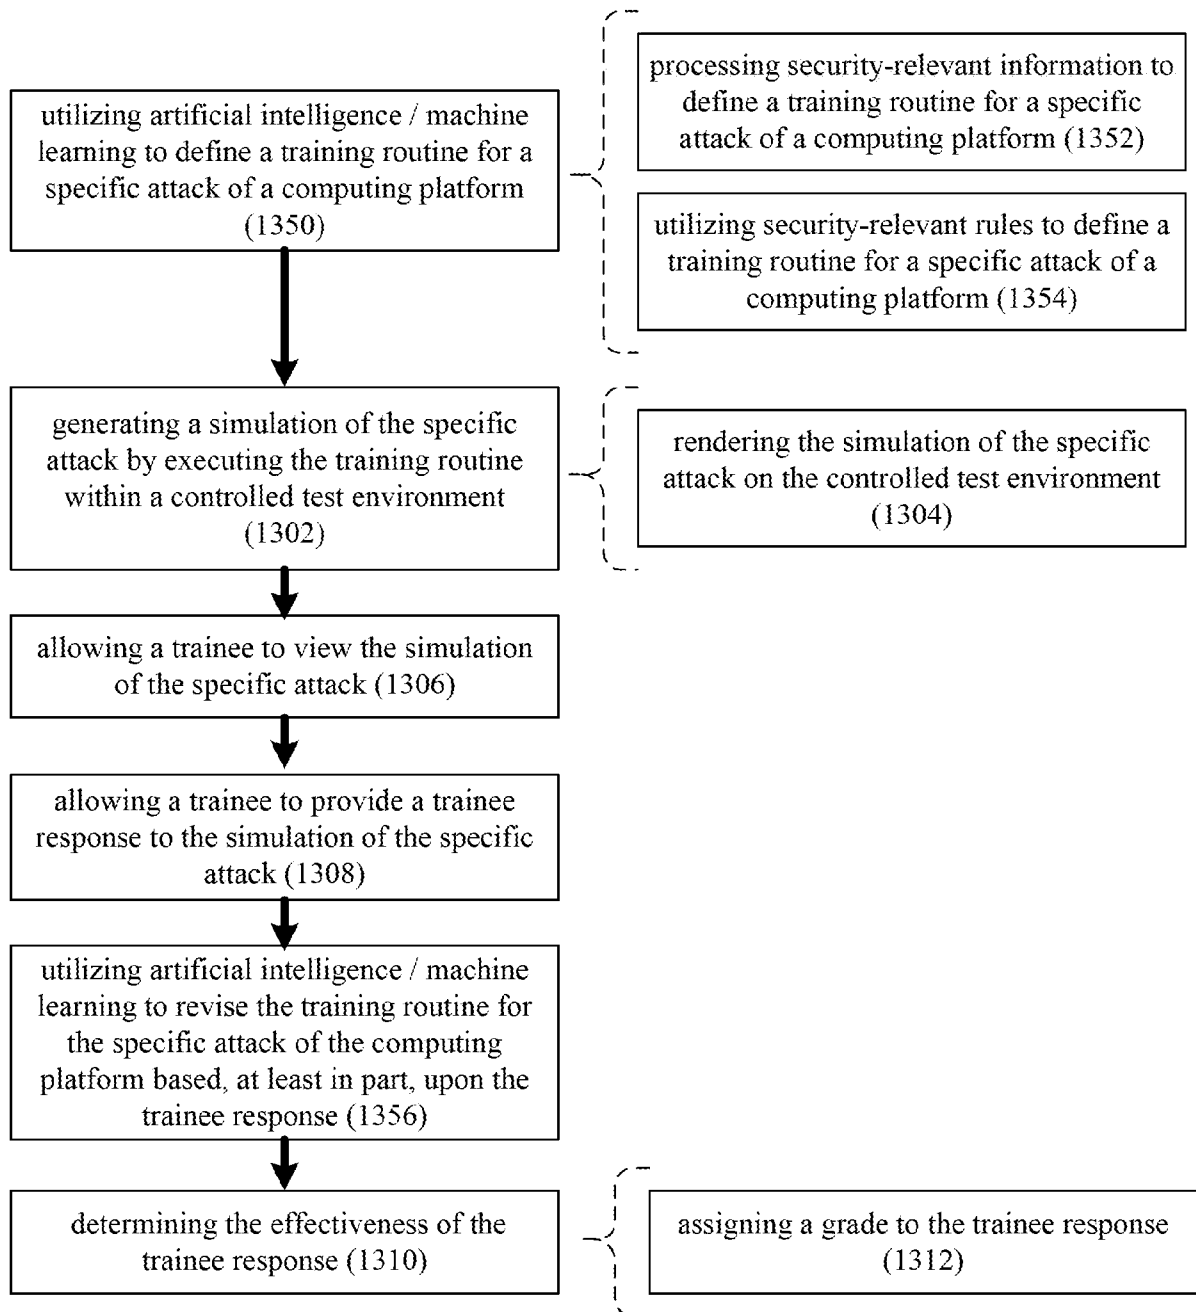


FIG. 26

10

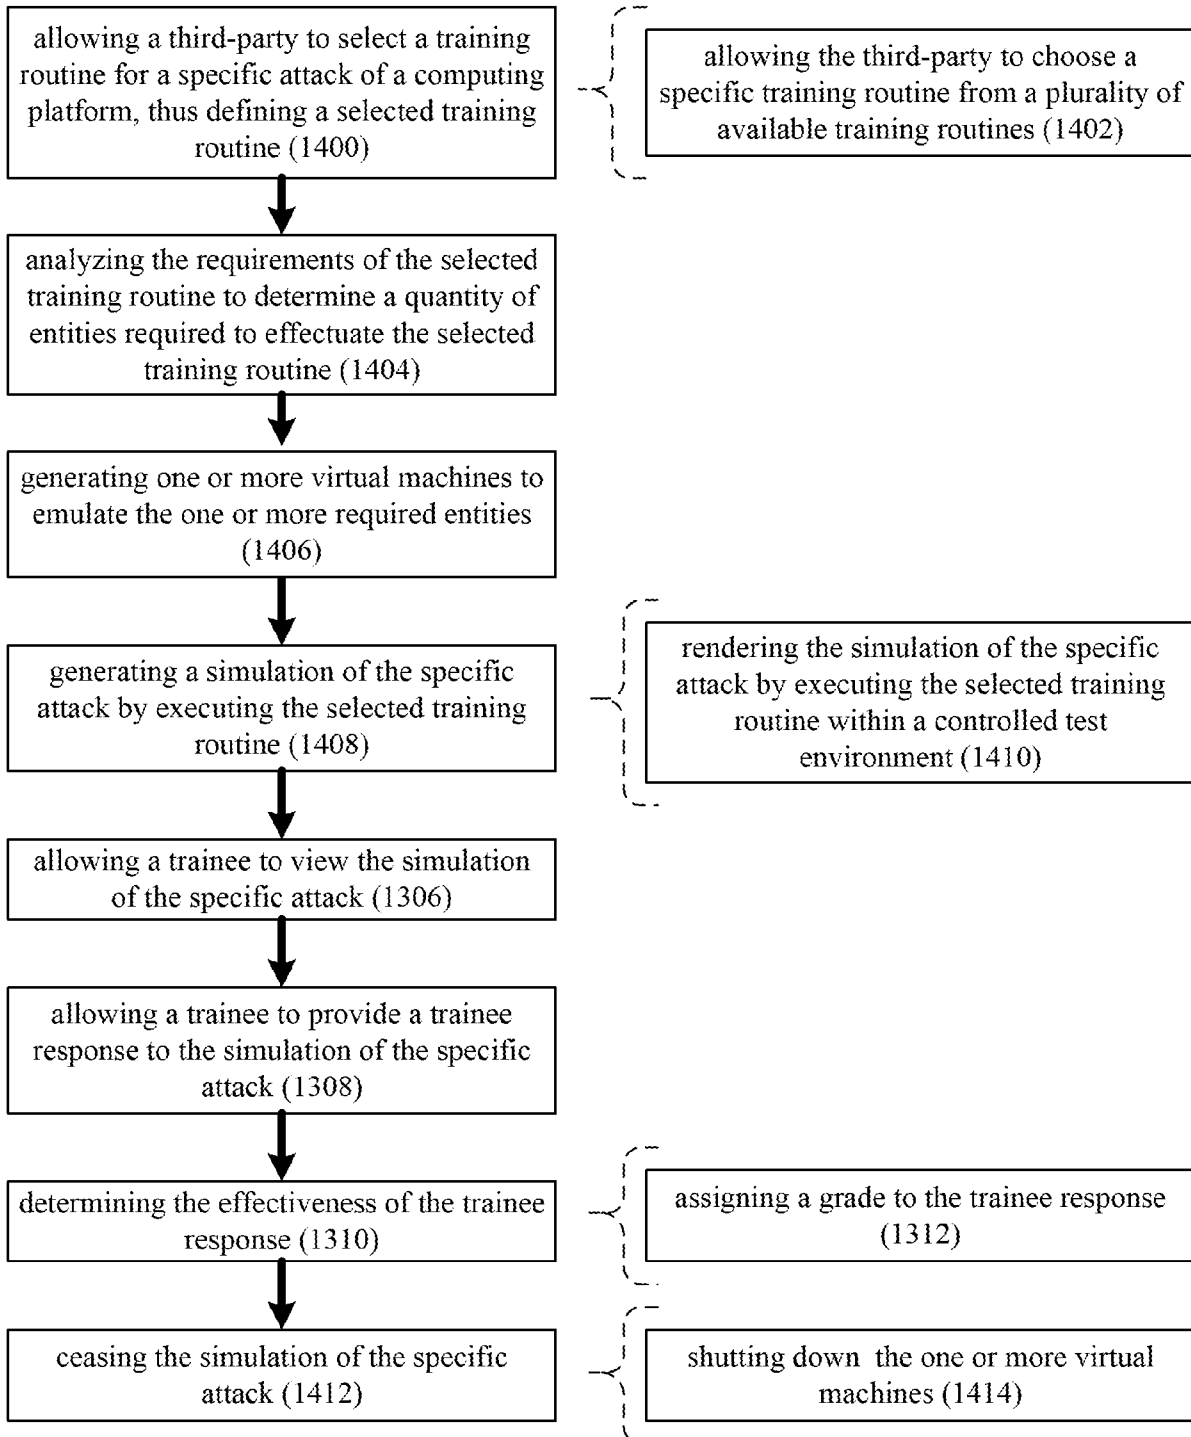


FIG. 27

10

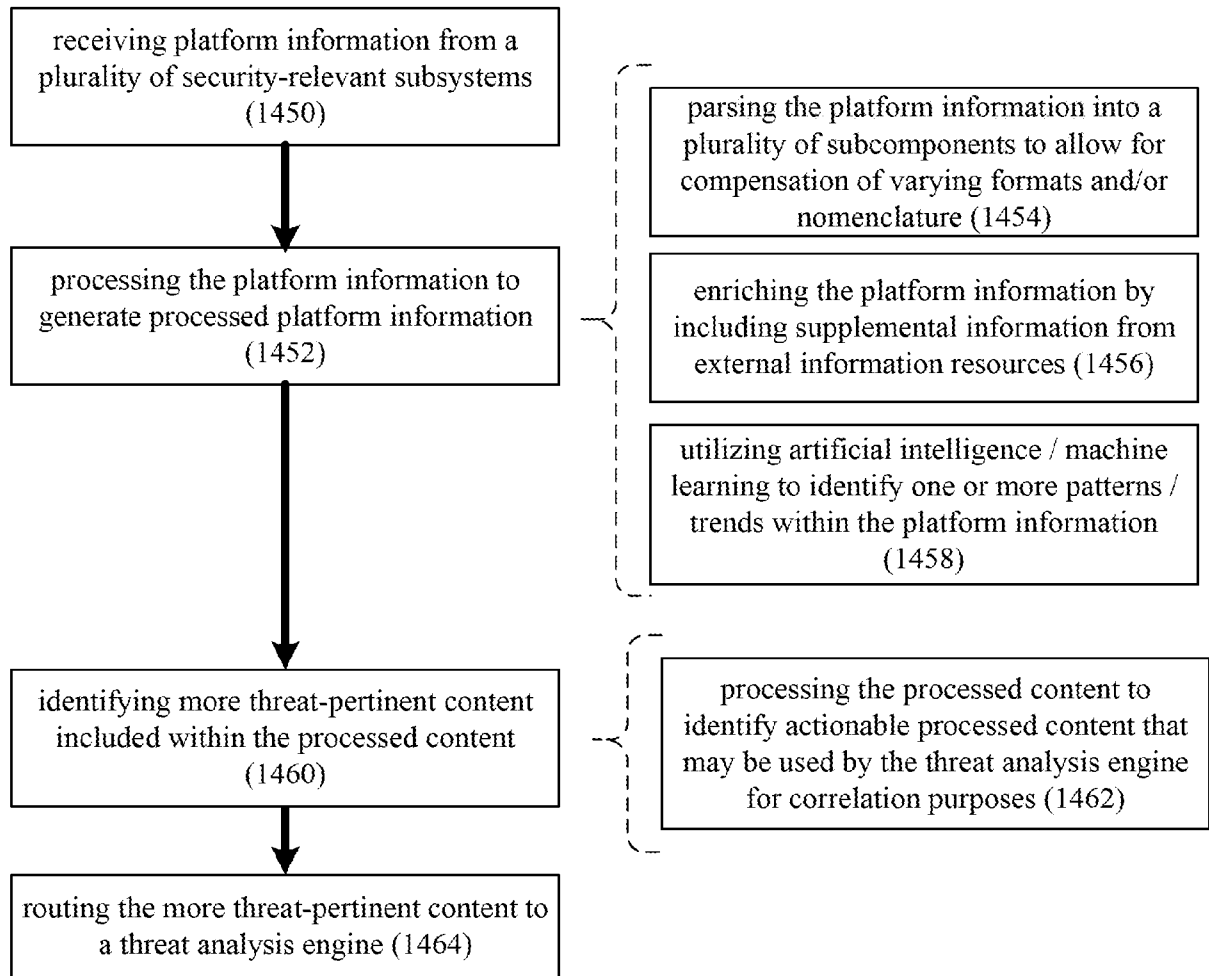


FIG. 28

10

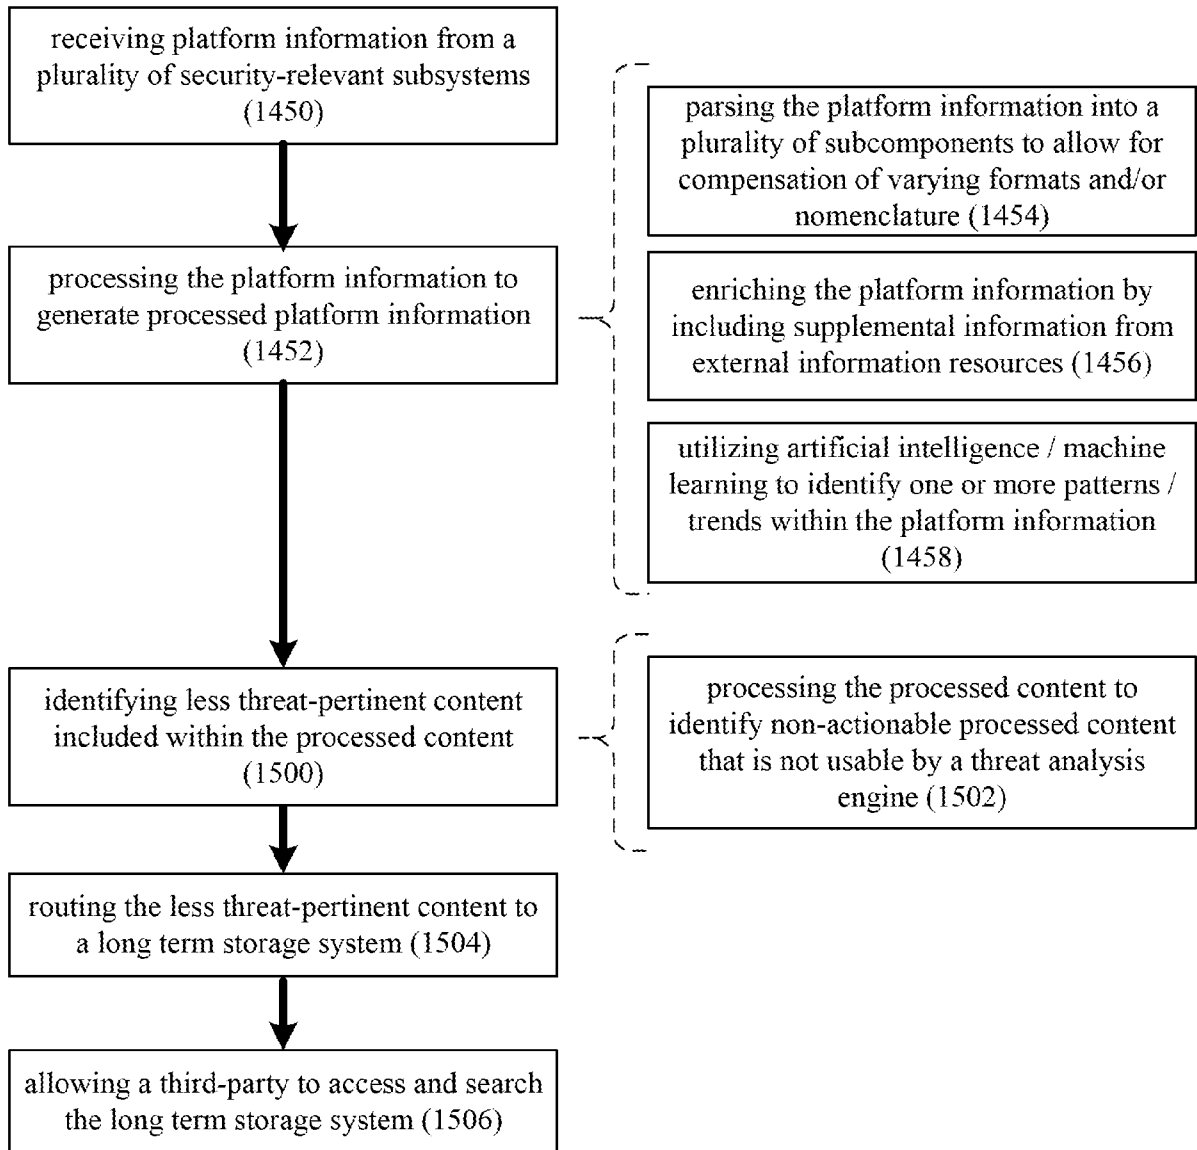


FIG. 29

10

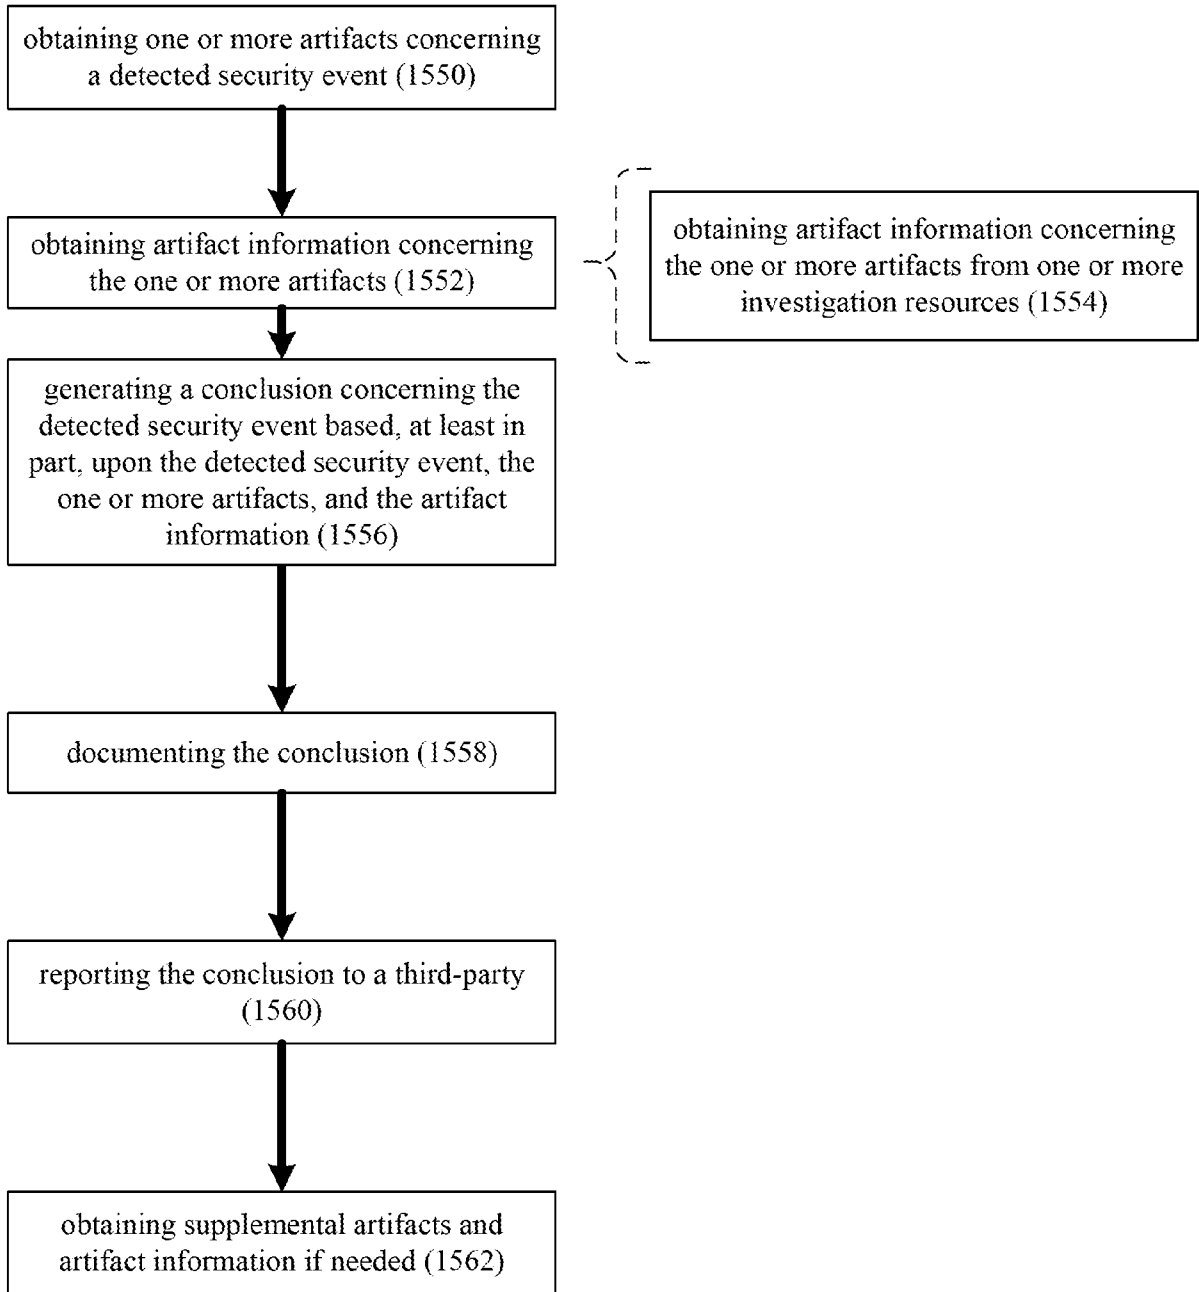


FIG. 30

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 19/35704

A. CLASSIFICATION OF SUBJECT MATTER IPC(8) - H04L 29/00 (2019.01) CPC - H04L 41/0213, H04L 41/0631, H04L 63/1408, H04L 63/1416, H04L 63/20		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) See Search History Document		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched See Search History Document		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) See Search History Document		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2017/0214701 A1 (HASAN), 27 July 2017 (27.07.2017), entire document, especially Abstract; para [0002], [0016], [0049]-[0050], [0073], [0078], [0087], [0094]	1-21
A	US 9,398,035 B2 (Cisco Technology, Inc.) 19 July 2016 (19.07.2016) entire document	1-21
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family		
Date of the actual completion of the international search 02 August 2019 (02.08.2019)		Date of mailing of the international search report 08 OCT 2019
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-8300		Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774