



(51) International Patent Classification:
H04L 29/06 (2006.01)

(21) International Application Number:
PCT/US2017/021961

(22) International Filing Date:
10 March 2017 (10.03.2017)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: VISA INTERNATIONAL SERVICE ASSOCIATION [US/US]; P.O. BOX 8999, SAN FRANCISCO, California 94128-8999 (US).

(72) Inventors: ZAWOAD, D.M. Shams; 748 Marlin Ave., #4, Foster City, California 94404 (US). DESCH, Philip; 900 Metro Center Boulevard, Foster City, California 94404 (US).

(74) Agent: RICKETT, Cynthia H. et al.; Kilpatrick, Townsend & Stockton LLP, Mailstop: IP Docketing - 22, 1100 Peachtree Street, Suite 2800, Atlanta, Georgia 30309 (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, 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: IDENTIFYING MALICIOUS NETWORK DEVICES

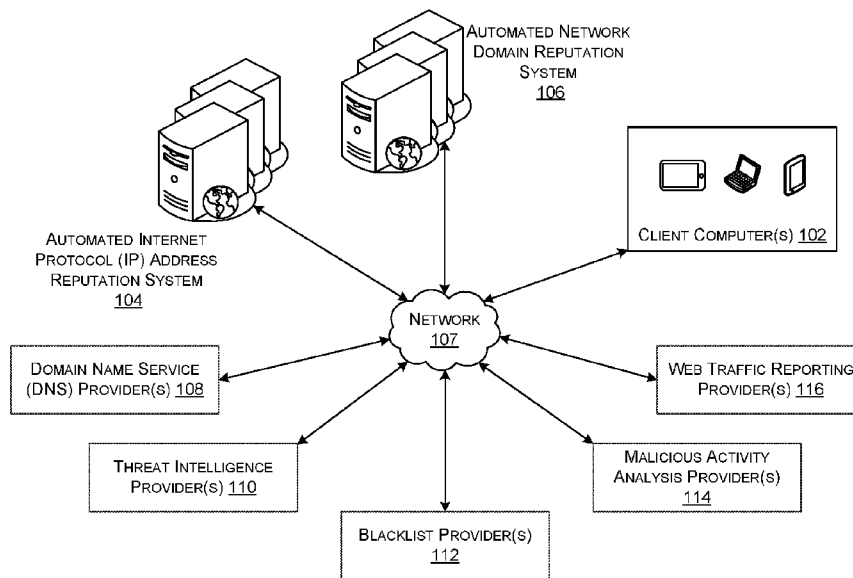


FIG. 1

100

(57) Abstract: Embodiments provide for maliciousness scores to be determined for IP addresses and/or network domains. For example, a request to evaluate malicious activity with respect to an IP address/network domain may be received. Multiple, and in some cases disparate, third-party systems may provide malicious activity information associated with the IP address and/or network domain. A feature set may be extracted from the malicious activity information and statistical values may be calculated from the extracted data and added to the feature set. The features set may be provided to a machine learning model as input and a maliciousness score/classification may be returned. A remedial action may be performed in accordance with the output of the machine learning model.



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))

IDENTIFYING MALICIOUS NETWORK DEVICES

BACKGROUND

[0001] A crucial task in cyber defense is to assess the threats posed by a particular network device (e.g., a website), such as one identified by an Internet protocol (IP) address and/or a network domain name. Communication involving a malicious website can compromise a network. However, automatically assessing the maliciousness of a website, or the severity level of potential maliciousness, is challenging due to the rapidly changing behavior of IP addresses and/or network domain names. Conventional cybersecurity systems are either too subjective or inaccurate in detecting the maliciousness of a website with a high degree of fidelity. Therefore, it is currently difficult to detect malicious and suspicious activity, or a potential of malicious or suspicious activity, of IP addresses and/or network domains involved in the communication. Accordingly, conventional cybersecurity systems present significant drawbacks with respect to detection and prevention capabilities and overall security posture.

BRIEF SUMMARY

[0002] Embodiments described here address the problem described above and other problems individually and collectively. A cybersecurity system can calculate maliciousness scores for network identifiers (e.g., an IP address, a network domain, etc.) for which a malicious activity evaluation has been requested. To calculate such scores, the cybersecurity system may request malicious activity information from a variety of third-party servers such as threat intelligence providers, DNS providers, blacklist providers, network traffic report providers, and/or one or more antivirus report providers. The cybersecurity system may utilize the received malicious activity information to extract a first set of features and calculate a second set of features. These features can be input into a machine-learning model that has been previously trained using similar features of historical malicious activity information and known security classifications. Utilizing the machine-learning model, a maliciousness score may be calculated for the network identifiers. The cybersecurity system, or another system, can perform various remedial actions (e.g., influence network traffic) in accordance with the maliciousness score.

[0003] Other embodiments are directed to systems and non-transitory computer readable media associated with methods described herein.

[0004] A better understanding of the nature and advantages of embodiments of the present invention may be gained with reference to the following detailed description and the
5 accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 is a block diagram of a system for determining a maliciousness score for a network identifier, according to some embodiments.

[0006] FIG. 2 is a block diagram of an exemplary cybersecurity system for determining
10 maliciousness scores for network identifiers, according to some embodiments.

[0007] FIG. 3 is an example computer architecture capable of implementing at least some embodiments of the cybersecurity system.

[0008] FIGS. 4A and 4B are flow diagrams illustrating an example processes for
15 determining related network identifiers from an initial set of network identifiers, in accordance with some embodiments.

[0009] FIG. 5 is an example computer architecture of the parameter determination engine of FIG. 3 that is capable of implementing at least some aspects of the cybersecurity system, according to some embodiments.

[0010] FIG. 6 is an example computer architecture of the decision engine of FIG. 3, that is
20 capable of implementing at least some aspects of the cybersecurity system, according to some embodiments.

[0011] FIG. 7 is a flow diagram illustrating an example process for training and updating a machine-learning model, according to some embodiments.

[0012] FIG. 8 shows a flow chart of an exemplary method for determining a maliciousness
25 score for one or more network identifiers, in accordance with some embodiments.

TERMS

[0013] Prior to discussing embodiments of the invention, description of some terms may be helpful in understanding embodiments of the invention.

[0014] The term “client computer” generally refers to a computer that requests information or a service. A client computer may comprise a computer (e.g., desktop computer), a mobile device (e.g., a smart phone, laptop, or tablet computer), or a wearable device (e.g., a smart watch or activity tracker). The client computer may include wireless communication capabilities (e.g., Wi-Fi, Bluetooth, or near-field communications). In some embodiments, a client computer may communicate with a server computer. In some embodiments, a first client computer may not be capable of communicating with a server computer unless a second client computer acts as a proxy, sending and receiving messages for the first client computer.

10 [0015] The term “server computer” may include a powerful computer or cluster of computers. For example, the server computer can be a large mainframe, a minicomputer cluster, or a group of servers functioning as a unit. The server computer may be coupled to one or more databases and may include any hardware, software, other logic, or combination of the preceding for servicing requests from one or more client computers. The server
15 computer may be a service provider computer that provides services (e.g., a reputation determination service) to one or more client computers. The server computer may comprise one or more computational apparatuses and may use any of a variety of computing structures, arrangements, and compilations for servicing the requests from one or more client computers.

[0016] The term “network identifier” is intended to refer to an internet protocol (IP) address or a network domain name. “Network identifiers” may refer to two or more IP
20 addresses, two or more network domain names, or two or more identifiers that include at least one IP address and at least one network domain name. An IP address may be a unique string of numbers separated by periods that identifies a computer using the Internet Protocol to communicate over a network. A network domain name may be an identification string that
25 defines a realm of administrative autonomy, authority or control within the Internet. Network domain names may be formed by the rules and procedures of a Domain Name System (DNS) and registered in the DNS. Network domain names may be used in various networking contexts and application-specific naming and addressing purposes. In general, a network domain name represents an Internet Protocol (IP) resource, such as a personal computer used
30 to access the Internet, a server computer hosting a web site, or the web site itself or any other service communicated via the Internet.

[0017] The term “domain name service (DNS) provider” generally refers to a computer that is registered to join a network domain name system. In some examples, the DNS provider runs special-purpose networking software, features a public IP address, and contains a database of network domain names and IP addresses for Internet hosts. The DNS provider
5 can be a root server that contains a complete database of Internet network domains and IP addresses, or the DNS provider may be configured to manage a subset of Internet network domain names. A DNS provider can be configured to receive requests for DNS information and to respond to such requests by transmitting DNS information (e.g., DNS resource records history) associated with one or more network domain names and/or IP addresses. DNS
10 information may include any suitable information maintained by the DNS provider, such as IP addresses associated with a particular network domain name and network domain names that host particular IP addresses. In some embodiments, DNS information may include autonomous system information associated with one or more IP addresses, one or more threat scores for one or more network domains, a number of malicious domains hosted by a
15 particular IP address, and the like.

[0018] The term “autonomous system” generally refers to a collection of connected IP routing prefixes under the control of one or more network operators on behalf of a single administrative entity or network domain name that presents a common routing policy to the Internet.

[0019] The term “threat intelligence provider” generally refers to threat intelligence management platform that is configured to gather, filter/analyze data, and provide threat intelligence information. In some embodiments, the threat intelligence provider can provide threat intelligence information via an electronic feed. Threat intelligence information may be customized by industry and/or by organization. Threat intelligence information can include
20 high-risk hosts, network domain names, malicious payloads and Internet Protocol (IP) addresses, a threat classification and/or score of an IP address and/or network domain name, malware/phishing identification information, information regarding malicious files associated with an IP address and/or a network domain, indications of criminal intent, and the like.
25

[0020] The term “blacklist provider” generally refers to a computer that is configured to
30 maintain and provide one or more IP-address-based and/or network-domain-name-based blacklists. A blacklist can be utilized to identify malicious sources such as email addresses,

users, passwords, uniform resource locators (URLs), IP addresses, network domain names, files, hashes, from which network traffic should be restricted and/or blocked because of an association with malicious activities. In some embodiments, a blacklist may include features such as a number of hits, an attack categorization score, a history-based score, and the like, where each feature may be associated with a particular malicious source. In some
5 embodiments, blacklist providers can include open-source blacklists and/or publically available blacklists such as Zeus Tracker, IP-Sum, Ransomware Tracker, Frodo Tracker, SSL Blacklist, Spam Filters, Bambenek, etc.

[0021] The term “malicious content analysis provider” generally refers to one or more
10 computers that are configured to provide forensic information associated with an IP address and/or network domain name. Forensic information can include malicious files downloaded from an IP address/network domain, malicious files communicating with an IP address/network domain, and/or malicious URLs tied to an IP address/network domain name. In some examples, forensic information can include a detection ratio and/or a last-scan date. In some embodiments, a malicious content analysis provider is a virus report manager that is
15 configured to aggregate antivirus products and/or scan engines to detect viruses and/or to verify against false positives. In some embodiments, a malicious content analysis provider may be a publically available virus detection service (e.g., VirusTotal™) that analyzes suspicious files and URLs and facilitates detection of viruses, worms, Trojans, and various
20 kinds of malware.

[0022] The term “web traffic reporting provider” generally refers to one or more computers that are configured to analyze performance of a website with respect to other websites. In some embodiments, a web traffic-reporting provider may be configured to score and/or rank a website’s performance and publically provide such scores/ranks. In some embodiments, a
25 web traffic-reporting provider may calculate various aspects of a particular website’s network traffic, such as an estimated average of daily unique visitors, an estimate number of page views over some period of time, etc. In some cases, the web traffic-reporting provider can monitor the frequency of visits and identity of visitors of a website. In some cases, the web traffic-reporting provider may assign a website a highest rank position based on determining
30 that the website has a highest combination of unique visitors and page views with respect to other websites ranked.

[0023] The term “malicious activity information” generally refers to any suitable information provided by a DNS provider, a threat intelligence provider, a blacklist provider, a malicious content analysis provider, and/or a web traffic reporting provider. Malicious activity information can include, but is not limited to: a resource record history, a security ranking, a security category score, a domain generation algorithm (DGA) score, a threat score, a popularity score, a page rank, a geographical score, a geographical diversity score, an attack severity score, a threat type severity score, a list of malicious files associated with an IP address and/or network domain name, a number of malicious files associated with an IP address and/or network domain, a list of malicious Uniform Resource Locators (URLs) associated with an IP address and/or network domain name, blacklist data, whitelist data, a history-based score, a number and/or identity of different autonomous systems associated with an IP address, metadata associated with the different autonomous systems, a stability score related to the IP address and/or network domain, or a network traffic ranking associated with the network domain.

[0024] The term “maliciousness score” is intended to refer to a value quantifying actual and/or potential malicious activities. For example, a maliciousness score can be a numerical value that corresponds to a classification label that identifies a security risk or classification. As a non-limiting example, classification labels can include “Critical,” “High,” “Medium,” and “Low,” where “Critical” classification labels indicate a highest risk severity and “Low” classification labels indicate a lowest risk severity. In some embodiments, assignment of a classification label may depend on evidence that network identifier is associated with a maliciousness score that falls within a particular range. For example, an IP address with a maliciousness score of 10 may be assigned a “Low” classification label, while another IP address with a maliciousness score of 100 may be assigned a “High” classification label.

[0025] A “machine-learning model,” may include a mathematical model that is configured to provide statistical outputs (e.g., estimates, probabilities, predictions, classifications) from a set of input variables (e.g., a feature set). In some examples, a machine-learning model may be generated and/or updated utilizing supervised machine-learning techniques. In supervised machine learning, a model is prepared through a training process where it is required to make predictions and is corrected when those predictions are incorrect or reinforced when those predictions are correct. In some embodiments, training machine-learning models may utilize

conventional supervised learning algorithms, such as J48, Naive Bayes, Logistic, Decision Table, RandomTree, etc.

DETAILED DESCRIPTION

5 [0026] The systems and methods described herein provide comprehensive analysis for identifying malicious IP addresses and/or network domain names. A cybersecurity system may utilize malicious activity information to generate a machine-learning model to calculate a maliciousness score for an IP address and/or network domain name.

10 [0027] One advantage of the systems and methods described herein is that malicious activity information from multiple sources, and statistical features calculated from such information, may be combined with machine learning techniques in order to score the maliciousness of an IP address and/or network domain name with a higher degree of fidelity than was previously achievable with conventional systems. Additionally, techniques provided herein for expanding the feature set by determining related IP addresses and/or network domain names can improve the accuracy of malicious activity analysis, resulting in more
15 accurate malicious activity detection. By utilizing the cybersecurity system(s) described herein, the multi-dimensional behavior of an IP address and/or a network domain name may be considered. Accordingly, the cybersecurity systems can provide enhanced detection and prevention capabilities that can be utilized to improve the overall security posture of an entity. Further advantages are described below.

20 I. CYBERSECURITY SYSTEM AND RELATED INFRASTRUCTURE

[0028] In some embodiments, the cybersecurity system can receive requests to identify malicious IP addresses and/or network domain names. The identification can involve determining maliciousness scores for such network identifiers (i.e., the IP addresses and/or network domain name). Such requests may be transmitted by client computer(s) 102 or any
25 suitable electronic device. In some embodiments, the cybersecurity system may evaluate the maliciousness of an IP addresses and/or a network domain without receiving a specific request to do so, such as in intermittent or regular intervals, as part of a process for updating maliciousness scores.

[0029] In response to receiving a request, the cybersecurity system may determine
30 additional network identifiers that have some relation to the network identifier for which a

malicious activity evaluation/maliciousness score has been requested. The process for determining these additional IP addresses and/or network domains will be discussed further with respect to FIGS. 4A and 4B.

5 [0030] In response to determining additional network identifiers, the cybersecurity system may request corresponding malicious activity information from a DNS provider, a threat intelligence provider, a blacklist provider, a malicious content analysis provider, a web traffic reporting provider, or any suitable combination of the above.

10 [0031] The cybersecurity system may receive at least some portion of the malicious activity information requested (e.g., from the third-party servers). The cybersecurity system may extract various features from the malicious activity information. Using the extracted features, the cybersecurity system may calculate additional features (e.g., statistical values such as standard deviations, means, medians, or the like). The extracted and/or calculated features can be utilized with the machine-learning model to determine a maliciousness score for a network identifier.

15 [0032] The cybersecurity system may associate the network identifier with a classification label corresponding to the maliciousness score. A classification label may indicate a degree of security risk (e.g., "High," "Medium," "Low," etc.). The cybersecurity system may provide the maliciousness score to the requesting client computer(s) and/or the cybersecurity system may store the maliciousness score in a data store that is suitable to store such
20 information.

[0033] Depending on the maliciousness score calculated, the cybersecurity system may perform, or cause another system to perform, various remedial actions to influence network traffic. For example, a network identifier having a maliciousness score that indicates low risk may be added to a whitelist to cause traffic associated with the network identifier to be
25 immediately allowed. Similarly, should the network identifier have a maliciousness score that indicates high risk, the network identifier may be added to a blacklist to cause traffic associated with the network identifier to be immediately rejected.

[0034] FIG. 1 is a block diagram of a system 100 for determining a maliciousness score for a network identifier, according to some embodiments. The client computer(s) 102, the
30 cybersecurity system 104, the DNS provider(s) 108, the threat intelligence provider(s) 110, the blacklist provider(s) 112, the malicious content analysis provider(s) 114, and the web

traffic reporting provider(s) 116 may individually be any suitable computing device (e.g., mainframes, desktop or laptop computers, tablets, mobile devices, wearable devices, etc.). The various provider(s) 108-116 are collectively referred to herein as “third-party servers.”

5 [0035] The various systems and/or computers depicted in FIG. 1 may be communicatively coupled by network 107. The network 107 may include any one or a combination of many different types of networks, such as cable networks, the Internet, wireless networks, cellular networks, and other private and/or public networks. It should be appreciated that the described techniques may be applied in other client/server arrangements, as well as in non-client/server arrangements (e.g., as locally stored applications, etc.).

10 [0036] In the system depicted in FIG. 1, each of the third-party servers may include one or more provider computers. Therefore, it is contemplated that multiple third-party servers (e.g., multiple DNS providers) can be utilized to obtain malicious activity information for one or more IP addresses and/or network domain names.

15 [0037] In some embodiments, the cybersecurity system 104 may be implemented as a distributed system, with various modules and functionality being provided across a set of server computers. The cybersecurity system 104 may be configured to operate as a service. For example, the cybersecurity system, operating as a service, may expose one or more application programming interfaces (APIs) to be utilized by remote systems and/or devices in order to stimulate the functionality provided by the cybersecurity system 104. The
20 cybersecurity system 104 may process request messages in TCP/IP format, HTTP format, or any suitable message format.

II. FUNCTIONALITY OF CYBERSECURITY SYSTEM

[0038] The cybersecurity system may include various engines that collectively perform operations to calculate maliciousness scores for network identifiers. The engines may include
25 a filtering engine, a parameter determination engine, and a decision engine. Depending on the maliciousness score, a classification label indicating a degree of risk (e.g., “Critical,” “High,” “Medium,” and “Low”) may be assigned to a network identifier.

[0039] FIG. 2 is a block diagram 200 of an exemplary cybersecurity system (e.g., the cybersecurity system 104) that determines maliciousness scores for a network identifiers
30 (e.g., IP addresses and/or network domain names), according to some embodiments. The

cybersecurity system 104 may include a plurality of engines that may carry out various embodiments. These engines may be software modules, hardware modules, or a combination thereof. If the engines are software modules, the engines can be embodied on a computer readable medium and processed by a processor in any of the computer systems described
5 herein. It should be noted that any engine, module, or data store described herein, may be a service responsible for managing data of the type required to make corresponding calculations. The engines may exist as part of the cybersecurity system 104, or the engines may exist as separate modules or services external to the cybersecurity system 104.

[0040] In the embodiment depicted in FIG. 2, the cybersecurity system 104 may include a
10 filtering engine 204, a parameter determination engine 206, and a decision engine 208. These engines will be discussed in more detail in the descriptions corresponding to FIGS. 3, 5, and 6.

[0041] In some embodiments, a client computer (e.g., the client computer(s) 102 of FIG. 1) can send a request message to the cybersecurity system 104 to request a maliciousness score
15 for one or more network identifiers. In some cases, the cybersecurity system 104 may automatically evaluate malicious activity without receiving a request message. For example, as part of a process for updating a maliciousness score, the cybersecurity system 104 may evaluate the maliciousness of a network identifier according to a schedule, in response to a threshold time period elapsing since a score was last calculated, etc.

[0042] The filtering engine 204 may maintain one or more whitelists that identify network identifiers that have previously been determined to be benign and/or one or more blacklists that identify network identifiers that have previously been determined to be hostile. The filtering engine 204 may be utilized to filter out network identifiers from further processing in order to preserve processing resources of the engine(s) downstream. By way of example, the
25 filtering engine 204 may receive a set of one or more network identifiers for which a maliciousness score is to be determined. The filtering engine 204 may compare the received network identifiers to the whitelist. If a particular network identifier is identified in a whitelist, the filtering engine 204 may ensure that the particular network identifier is not forwarded to the parameter determination engine 206 for further processing.

[0043] Network identifiers not contained in the whitelist may be forwarded to the
30 parameter determination engine 206. The parameter determination engine 206 may collect

malicious activity information from various malicious activity sources (e.g., the third-party servers of FIG. 1), extract and/or calculate various parameters (features) from the collected malicious activity information, and provide the extracted and/or calculated features to the decision engine 208.

5 [0044] The decision engine 208 may generate and maintain one or more machine-learning models. The decision engine 208 may utilize the features provided by the parameter determination engine 206 and one or more machine-learning models to calculate a maliciousness score for a network identifier. In some examples, the maliciousness score may correspond to a classification label (e.g., the classification labels 210). Upon determining a
10 maliciousness score, or at another suitable time, the decision engine 208 may perform, or cause another system to perform, one or more remedial actions. Such remedial actions may include, but are not limited to, restricting/allowing network traffic, providing a notification of a condition and/or risk severity, providing one or more maliciousness scores, modifying a filter list such as a blacklist and/or whitelist, or assigning/modifying an association between
15 an IP address/network domain name and a reputation label.

III. ARCHITECTURE CYBERSECURITY SYSTEM

[0045] FIG. 3 is an example computer architecture 300 capable of implementing at least some embodiments of the cybersecurity system 104. The cybersecurity system 104 may include a processor 304 and a computer readable medium 310 coupled to the processor 304,
20 the computer readable medium 310 comprising code, executable by the processor 304 for performing the functionality described herein. It should be appreciated that any functionality described with respect to the engines of FIG. 3 may be combined to be performed by a single engine or may be performed by an engine that is external to the cybersecurity system 104.

[0046] FIG. 3 shows cybersecurity system 104 communicatively coupled to the data store
25 312. The data store 312 may be configured as depicted in FIG. 2, or the data store 312 may be provided, in whole or in part, as part of the cybersecurity system 104. The data store 312 may be a conventional, fault tolerant, relational, scalable, secure database such as Oracle™ or Sybase™. The data store 312 may be implemented using various data structures, such as an array, hash map, (linked) list, structured text file (e.g., XML), table, and/or the like. Such data
30 structures may be stored in memory and/or in structured files. In some examples, the data store 312 may be configured to store maliciousness scores and/or reputation labels associated

with one or more IP addresses and/or network domain names. Additionally, or alternatively, the data store 312 may be configured to store malicious activity information received or obtained by one or more third-party servers of FIG. 1.

[0047] The processor 304 may be coupled to a system memory 306 and an external communication interface 308. The computer readable medium 310 may also be operatively coupled to the processor 304.

[0048] The computer readable medium 310 may comprise a number of software and/or hardware engines including a filtering engine 314 (e.g., the filtering engine 204 of FIG. 2), a parameter determination engine 316 (e.g., the parameter determination engine 206 of FIG. 2), a decision engine 318 (e.g., the decision engine 208 of FIG. 2), and a request processing engine 320. More or fewer software/hardware engines may be utilized to perform the functionality described herein. The engines depicted in FIG. 3 may be included in a software application stored in the computer readable medium 310 and operating on the cybersecurity system 104.

15 A. Request Processing Engine

[0049] The request processing engine 320 may receive and process requests for a malicious activity evaluation. A request message may include one or more network identifiers for which a maliciousness score is requested. In some cases, rather than received a request, the request processing engine 320 may cause the processor 304 to determine one or more IP addresses and/or one or more network domains from a data store (e.g., the data store 312). By way of example, the request processing engine 320 may determine that a maliciousness score of a network identifier is to be determined based on a time since a last maliciousness score for the network identifier was assigned. Accordingly, network identifiers having maliciousness scores over a threshold age may be reevaluated for malicious activity.

[0050] In some embodiments, the request processing engine 320 may cause maliciousness scores to be recalculated intermittently, or according to a schedule maintained by the request processing engine 320. The schedule may be predetermined or the schedule may be user-defined. In some cases, the schedule may be stored in the data store 312, or another suitable storage location, and accessible to the request processing engine 320. In some embodiments, the request processing engine 320 may receive and process scheduling creation and/or modification requests (e.g., from the client computer(s) 102 of FIG. 1).

[0051] The request processing engine 320 may identify a set of network identifiers that relate to the one or more network identifiers for which a maliciousness evaluation has been initiated. To facilitate the identification of related network identifiers, the request processing engine 320 may request and/or receive passive DNS information from one or more DNS providers (e.g., the DNS provider(s) 108 of FIG. 1). “Passive DNS information” is intended to refer to DNS information that has been constructed to provide at least a partial view of the historical DNS information available in a global Domain Name System. In some examples, the request processing engine 320 can obtain passive DNS information from a centralized data store (e.g., the data store 312 or another suitable data store configured to store such information).

B. Identification of related network identifiers

[0052] Identifying a set of related network identifiers may occur in various ways. Two examples of identifying a set of related network identifiers are provided below with respect to FIGS. 4A and 4B. Specifically, FIGS. 4A and 4B are flow diagrams 400A and 400B illustrating an example processes for determining related network identifiers from an initial set of network identifiers, in accordance with some embodiments. Specifically, the flow diagram 400A illustrates an example process for determining related IP addresses and network domains associated with an IP address (IP_i). The IP address IP_i may be one of a set of one or more of candidate IP addresses (IP_c). The process described in FIG. 4A may be utilized in examples in which a maliciousness score is requested for one or more IP addresses.

[0053] At 402, the IP address, IP_i (one of the candidate IP addresses of the set IP_c), may be received and/or obtained by the request processing engine 320 of FIG. 3 as part of the set IP_c . The set of one or more IP addresses IP_c may be received in a request, or the set IP_c may be determined by schedule or by identifying IP addresses for which a maliciousness score is over a threshold age.

[0054] At 404, given the IP address (IP_i), the request processing engine 320 may determine a domain set RD_i^1 . In some examples, passive DNS information of a DNS provider may be utilized to determine a domain set RD_i^1 that includes one or more network domains that are associated with the IP address, IP_i .

[0055] At 406, the request processing engine 320 may determine a related IP address set RI_i utilizing the passive DNS information. The set RI_i may include one or more IP addresses that are pointed to by a network domain name of the domain set RD_i^1 . As a non-limiting example, 195.154.209.132 may be the IP address IP_i . IP_i currently hosts four domains.

5 Accordingly, the set $RD_i^1 = \{\text{par81-020.ff.avast.com, ipsc-par.sl.ff.avast.com, sla-par.ff.avast.com, and slw-par.ff.avast.com}\}$. One particular domain, ipsc-par.sl.ff.avast.com, may be tied to various other IPs, which comprise the RI_i set $\{62.210.142.144,$
62.210.142.179, 62.210.142.181, 62.210.142.182, 62.210.142.183, 62.210.142.187,
62.210.142.189, 62.210.146.141, 62.210.146.154, 62.210.151.43, 62.210.152.170,
10 62.210.189.3}.

[0056] At 408, the request processing engine 320 may determine a second set of related domains RD_i^2 utilizing the passive DNS information. The set RD_i^2 may include one or more network domains that are individually associated with the IP addresses of the set RI_i . The set RI_i may include one or more IP addresses, which are pointed to by a network domain name of
15 the domain set RD_i^2 .

[0057] The flow diagram 400B of FIG. 4B illustrates an example process for determining related IP addresses and network domains associated with a network domain name (D_d). The network domain D_d may be one of a set of one or more of candidate network domain names (D_c). The process described in FIG. 4B may be utilized in examples in which a maliciousness
20 score is requested for one or more network domain names.

[0058] At 410, the network domain, D_d , may be received and/or obtained by the request processing engine 320 of FIG. 3 as described above. The network domain, D_d , may be received and/or obtained as part of the set D_c . The set of one or more network domains D_c may be received in a request, or the set of one or more network domains may be determined
25 by schedule or by identifying network domains for which a maliciousness score is over a threshold age.

[0059] At 412, the request processing engine 320 may determine a related domain set CD_d utilizing passive DNS information. The domain set CD_d may include one or more network domains that have been accessed by a same user within a period of time before/after the user
30 accessed the domain D_d .

[0060] At 414, the request processing engine 320 may cause the processor 304 to determine a related domain set RD_d utilizing passive DNS information. The related domain set RD_d may include one or more network domains that have been frequently (e.g., over a threshold frequency) requested around the same time (e.g., with a period of time) as the domain D_d , and
5 that are not frequently (e.g., under a threshold frequency) associated with other network domains.

[0061] At 416, the request processing engine 320 may determine a domain set ND_d utilizing passive DNS information. The domain set ND_d may include one or more network domains that share a same name server as the domain name D_d .

10 [0062] At 418, the request processing engine 320 may determine a domain set ED_d utilizing passive DNS information. The domain set ED_d may include one or more network domain names that are registered with the same email address as the domain name D_d .

[0063] At 420, the request processing engine 320 may determine a set of one or more IP addresses, RI_d , utilizing passive DNS information. The set RI_d may include one or more IP
15 addresses, which are historically pointed to by individual domains of the domain name D_d .

[0064] At 422, the request processing engine 320 may determine a domain set PD_d utilizing passive DNS information. The domain set PD_d may include a set of one or more network domains that are associated with one or more of the IP addresses included in the set RI_d .

[0065] Upon concluding either process described in connection with FIGS. 4A and 4B, the
20 request processing engine 320 may forward any suitable combination of the received network identifier sets (e.g., IP_1 and/or D_d) and the determined network identifier set(s) (e.g., RI_i , RD_i^1 , RD_i^2 , CD_d , RD_d , ND_d , ED_d , RI_d , PD_d) to the filtering engine 314 for further processing.

C. Filtering Engine

[0066] The filtering engine 314 may receive (e.g., from the request processing engine 320)
25 one or more network identifiers. The received network identifiers may include any suitable number and combination of the sets received and/or determined by the request processing engine 320 in the manner described above. The filtering engine 314 may access one or more filtering lists (e.g., a whitelist, a blacklist, etc.). The filtering list may be stored in data store 312 and maintained by the request processing engine 320. Additionally, or alternatively, the

filtering engine 314 may obtain a whitelist and/or blacklist from a publically available website.

[0067] The filtering engine 204 may determine whether the received network identifiers are identified in one or more whitelists. If a network identifier is identified in a whitelist or a
5 blacklist, the filtering engine 204 may be configured to ensure that the network identifier is not forwarded to the parameter determination engine 206 for further processing. By way of example, the filtering engine 204 may remove the network identifier that is identified in the
whitelist/blacklist from the set to which the network identifier belongs. Thus, the filtering engine 204 may be utilized to filter out previously determined benign or hostile network
10 identifiers in order to preserve processing resources of the engine(s) downstream (e.g., the parameter determination engine 316, the decision engine 318, etc.).

D. Parameter Determination Engine

[0068] The parameter determination engine may generally be used to extract and/or calculate features from malicious activity information. The engines provided by the
15 parameter determination engine may individually be responsible for processing malicious activity information from a particular type of third-party server. FIG. 5 is an example computer architecture of the parameter determination engine 316 of FIG. 3 that is capable of implementing at least some aspects of the cybersecurity system 104 of FIG. 3, according to some embodiments.

[0069] The engines of the parameter determination engine 316 may be software modules, hardware modules, or a combination thereof. If the engines are software modules, the engines can be embodied on a computer readable medium and processed by a processor in any of the computer systems described herein. It should be noted that any engine, module, or data store described herein, may be a service responsible for managing data of the type required to
20 make corresponding calculations. The engines depicted in FIG. 5 may exist as part of the parameter determination engine 316, or the engines may exist as separate modules or services external to the parameter determination engine 316.

[0070] In the embodiment depicted in FIG. 5, the parameter determination engine 316 may include a threat intelligence engine 502, a malicious activity content engine 504, a blacklist
25 engine 506, a DNS engine 508, a web traffic engine 510, and a data processing engine 512.

[0071] The threat intelligence engine 502 may be configured to transmit requests to one or more threat intelligence providers, such as the threat intelligence provider(s) 110 of FIG. 1, for malicious activity information associated with a one or more network identifiers. The threat intelligence engine 502 may receive malicious activity information from the threat intelligence provider(s) 110. Received malicious activity information may be stored in the data store 312, or within another suitable storage location. The stored malicious activity information may be associated with a network identifier in order to provide a mechanism for retrieving the malicious activity information after it is stored. According to some embodiments, the malicious activity information received from the threat intelligence providers may include threat scores and/or malicious files associated with a network identifier.

[0072] The malicious activity content engine 504 may transmit requests for malicious activity information to one or more malicious activity analysis providers, such as the malicious activity analysis providers(s) 114 of FIG. 1. The malicious activity content engine 504 may receive malicious activity information from the malicious activity analysis provider(s) 114. The malicious activity content engine 504 may store such malicious activity information in the data store 312, or another suitable storage location. The stored malicious activity information may be associated with a network identifier in order to provide a mechanism for retrieving the malicious activity information after it is stored. According to some embodiments, the threat intelligence engine 502 may extract forensics information (e.g., features) from the malicious activity information. Such features may include malicious files downloaded from an IP address, malicious files communicating with an IP address, and/or malicious URLs tied to an IP address.

[0073] The blacklist engine 506 may transmit requests for malicious activity information to one or more blacklist providers, such as the blacklist provider(s) 112 of FIG. 1. The blacklist engine 506 may receive malicious activity information from the blacklist provider(s) 112. The blacklist engine 506 may store such malicious activity information in the data store 312, or another suitable storage location. The stored malicious activity information may be associated with a network identifier in order to provide a mechanism for retrieving the malicious activity information after it is stored.

[0074] According to some embodiments, the blacklist engine 506 may maintain a Unified Blacklist Database (UBDB) that includes blacklist information collected from a number of obtained blacklists. By way of example, the blacklist engine 506 may access various open source blacklists, such as from Zeus Tracker, Ransomware Tracker, Frodo Tracker, SSL
5 Blacklist, Spam Filters, and Google safe browsing. The various open source blacklists may be merged and stored in the UBDB for later consultation. The blacklist engine 506 may extract current hit and/or historic hits for an network identifier listed in the Unified Blacklist Database (UBDB).

[0075] The DNS engine 508 may transmit requests for malicious activity information to
10 one or more DNS providers, such as the DNS provider(s) 108 of FIG. 1. The DNS engine 508 may receive malicious activity information from the DNS provider(s) 108. The DNS engine 508 may store such malicious activity information in the data store 312, or another suitable storage location. The stored malicious activity information may be associated with a network identifier in order to provide a mechanism for retrieving the malicious activity information
15 after it is stored. According to some embodiments, the DNS engine 508 may receive DNS information from a DNS provider. DNS information received may include any suitable combination of resource record (RR) history for IPs and domains, autonomous system information associated with a multiple IP addresses, threat scores for network domain names, security scores associated with a domain, co-occurrence domain sets (e.g., the set CD_a of FIG.
20 4B), related domain sets, malicious files associated with a domain, and number of malicious domains hosted on the IP.

[0076] The web traffic engine 510 may transmit requests for malicious activity information to one or more web traffic reporting providers, such as the web traffic reporting provider(s)
25 116 of FIG. 1. The web traffic engine 510 may receive such malicious activity information from the web traffic reporting provider(s) 116. The web traffic engine 510 may store malicious activity information in the data store 312, or another suitable storage location. The stored malicious activity information may be associated with a network identifier in order to provide a mechanism for retrieving the malicious activity information at a time subsequent to storage. According to some embodiments, the malicious activity information received from
30 the web traffic reporting provider(s) 116 may include a popularity ranking or score assigned to an IP address and/or network domain name by the third-party.

[0077] The feature processing engine 512 may collect and/or access the malicious activity information received by the various third-party servers discussed above. In some examples, the feature processing engine 512 can extract suitable combinations of features (e.g., external score-based features, forensics features, blacklist-based features, behavior-based features, and network-based, discussed below) from malicious activity information provided by the third-party servers. After extracting the features from the malicious activity information, the feature processing engine 512 may compute a number of statistical features from the extracted features (e.g., summations, means, medians, etc.). The feature processing engine 512 may transmit the extracted features and computed statistical features to the decision engine 318 of FIG. 3 as part of a feature set.

[0078] In some embodiments, the feature processing engine 512 may calculate a freshness score for any of the features discussed below. A freshness score may be utilized to weigh features extracted/computed from recent malicious activity information more heavily than features extracted/computed from older malicious activity information. The features may be weighed according to a time of receipt associated with the malicious activity information. In still further examples, the features may be weighed according to a timestamp assigned by a third-party server and provided with the malicious activity information. A weight value may be determined for features of received malicious activity information based on calculating $(-Log(x^3)) + 6 * 5$, where x is the difference between the time of receipt or timestamp date of the malicious activity information and a current date in days. This weight values allows for more recent malicious activity information to have a greater impact than that of older malicious activity information.

a) Example features

[0079] As discussed above, the feature processing engine 512 may extract or compute external-based features, forensics features, blacklist-based features, behavior-based features, and network-based from the malicious activity information received. Such features may be used to compute additional features (e.g., statistical values such as summations, means, medians, or the like, computed from the extracted features). The extracted features and computed statistical values may be used later as input for a machine-learning model. Such features are discussed further below.

a) *External-Based Features related to an IP address*

[0080] In some embodiments, external-based features listed below may be used to determine a maliciousness score for individual IP addresses. The external-based features may be extracted or computed by the feature processing engine 512. These external-based features
5 may include threat scores received from the threat intelligence provider(s) 110 of FIG. 1. The threat scores may include one or more threat scores associated with each IP address of the RI_i , RD_i^1 , and RD_i^2 sets discussed in connection with FIG. 4A. The external-based features may be weighed according to a freshness score such that more recent features effect a maliciousness score more heavily than older features.

10 [0081] Specific examples of external-based features extracted/computed by the feature processing engine 512 include, but are not limited to:

- A threat score provided by a threat intelligence source of an IP address of the set IP_i .
- A threat score of an IP address IP_i calculated from a publicly available IP
15 report.
- At least one of a summation, mean, median, or standard deviation of at least one of the following:
 - One or more threat scores provided by one or more threat intelligence sources for each IP belonging to the set RI_i .
 - 20 ○ One or more threat scores calculated from a publicly available IP report for each IP belonging to the set RI_i .
 - One or more threat scores provided by one or more DNS providers (e.g., the DNS provider(s) 108 of FIG. 1) for each domain belonging to the set RD_i^1 .
 - 25 ○ One or more threat scores provided by one or more DNS providers (e.g., the DNS provider(s) 108) for each domain belonging to the set RD_i^2 .
 - One or more threat scores calculated from a publicly available domain report for each domain belonging to the set RD_i^1 .
 - 30 ○ One or more threat scores calculated from a publicly available domain report for each domain belonging to the set RD_i^2 .

- One or more popularity rankings/scores (e.g., provided by the web traffic reporting provider(s) 116 of FIG. 1) for each domain belonging to the set RD_i^1 .
- One or more popularity rankings/scores (e.g., provided by the web traffic reporting provider(s) 116 of FIG. 1) for each domain belonging to the set RD_i^2 .

b) External-Based Features related to a network domain

[0082] In some embodiments, external-based features listed below may be used to determine a maliciousness score for individual network domain name D_d or the network domain name sets CD_d , RD_d , ND_d , ED_d , PD_d sets described above in connection with FIG. 4B. The external-based features may be extracted or computed by the feature processing engine 512. The external-based features used to determine a maliciousness score of a network domain name may be weighed according to a freshness score such that more recent features effect a maliciousness score more heavily than older features.

[0083] Specific examples of external-based features extracted/computed by the feature processing engine 512 may include at least one of:

- One or more security ranks of a network domain provided by one or more DNS service providers (e.g., the DNS provider(s) 108 of FIG. 1).
- One or more threat scores of a network domain calculated from one or more publicly available domain reports.
- One or more security category scores of a network domain name calculated from malicious activity information provided by one or more DNS service providers (e.g., the DNS provider(s) 108).
- One or more content category scores of a network domain name calculated from malicious activity information provided by one or more DNS service providers (e.g., the DNS provider(s) 108).
- One or more Domain Generation Algorithm (DGA) scores of a network domain name provided by one or more DNS service providers (e.g., the DNS provider(s) 108). The DGA scores, in some examples, may quantify a

probability value indicating the likelihood that a network domain name was generated by an algorithm rather than a human.

- One or more page ranks of a network domain provided by one or more DNS service providers. In some examples, a page rank may reflect the popularity of a network domain name according to Google's PageRank algorithm.
- One or more popularity scores, which in some examples, may be a number of unique client IP addresses that have visited a network domain name, relative to the all requests to all network domain names. In some examples, this information may be provided by a DNS service provider.
- One or more Geoscores of a network domain provided by one or more DNS service providers. In some examples, a Geoscore may represent how far apart the different physical locations serving the network domain name are with respect to one another.
- One or more TLD Geo-diversity scores. In some cases, a TLD Geo-diversity score may be represented as a percentage of clients visiting the network domain name.
- One or more attack severity scores of a network domain name. In some examples, an attack severity score may be calculated from the malicious activity information provided by a DNS service provider.
- One or more severity scores for the threat type of a network domain name. In some examples, the severity score for a threat type may be calculated from the maliciousness activity information provided by a DNS service provider.

c) Example forensic features related to an IP address

- [0084] In some embodiments, forensic features listed below may be used to determine a maliciousness score for one or more IP addresses. In some examples, the feature processing engine 512 may extract malicious URLs, malicious files, and/or malicious communicating files from malicious activity information. The related malicious URL set U_i for the IP address IP_i may include one or more URLs that have been historically tied with the IP address IP_i .
- The related malicious file set MF_i , corresponding to IP_i , may include a set of one or more files that are identified as possibly malicious files and that are tied with the IP address IP_i .
- The related malicious communicating file set CF_i for IP_i , may include one or more possibly

malicious files that have been communicated to/from the IP address IP_i . The forensic features used to determine a maliciousness score of an IP address may be weighed according to a freshness score, such that more recent features effect a maliciousness score more heavily than older features.

5 [0085] In some embodiments, the feature processing engine 512 may identify suspicious files and/or URLs related to an IP address IP_i from the received malicious activity information. In some examples, a number of suspicious URLs, and a number of files in the MF_i and CF_i sets may be identified. Additionally, or alternatively, the malicious activity information may provide a list of malware tied to an IP address IP_i . The feature processing engine 512 may assign each of the possibly malicious URLs and files a threat score based on 10 a last-detected date and a score provided by threat intelligence source. The last detection date may be used to measure the freshness score of the evidence. Statistical features such as a summation, a mean, a median, and a standard deviation may be calculated from the threat score for each of the sets U_i , MF_i , and CF_i .

15 [0086] In accordance with at least one embodiment, the feature processing engine 512 may extract at least one forensics feature for an IP address IP_i . Specific examples of forensic features extracted/computed by the feature processing engine 512 may include at least one of:

- A total number of elements in the related malicious file set MF_i .
- 20 • A total number of elements in the related communicating file set CF_i .
- A total number of elements in the related malicious URL set U_i .
- A percentage of malicious domains hosted by the IP address IP_i .
- At least one of a summation, a mean, a median, or a standard deviation of at least one of the following:
 - 25 ○ One or more threat scores for the file set MF_i .
 - One or more threat scores for the file set CF_i .
 - One or more threat scores for the URL set U_i .

d) Example forensic features related to a network domain

30 [0087] In some embodiments, forensic features listed below may be used to determine a maliciousness score for one or more network domain names. The feature processing engine

512 may extract malicious URLs, malicious files, and/or malicious communicating files from malicious activity information. The related malicious URL set U_d for a network domain D_d may include one or more URLs that have been historically tied with the network domain D_d . The related malicious file set MF_d corresponding to D_d may include a set of one or more files that are identified as possibly malicious files and that are tied with the network domain D_d . The related malicious communicating file set CF_d corresponding to D_d may include one or more possibly malicious files that have been communicated to/from the network domain D_d . The feature processing engine 512 may calculate a percentage of malicious domains hosted by one or more IP addresses belonging to the set RI_d . The forensic features used to determine a maliciousness score of a network domain may be weighed according to a freshness score, such that more recent features effect a maliciousness score more heavily than older features.

[0088] A number of suspicious URLs, and a number of files in the MF_d and CF_d sets may be identified from received malicious activity information. The malicious activity information may provide a list of malware tied to the network domain D_d . The feature processing engine 512 may assign each of the possibly malicious URLs and/or files a threat score based on a last-detected date and a score provided by threat intelligence source. The last detection date may be used to measure the freshness score of the evidence. Statistical features such as a summation, a mean, a median, and a standard deviation may be calculated from the threat score for each of the sets U_d , MF_d , and CF_d .

[0089] Specific examples of forensic features extracted/computed by the feature processing engine 512 may include at least one of the following:

- A total number of elements in the related malicious file set MF_d .
- A total number of elements in the related communicating file set CF_d .
- A total number of elements in the related malicious URL set U_d .
- At least one of a summation, a mean, a median, and a standard deviation of at least one of the following:
 - One or more threat scores for one or more files of the set MF_d .
 - One or more threat scores for one or more files of the set CF_d .
 - One or more threat scores for one or more URLs of the set U_d .
 - A percentage of malicious network domains tied to an IP address of the set RI_d .

e) Example blacklist-based features related to an IP address

[0090] In some embodiments, blacklist-based features listed below may be used to determine a maliciousness score for an IP address IP_i , described above in connection with FIG. 4A. The black-list features may be extracted or computed by the feature processing engine 512. The blacklist-based features may be weighted according to a freshness score such that more recent features effect a maliciousness score more heavily than older features. In some examples, malicious activity information may be obtained by one or more blacklist providers (e.g., the blacklist provider(s) 112 of FIG. 1). Additionally, or alternatively, malicious activity information may be obtained from a unified blacklist databased (UBDB) stored in the data store 312 of FIG. 3. The UBDB may contain malicious activity information previously obtained from one or more open source blacklists.

[0091] The feature processing engine 512 may calculate a number of hits for an IP address IP_i . The feature processing engine 512 may calculate an attack severity score based on the particular blacklist that identified the IP address. In some examples, a higher severity score may be assigned to an IP address that is used for command and control (C2) activity than an a severity score assigned to an IP address that has been used to send spam emails. To calculate an attack severity score, the feature processing engine 512 may utilize malicious activity information provided by one or more threat intelligence sources (e.g., the threat intelligence sources 110 of FIG. 1). In some examples, an attack severity score may be calculated from the malicious activity information received from the threat intelligence source(s) utilizing various categories of attack with which an IP address is associated. From the UBDB, a history-based score may be calculated for the IP address IP_i .

[0092] Specific examples of blacklist-based features extracted/computed by the feature processing engine 512 may include one or more of the following:

- A total number of hits within various blacklists for the IP address IP_i .
- An attack severity score the IP address IP_i .
- A history-based score for the IP address IP_i .
- At least one of a summation, a mean, a median, or a standard deviation of at least one of the following:
 - A hit count for all IP addresses belonging to the set RI_i .

- One or more attack severity scores for all IP addresses belonging to the set RI_i .
- One or more history-based scores for all IPs belonging to the set RI_i .

f) Example blacklist-based features related to a network domain

- 5 [0093] The blacklist-based features listed below may be used to determine a maliciousness score for a network domain D_d , described above in connection with FIG. 4B. The blacklist-based features may be extracted or computed by the feature processing engine 512. In some embodiments, the blacklist-based features may be weighted according to a freshness score such that more recent features effect a maliciousness score more heavily than older features.
- 10 In some examples, malicious activity information may be obtained by one or more blacklist providers (e.g., the blacklist provider(s) 112 of FIG. 1). Additionally, or alternatively, malicious activity information may be obtained from a unified blacklist databased (UBDB) stored in the data store 312 of FIG. 3. The UBDB may contain malicious activity information previously obtained from one or more open source blacklists.
- 15 [0094] The feature processing engine 512 may calculate a number of hits for the network domain D_d . The feature processing engine 512 may calculate an attack severity score based on the particular blacklist that identified the network domain. In some examples, a higher severity score may be assigned to a network domain that is used for command and control (C2) activity than a severity score assigned to a network domain that has been used to send
- 20 spam emails. To calculate an attack severity score, the feature processing engine 512 may utilize malicious activity information provided by one or more threat intelligence sources (e.g., the threat intelligence sources 110 of FIG. 1). In some examples, an attack severity score may be calculated from the malicious activity information received from the threat intelligence source(s) utilizing various categories of attack with which a network domain is
- 25 associated. From the UBDB, a history-based score may be calculated for a network domain.

[0095] Specific examples of external-based features extracted/computed by the feature processing engine 512 may include one or more of the following:

- A total number of hits within various blacklists for a network domain D_d .
- An attack severity score of a network domain D_d .
- A history-based score of a network domain D_d .

30

g) Example network-based features related to an IP address

[0096] The network-based features listed below may be used to determine a maliciousness score for an IP address IP_i , described above in connection with FIG. 4A. Such features may
5 be extracted or computed by the feature processing engine 512. The network-based features may be weighted according to a freshness score such that more recent features effect a maliciousness score more heavily than older features. In some examples, malicious activity information may be obtained by one or more blacklist providers (e.g., the blacklist provider(s) 112 of FIG. 1). Additionally, or alternatively, malicious activity information may
10 be obtained from a unified blacklist database (UBDB) stored in the data store 312 of FIG. 3. The UBDB may contain malicious activity information previously obtained from one or more open source blacklists.

[0097] It should be appreciated that adversaries may often use "bullet-proof" hosting services to launch attacks in order to avoid law enforcement and other legal repercussions.
15 Additionally, certain autonomous systems (ASs) are known for to have a higher incidence of malicious activity than other ASs. Accordingly, in some embodiments, threats associated with an AS may be used to determine the severity level of an IP address IP_i . If an IP address is tied with an AS which has a history of malicious activity, that IP address may be considered to be a higher severity than other IP addresses that have no association to an AS
20 with a history of malicious activity. To determine the threat associated with an AS, the feature processing engine 512 may calculate threat scores of one or more IP addresses of the set IP_i utilizing malicious activity information provided by various threat intelligence sources and blacklist sources. The threat scores of an IP address IP_i may further depend on malicious activity information provided by one or more DNS providers.

[0098] In some embodiments, the number of ASs associated with an IP directly and indirectly, the number of registration records, and the frequency of new registrations of the ASs can provide information about the life cycle of an IP address. The feature processing engine 512 may calculate the number of ASs associated with an IP address (directly or indirectly), the number of registration records, and the frequency of new registrations of the
30 ASs. In some cases, the higher these numbers are for an IP address, the higher the possibility that the IP may be considered risky.

[0099] Specific examples of network-based features extracted/computed by the feature processing engine 512 may include one or more of the following:

- A number of different Autonomous Systems (ASs) of the set AI_i .
- A number of registration entries corresponding to an AS belonging to the set AI_i .
- An average frequency of new registrations associated with an AS belonging to the set AI_i .
- The number of suspicious activities associated with an AS belonging to the set AI_i .
- At least one of a summation, a mean, a median, or a standard deviation for at least one of the following:
 - One or more threat scores provided by an intelligence source for each IP address belonging to the set AI_i .
 - One or more threat scores calculated from a publicly available IP reports for one or more IP addresses belonging to the set AI_i .
 - One or more blacklist hit counts for one or more IP addresses belonging to the set AI_i .
 - One or more attack severity scores for one or more IP addresses belonging to the set AI_i .
 - One or more history-based scores for one or more IP addresses belonging to the set AI_i .

h) Example network-based features related to a network domain

[0100] The network-based features listed below may be used to determine a maliciousness score for a network domain D_a , described above in connection with FIG. 4B. The network-based features used to determine a maliciousness score of one or more network domains may be weighted according to a freshness score such that more recent features effect a maliciousness score more heavily than older features. In some examples, malicious activity information may be obtained by one or more blacklist providers (e.g., the blacklist provider(s) 112 of FIG. 1). Additionally, or alternatively, malicious activity information may be obtained from a unified blacklist databased (UBDB) stored in the data store 312 of FIG. 3.

The UBDB may contain malicious activity information previously obtained from one or more open source blacklists.

[0101] As described above, adversaries may often use "bullet-proof" hosting services to launch attacks in order to avoid law enforcement and other legal repercussions and certain autonomous systems (ASs) are known for to have a higher incidence of malicious activity than other ASs. Therefore, it is important to identify the risks associated with the network where a network domain is hosted. Accordingly, in some embodiments, threats associated with the ND_d , ED_d , PD_d , RI_d , and AS sets may be used to determine the severity level of a network domain. As an example, network domain names that are associated with a name server, AS, or an email address that has a history of malicious activities, may be considered to be more risky to an infrastructure than other domain names that do not have such ties. In some cases, network domains that are associated with to higher number of ASs may be considered more risky than netowrk domains tied to fewer ASs.

[0102] To determine the threat associated with the network identifiers belonging to the sets ND_d , ED_d , PD_d , RI_d , and AS, the feature processing engine 512 may calculate threat scores of a network domain D_d utilizing malicious activity information provided by various threat intelligence sources and/or blacklist sources. The threat scores of a network domain D_d may further depend on malicious activity information provided by one or more DNS providers.

[0103] Specific examples of network-based features extracted/computed by the feature processing engine 512 may include one or more of the following:

- The number of different ASs belonging to the set AS.
- The number of registration entries found for one or more of the ASs belonging to the set AS.
- The average frequency of new registration of one or more of the ASs belonging to the set AS.
- The number of suspicious activities tied to an AS belonging to the set AS.
- One or more AS reputation scores provided by one or more DNS providers.
- One or more RIP scores that indicate a ranking for one or more IP addresses related to a network domain D_d . In some examples, the RIP scores may be provided by one or more DNS providers.

- One or more stability scores for one or more IP addresses related to a network domain D_d . In some examples, the stability scores may be provided by one or more DNS providers.
- One or more prefix ranks of a network domain D_d given associated IP prefixes and a reputation score of these prefixes.
- The probability of a network domain D_d being a fast flux candidate
- At least one of a summation, a mean, a median, or a standard deviation for at least one of the following:
 - One or more security ranks for one or more network domains of the set ND_d
 - One or more security ranks for one or more network domains of the set ED_d
 - One or more security ranks for one or more network domains of the set PD_d
 - One or more DGA scores for one or more network domains of the set ND_d
 - One or more DGA scores for one or more network domains of the set ED_d
 - One or more DGA scores for one or more network domains of the set PD_d
 - One or more attack severity scores for one or more network domains of the set ND_d
 - One or more attack severity scores for one or more network domains of the set ED_d
 - One or more attack severity scores for one or more network domains of the set PD_d
 - One or more threat type severity scores for one or more network domains of the set ND_d
 - One or more threat type severity scores for one or more network domains of the set ED_d
 - One or more threat type severity scores for one or more network domains of the set PD_d

- One or more blacklist hit counts for one or more network domains belonging to the set ND_d
- One or more blacklist hit counts for one or more network domains belonging to the set ED_d
- 5 ○ One or more blacklist hit counts for one or more network domains belonging to the set PD_d
- One or more threat scores provided by a threat intelligence sources for each IP address belongs to RI_d .

10 *i) Example behavior-based features related to a network domain*

[0104] The behavior-based features listed below may be used to determine a maliciousness score for a network domain D_d , described above in connection with FIG. 4B. The behavior-based features used to determine a maliciousness score of a network domain may be weighted according to a freshness score such that more recent features effect a maliciousness score
15 more heavily than older features.

[0105] The feature processing engine 512 may receive malicious activity information from one or more DNS providers. The received malicious activity information may include a list of co-occurrence domains, CD_d , and a list of related domains, RD_d . As discussed above in connection with FIG. 4A, the network domains belonging to the set RD_d include the network
20 domains requested around the same time as the network domain D_d . Therefore, the number of malicious domains in CD_d and RD_d can provide valuable information about the malicious behavior of a network domain. For example, a high number of DGA domain names in the CD_d and RD_d sets can indicate the presence of malware. The feature processing engine 512 may calculate statistical values from the received malicious activity information.

25 [0106] Specific examples of behavior-based features extracted/computed by the feature processing engine 512 may include at least one of the following:

- One or more security ranks for the network domains belonging to the set CD_d .
- One or more security ranks for the network domains belonging to the set RD_d .
- One or more DGA scores for the network domains belonging to the set CD_d .
- 30 • One or more DGA scores for the network domains belonging to the set RD_d .

- One or more security category scores for the network domains belonging to the set CD_d .
- One or more security category scores for the network domains belonging to the set RD_d .
- 5 • One or more attack severity scores for the network domains belonging to the set CD_d .
- One or more attack severity scores for the network domains belonging to the set RD_d .
- One or more threat type severity scores for the network domains belonging to the set CD_d .
- 10 • One or more threat type severity scores for the network domains belonging to the set RD_d .
- One or more blacklist hit counts for one or more network domains belonging to the set CD_d .
- 15 • One or more black list hit counts for one or more network domains belonging to the set RD_d .

E. Decision Engine

[0107] A decision engine may be utilized to access a machine-learning algorithm and the features provided by the parameter determination engine 316 to determine a maliciousness score for a network identifier. FIG. 6 is an example computer architecture of the decision engine 318 of FIG. 3, which is capable of implementing at least some aspects of the cybersecurity system of FIG. 3, according to some embodiments.

[0108] The decision engine 208 may include a plurality of engines that may carry out various embodiments. These engines may be software modules, hardware modules, or a combination thereof. If the engines are software modules, the engines can be embodied on a computer readable medium and processed by a processor in any of the computer systems described herein. It should be noted that any engine, module, or data store described herein, may be a service responsible for managing data of the type required to make corresponding calculations. The engines depicted in FIG. 6 may be exist as part of the decision engine 318, or the engines may exist as separate modules or services external to the decision engine 318.

[0109] In the embodiment depicted in FIG. 6, the decision engine 318 may include a model training engine 602, a scoring engine 604, and a remedial action engine 606. Such engines will be discussed in more detail below.

5

Model training and updating

[0110] The model training engine 602 may train and/or update one or more machine-learning models. The model training engine 602 may employ supervised-learning techniques to train and/or update one or more machine learning models.

[0111] The model training engine 602 may receive a feature set provided by the parameter
10 determination engine 316 of FIG. 3. In some examples, communication between the parameter determination engine 316 and the decision engine 318 may occur through a persistent message queue for failover recovery purposes. To train the machine-learning model, the model training engine 602 may obtain training data including a list of network identifiers (e.g., one or more IP addresses and/or one or more network domain names) with
15 an preassigned maliciousness scores. In some examples, the training data may include historical maliciousness activity information obtained from one or more malicious activity sources (e.g., the third-party servers described in connection to FIG. 1).

[0112] Using the training data, the model training engine 602 may generate and train one or
20 more machine-learning models utilizing the features set and conventional supervised-learning algorithms, such as J48, Naive Bayes, Logistic, Decision Table, RandomTree, etc. The model training engine 602 may score each of the trained machine learning models according to its output accuracy. By way of example, the model training engine 602 may input parameters of the feature set corresponding to a single network identifier and determine a degree of
25 accuracy of the output provided by the model. The output of the model may be compared to the preassigned maliciousness score corresponding to the network identifier. Accordingly, the model training engine 602 may score each model according to a degree of accuracy of the output of the model to the preassigned maliciousness scores provided in the training data. In some examples a highest scored model may be selected to provide maliciousness scores for network identifiers.

[0113] Multiple machine-learning models may be retrained with the updated historical malicious activity information and the model training engine 602 may select a machine learning model (e.g., the machine learning model 704 or another machine learning model) based on a degree of accuracy between the respective output of a model and the updated
5 historical malicious activity information.

[0114] The scoring engine 604 may be configured to receive (e.g., from the parameter determination engine 316 of FIG. 3) a feature set associated with one or more network identifiers. The feature set may include any combination of the features described above in connection with the parameter determination engine 316). The scoring engine 604 may
10 utilize the selected machine-learning model to output a maliciousness score for the network identifier corresponding to the feature set. A maliciousness score can be a numerical value. In some cases, the maliciousness score can corresponds to a classification label that identifies a security risk. As a non-limiting example, reputation labels may include "Critical," "High," "Medium," and "Low," where "Critical" reputation labels indicate a highest risk severity and
15 "Low" reputation labels indicate a lowest risk severity. The output may be stored as historical malicious activity information to be used to update/train machine-learning algorithms.

[0115] FIG. 7 is a flow diagram 700 illustrating an example process for training and updating a machine-learning model, according to some embodiments. The process depicted in FIG. 7 may be performed by the model training engine 602.

20 [0116] At 702, the model training engine 602 may select machine-learning model 704 from a set of trained machine learning models according to a score. The score may be initially determined in the manner described above. The machine-learning model 704 may be stored in data store 312 of FIG. 3 or another location accessible to scoring engine 606 of FIG. 6.

[0117] At 706, the scoring engine 606 may access the machine learning model. Utilizing
25 common feature types as those used to train the machine-learning model 704, the scoring engine 606 may input such features into the machine-learning model 704 to calculate a maliciousness score. By way of example, the scoring engine 606 may input the features associated with an IP address into the machine-learning model 704 and receive a maliciousness score for the IP address.

[0118] At 708, the maliciousness scores may be stored in the maliciousness scores data store 710. The maliciousness scores data store 710 may be included as part of the data store 312 of FIG. 3.

[0119] At 712, the maliciousness scores of the maliciousness scores data store 710 may be provided to the model training engine 602. The model training engine 602 may utilize the maliciousness scores (separately, or in conjunction with the historical malicious activity information) to retrain the machine-learning model 704. In some examples, the maliciousness scores may be stored as part of the historical malicious activity information for future training activities.

[0120] The process depicted in FIG. 7 may be performed any suitable number of times. Accordingly, as time progresses, the machine learning model 704 increases the efficiency and accuracy of determining maliciousness score by utilizing an ever-expanding training data set.

Freshness considerations

[0121] The rapidly evolving nature of network identifiers over time can present challenges to an cybersecurity system such as those described herein. Additionally, network identifiers which are found to be malicious more recently potentially pose a greater security threat to an infrastructure than a network identifier that was determined to be malicious long ago. Accordingly, the machine learning models described herein may, in some embodiments, weigh features extracted and/or determined from recent malicious activity information heavier than features extracted and/or determined from older malicious activity information. In some examples, the features may be weighed according to a time of receipt associated with the malicious activity information. In still further examples, the features may be weighed according to a timestamp assigned by a third-party server and provided with the malicious activity information. A weight value may be determined for features of received malicious activity information based on calculating $(-Log(x^3)) + 6 * 5$, where x is the difference between the time of receipt or timestamp date of the malicious activity information and a current date in days. Thus, allowing more recent malicious activity information to have a greater impact than that of older malicious activity information.

Performing Remedial Actions

[0122] One function of the decision engine 318 is to perform one or more remedial actions in accordance with a determined maliciousness score. The remedial action engine 608, a

module of the decision engine 318, may be configured to perform such actions. Remedial actions may include, but are not limited to, restricting/allowing network traffic, notifying a user and/or remote device of network identifier with a reputation label corresponding to a risk severity, providing one or more maliciousness scores, modifying a filter list such as a
5 blacklist or whitelist, or assigning/modifying an association between network identifier and a reputation label or maliciousness score. It is contemplated that a single remedial action may include any suitable number and combination of the particular remedial actions listed above. In some examples, the remedial action is not necessarily performed by the remedial action engine 608, but rather the remedial action engine 608 may stimulate a separate device to
10 perform at least some portion of the remedial action(s).

[0123] As a non-limiting example, the remedial action engine 606 may receive a maliciousness score provided by the scoring engine 604. The remedial action engine 606 may consult a protocol set to determine one or more remedial actions to initiate in response to a particular maliciousness score. By way of example, the protocol set may include rules,
15 procedures, and/or mappings that indicate a particular set of one or more remedial actions that are to be performed in response to a particular maliciousness score.

[0124] By way of example, the protocol set may indicate that network identifiers that have been associated with a “Critical” reputation label (or maliciousness score over a threshold value) should be added to a particular blacklist. As another example, the protocol set may
20 indicate that network identifiers that have been associated with a “Low” reputation label (or maliciousness score under a threshold value) should be added to a particular whitelist. The protocol set may indicate particular entities should be notified when particular network identifiers have been associated with a “High” reputation label. The possible variations available within such a protocol set are numerous and may be contemplated by one having
25 skill in the art of network traffic administration.

IV. METHOD FOR DETERMINING MALICIOUSNESS SCORES

[0125] A method for determining a maliciousness score for one or more network identifiers (IP addresses and/or network domains) is described below with reference to FIG. 8. This method can be implemented by the computer systems described above with respect to FIGS.
30 1, 2, 3, 5, and 6, for example.

[0126] FIG. 8 shows a flow chart of an exemplary method 800 for determining a maliciousness score for one or more network identifiers, in accordance with some embodiments. The method can be performed by a cybersecurity system (e.g., the cybersecurity system 104 of FIG. 2).

5 [0127] At 802, a request to evaluate malicious activity with respect to a set of one or more network identifiers may be received (e.g., by the request processing engine 320 of FIG. 3). In some examples, the request may have been initiated by the client computer(s) 102 of FIG. 1. In other embodiments, the set of one or more network identifiers may be obtained from a data store (e.g., the data store 312), rather than being received as in the method 800.

10 [0128] At 804, at least one of a set of one or more related network domains or a set of one or more related IP addresses associated with the individual network identifiers of the set of network identifiers may be determined (e.g., by the parameter determination engine 316 or modules of the parameter determination engine 316). If the set of network identifiers includes at least one IP address, related domains and related IP addresses may be determined for the IP
15 address in accordance with the process described above with respect to FIG. 4A. If the set of network identifiers includes at least one network domain, then related domains and related IP addresses may be determined for the network domain in accordance with the process described above with respect to FIG. 4B. If the set of one or more network domains comprises both an IP address and a network domain, then the corresponding processes of
20 FIGS. 4A and 4B may be performed to determine related network domains and related IP addresses.

[0129] At 806, requests to provide malicious activity information may be sent (e.g., by one or more engines of FIG. 5) to a plurality of third-party servers (e.g., the third-party servers of FIG. 1). The third-party servers may comprise any suitable combination of the DNS
25 provider(s) 108, the threat intelligence provider(s) 110, the blacklist provider(s) 112, the malicious activity analysis provider(s) 114, and/or the web traffic reporting provider(s) 116. The requests may relate to the set of one or more network identifiers, the set of one or more related network domains, and the set of one or more related IP addresses, although any suitable combination of the above may be utilized.

30 [0130] At 808, at least a portion of the malicious activity information requested from the plurality of third-party servers may be received (e.g., by the corresponding engines of FIG. 5,

namely the threat intelligence engine 502, the malicious activity content engine 504, the blacklist engine 506, the DNS engine 508, and the web traffic engine 510). In some embodiments, the malicious activity information may be stored by each engine of FIG. 5, or the engines of FIG. 5 may cause the malicious activity information to be stored (e.g., in data store 312). The malicious activity information received by each third-party server may be stored any suitable storage container and associated with a corresponding network identifier for later retrieval.

[0131] At 810, a first set of features may be extracted (e.g., by the feature processing engine 512 of FIG. 5) from the malicious activity information received at 808. The first set of features may include any suitable combination of the features described above in connection with the feature processing engine 512. In some embodiments, the first set of features may exclude the statistical values related to a summation, a mean, a median, and standard deviation described above.

[0132] At 812, a second set of features may be calculated (e.g., by the feature processing engine 512) from the first set of features extracted at 810. The second set of features may include any suitable statistical values (e.g., summation, mean, median, standard deviation, counts, hits, etc. such as those calculated from external-based features) described above. The second set of features may include at least one statistical value calculated from a common feature of the first set of features received from each of the plurality of third-party servers. For example, a summation may be calculated by utilizing a same feature (e.g., a threat score associated with an IP address) of one or more IP addresses, such that the feature values associated with each of the one or more IP addresses may be combined to provide the summation.

[0133] At 814, one or more maliciousness scores may be calculated (e.g., by the scoring engine 604) for the set of one or more network identifiers utilizing a machine learning model. The maliciousness scores may correspond to a reputation label (e.g., "Critical," "High," "Medium," "Low," etc.). The machine-learning model may utilize the first set of features and the second set of features as input. In some examples, the machine-learning model may be trained utilizing historical malicious activity information from one or more of the plurality of third-party servers

[0134] At 816, a remedial action may be performed (e.g., by the remedial action engine 606 of FIG. 6) based on the one or more maliciousness scores calculated at 814. In some embodiments, a portion or the entire remedial action may be performed by a device that is separate from the cybersecurity system 104. Remedial actions may include, but are not limited to, restricting/allowing network traffic, notifying a user and/or remote device of a condition and/or risk severity, providing one or more maliciousness scores, modifying a filter list such as a blacklist or whitelist, or assigning/modifying an association between an IP address/network domain name and a reputation label. The remedial action performed at 816 may include any suitable number and combination of the remedial actions listed above.

10 V. TECHNICAL BENEFITS

[0135] The systems and methods provided herein can be used in a variety of security controls such as firewalls, Intrusion Detection Systems (IDSs), etc.. Due to the low fidelity of conventional IP reports, malicious activity can go undetected. If undetected, Security Incidents and Events Monitoring systems (SIEMs) cannot be alerted for possibly malicious IP addresses seen in the network. Utilizing the systems and methods provided herein, multiple and disparate malicious activity information sources may be utilized to analyze malicious activity information related to one or more network identifiers. By utilizing the sources and feature sets provided herein, a risk associated with malicious activity may be more accurately determined and thus, higher fidelity alerts may be provided to other cyberdefense systems such as SIEMs and/or IDSs.

[0136] In some embodiments, the decision engine described herein runs behind a web server and supports an application programming interface (API). Accordingly, the functionality for determining a maliciousness score and/or reputation labels associated with one or more network identifiers may be provided as a service. As a result, the systems described herein can be easily integrated with conventional SIEMs and/or IDSs using the API.

[0137] The machine learning aspects provided herein, enable the automated IP address reputation system and the automated network domain system to produce higher fidelity maliciousness scores and alerts over time. Additionally, the particular features used to determine a maliciousness score provide for a more robust malicious activity analysis. Similarly, the processes of FIGS. 4A and 4B enable the automated IP address reputation

system and the automated network domain system to expand the scope of a request maliciousness analysis to include related IP addresses and network domains. By analyzing the related IP addresses and related network domains in addition to the original network identifier set, the likelihood of detecting malicious activity is increased over conventional systems.

VI. COMPUTER SYSTEM

[0138] Any of the computer systems mentioned herein may utilize any suitable number of subsystems. In some embodiments, a computer system includes a single computer apparatus, where the subsystems can be the components of the computer apparatus. In other 10 embodiments, a computer system can include multiple computer apparatuses, each being a subsystem, with internal components. A computer system can include desktop and laptop computers, tablets, mobile phones and other mobile devices.

[0139] The subsystems may be interconnected via a system bus. Additional subsystems such as a printer, a keyboard, one or more storage device(s), a monitor, which is coupled to a display adapter, and others may be utilized. Peripherals and input/output (I/O) devices, which 15 couple to I/O controller, can be connected to the computer system by any number of means known in the art such as input/output (I/O) port (e.g., USB, FireWire[®]). For example, I/O port or external interface (e.g. Ethernet, Wi-Fi, etc.) can be used to connect a computer system to a wide area network such as the Internet, a mouse input device, or a scanner. The 20 interconnection via system bus may allow the central processor to communicate with each subsystem and to control the execution of instructions from system memory or storage device(s) (e.g., a fixed disk, such as a hard drive, or optical disk), as well as the exchange of information between subsystems. The system memory and/or the storage device(s) may embody a computer readable medium. Another subsystem is a data collection device, such as 25 a camera, microphone, accelerometer, and the like. Any of the data mentioned herein can be output from one component to another component and can be output to the user.

[0140] A computer system can include a plurality of the same components or subsystems, e.g., connected together by external interface or by an internal interface. In some 30 embodiments, computer systems, subsystem, or apparatuses can communicate over a network. In such instances, one computer can be considered a client and another computer a

server, where each can be part of a same computer system. A client and a server can each include multiple systems, subsystems, or components.

[0141] It should be understood that any of the embodiments of the present invention can be implemented in the form of control logic using hardware (e.g. an application specific
5 integrated circuit or field programmable gate array) and/or using computer software with a generally programmable processor in a modular or integrated manner. As used herein, a processor includes a single-core processor, multi-core processor on a same integrated chip, or multiple processing units on a single circuit board or networked. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will know and appreciate other
10 ways and/or methods to implement embodiments of the present invention using hardware and a combination of hardware and software.

[0142] Any of the software components or functions described in this application may be implemented as software code to be executed by a processor using any suitable computer language such as, for example, Java, C, C++, C#, Objective-C, Swift, or scripting language
15 such as Perl or Python using, for example, conventional or object-oriented techniques. The software code may be stored as a series of instructions or commands on a computer readable medium for storage and/or transmission. A suitable non-transitory computer readable medium can include random access memory (RAM), a read only memory (ROM), a magnetic medium such as a hard-drive or a floppy disk, or an optical medium such as a compact disk
20 (CD) or DVD (digital versatile disk), flash memory, and the like. The computer readable medium may be any combination of such storage or transmission devices.

[0143] Such programs may also be encoded and transmitted using carrier signals adapted for transmission via wired, optical, and/or wireless networks conforming to a variety of protocols, including the Internet. As such, a computer readable medium according to an
25 embodiment of the present invention may be created using a data signal encoded with such programs. Computer readable media encoded with the program code may be packaged with a compatible device or provided separately from other devices (e.g., via Internet download). Any such computer readable medium may reside on or within a single computer product (e.g. a hard drive, a CD, or an entire computer system), and may be present on or within different
30 computer products within a system or network. A computer system may include a monitor, printer, or other suitable display for providing any of the results mentioned herein to a user.

[0144] Any of the methods described herein may be totally or partially performed with a computer system including one or more processors that can be configured to perform the steps. Thus, embodiments can be directed to computer systems configured to perform the steps of any of the methods described herein, potentially with different components
5 performing a respective steps or a respective group of steps. Although presented as numbered steps, steps of methods herein can be performed at a same time or in a different order. Additionally, portions of these steps may be used with portions of other steps from other methods. Also, all or portions of a step may be optional. Additionally, any of the steps of any of the methods can be performed with modules, units, circuits, or other means for performing
10 these steps.

[0145] The specific details of particular embodiments may be combined in any suitable manner without departing from the spirit and scope of embodiments of the invention. However, other embodiments of the invention may be directed to specific embodiments relating to each individual aspect, or specific combinations of these individual aspects.

15 [0146] The above description of example embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form described, and many modifications and variations are possible in light of the teaching above.

20 [0147] A recitation of "a", "an" or "the" is intended to mean "one or more" unless specifically indicated to the contrary. The use of "or" is intended to mean an "inclusive or," and not an "exclusive or" unless specifically indicated to the contrary. Reference to a "first" component does not necessarily require that a second component be provided. Moreover, reference to a "first" or a "second" component does not limit the referenced component to a particular location unless expressly stated.

25 [0148] All patents, patent applications, publications, and descriptions mentioned herein are incorporated by reference in their entirety for all purposes. None is admitted to be prior art.

WHAT IS CLAIMED IS:

1

1

1. A computer-implemented method, comprising performing by a cybersecurity system:

2

3

receiving a request to evaluate malicious activity with respect to a set of one or more network identifiers, wherein individual network identifiers of the set of one or more network identifiers comprise an internet protocol (IP) address or a network domain;

4

5

6

determining at least one of: a set of one or more related network domains or a set of one or more related IP addresses associated with the individual network identifiers of the set of network identifiers;

7

8

9

sending, to a plurality of third-party servers, requests to provide malicious activity information related to the set of one or more network identifiers, the set of one or more related network domains, and the set of one or more related IP addresses;

10

11

12

receiving, from the plurality of third-party servers, at least a portion of the malicious activity information requested;

13

14

extracting a first set of features from the received malicious activity information;

15

16

calculating a second set of features based on the received malicious activity information, wherein the second set of features includes at least one statistical value calculated from a common feature of the first set of features received from each of the plurality of third-party servers;

17

18

19

20

calculating one or more maliciousness scores for the set of one or more IP addresses utilizing a machine learning model, wherein the machine learning model utilizes the first set of features and the second set of features as input, and wherein the machine learning model is trained utilizing historical malicious activity information from one or more of the plurality of third-party servers; and

21

22

23

24

25

performing a remedial action based on the one or more maliciousness scores.

1

2

2. The computer-implemented method of claim 1, wherein the one or more network identifiers comprise at least one IP address, and wherein determining the set of one or more network domains and the set of one or more related IP addresses comprises performing by the cybersecurity system:

3

4

5 requesting, from a domain name service (DNS) provider computer, DNS
6 information, the DNS information identifying one or more network domains that individually
7 host the set of one or more IP addresses;

8 determining, from the DNS information, the set of one or more related IP
9 addresses based on the identified one or more network domains; and

10 determining, from the DNS information, one or more additional domains
11 related to the set of one or more related IP addresses.

1 3. The computer-implemented method of claim 1, wherein the one or
2 more network identifiers comprise at least one network domain, and wherein determining the
3 set of one or more network domains and the set of one or more related IP addresses
4 comprises performing by the cybersecurity system:

5 identifying at least one of: a co-occurrence network domain set, a related
6 network domain set, a shared-name network domain set, an email network domain set, a IP
7 address related network domain set, or a related IP address set.

1 4. The computer-implemented method of claim 1, further comprising
2 performing by the cybersecurity system:

3 maintaining a whitelist filter set;

4 comparing the set of one or more network identifiers, the set of one or more
5 related network domains, and the set of one or more related IP addresses to the white-list
6 filter set; and

7 removing a particular IP address from at least one of: the set of one or more
8 network identifiers, the set of one or more related network domains, or the set of one or more
9 related IP addresses based on the comparison, wherein removing the particular network
10 identifier from a set prevents further processing of the particular network identifier.

1 5. The computer-implemented method of claim 1, wherein the first set of
2 features comprises a time of receipt associated with the malicious activity information, and
3 wherein the machine learning model weighs the malicious activity information more heavily
4 than previously-received malicious activity information based on the time of receipt.

1 6. The computer-implemented method of claim 1, wherein the first set of
2 features comprises data obtained directly from the maliciousness activity information, and

3 wherein the second set of features comprises one or more statistical values corresponding to
4 at least one of the first set of features.

1 7. The computer-implemented method of claim 1, further comprising
2 performing by the cybersecurity system:
3 training the machine learning model by:
4 obtaining historical malicious activity information, the historical
5 malicious activity information being associated with predetermined maliciousness
6 scores;
7 extracting features from the historical malicious activity information;
8 calculating a plurality of statistical values based on the extracted
9 features;
10 generating a plurality of machine learning models utilizing the
11 extracted features and the plurality of statistical values;
12 for each of the plurality of machine learning models, calculating
13 maliciousness scores for the historical malicious activity information; and
14 selecting the machine learning model from the plurality of machine
15 learning models based on determining that the maliciousness scores generated by the
16 machine learning model most closely correspond to the predetermined maliciousness
17 scores.

1 8. The computer-implemented method of claim 1, further comprising
2 performing by the cybersecurity system:
3 determining a reputation score associated with a third-party server of the
4 plurality of servers, wherein the maliciousness score is calculated further based on the
5 reputation score.

1 9. The computer-implemented method of claim 1, wherein the one or
2 more remedial actions comprises at least one of: restricting network traffic associated with
3 the set of one or more IP addresses or providing a notification to a remote device to cause the
4 remote device to restrict network traffic associated with the set of one or more IP addresses.

1 10. A computer product, comprising:
2 one or more processors; and

3 a computer readable medium storing a plurality of instructions that, when
4 executed, cause the one or more processors to:

- 5 receive a request to evaluate malicious activity with respect to a set of
6 one or more network identifiers, wherein individual network identifiers of the set of
7 one or more network identifiers comprise an internet protocol (IP) address or a
8 network domain;
- 9 determine at least one of: a set of one or more related network domains
10 or a set of one or more related IP addresses associated with the individual network
11 identifiers of the set of network identifiers;
- 12 send, to a plurality of third-party servers, requests to provide malicious
13 activity information related to the set of one or more network identifiers, the set of
14 one or more related network domains, and the set of one or more related IP addresses;
- 15 receive, from the plurality of third-party servers, at least a portion of
16 the malicious activity information requested;
- 17 extracting a first set of features from the received malicious activity
18 information;
- 19 calculate a second set of features based on the received malicious
20 activity information, wherein the second set of features includes at least one statistical
21 value calculated from a common feature of the first set of features received from each
22 of the plurality of third-party servers;
- 23 calculate one or more maliciousness scores for the set of one or more
24 IP addresses utilizing a machine learning model, wherein the machine learning model
25 utilizes the first set of features and the second set of features as input, and wherein the
26 machine learning model is trained utilizing historical malicious activity information
27 from one or more of the plurality of third-party servers; and
- 28 perform a remedial action based on the one or more maliciousness
29 scores.

1 11. The computer product of claim 10, wherein the one or more network
2 identifiers comprise at least one IP address, and wherein determining the set of one or more
3 network domains and the set of one or more related IP addresses comprises further
4 instructions that, when executed, cause the one or more processors to:

5 request, from a domain name service (DNS) provider computer, DNS
6 information, the DNS information identifying one or more network domains that individually
7 host the set of one or more IP addresses;

8 determine, from the DNS information, the set of one or more related IP
9 addresses based on the identified one or more network domains; and

10 determine, from the DNS information, one or more additional domains related
11 to the set of one or more related IP addresses.

1 12. The computer product of claim 10, wherein the one or more network
2 identifiers comprise at least one network domain, and wherein the determining the set of one
3 or more network domains and the set of one or more related IP addresses comprises further
4 instructions that, when executed, cause the one or more processors to:

5 identify at least one of: a co-occurrence network domain set, a related network
6 domain set, a shared-name network domain set, an email network domain set, a IP address
7 related network domain set, or a related IP address set.

1 13. The computer product of claim 10, comprising further instructions that,
2 when executed, cause the one or more processors to:

3 maintain a whitelist filter set;

4 compare the set of one or more network identifiers, the set of one or more
5 related network domains, and the set of one or more related IP addresses to the white-list
6 filter set; and

7 remove a particular IP address from at least one of: the set of one or more
8 network identifiers, the set of one or more related network domains, or the set of one or more
9 related IP addresses based on the comparison, wherein removing the particular network
10 identifier from a set prevents further processing of the particular network identifier.

1 14. The computer product of claim 10, wherein the first set of features
2 comprises a time of receipt associated with the malicious activity information, and wherein
3 the machine learning model weighs the malicious activity information more heavily than
4 previously-received malicious activity information based on the time of receipt.

1 15. The computer product of claim 10, wherein the first set of features
2 comprises data obtained directly from the maliciousness activity information, and wherein the

3 second set of features comprises one or more statistical values corresponding to at least one
4 of the first set of features.

1 16. The computer product of claim 10, further comprising additional
2 instructions that, when executed, to cause the one or more processors to:
3 train the machine learning model by:
4 obtaining historical malicious activity information, the historical
5 malicious activity information being associated with predetermined maliciousness
6 scores;
7 extracting features from the historical malicious activity information;
8 calculating a plurality of statistical values based on the extracted
9 features;
10 generating a plurality of machine learning models utilizing the
11 extracted features and the plurality of statistical values;
12 for each of the plurality of machine learning models, calculating
13 maliciousness scores for the historical malicious activity information; and
14 selecting the machine learning model from the plurality of machine
15 learning models based on determining that the maliciousness scores generated by the
16 machine learning model most closely correspond to the predetermined maliciousness
17 scores.

1 17. The computer product of claim 10, further comprising additional
2 instructions that, when executed, to cause the one or more processors to:
3 determine a reputation score associated with a third-party server of the
4 plurality of servers, wherein the maliciousness score is calculated further based on the
5 reputation score.

1 18. The computer product of claim 10, wherein the one or more remedial
2 actions comprises at least one of: restricting network traffic associated with the set of one or
3 more IP addresses or providing a notification to a remote device to cause the remote device to
4 restrict network traffic associated with the set of one or more IP addresses.

1 19. A system comprising:
2 the computer product of claim 10; and

3 one or more processors for executing instructions stored on the computer
4 readable medium.

1 20. A system comprising means for performing any of the methods above.

1 21. A system configured to perform any of the above methods.

1 22. A system comprising modules that respectively perform the steps of
2 any of the above methods.

1

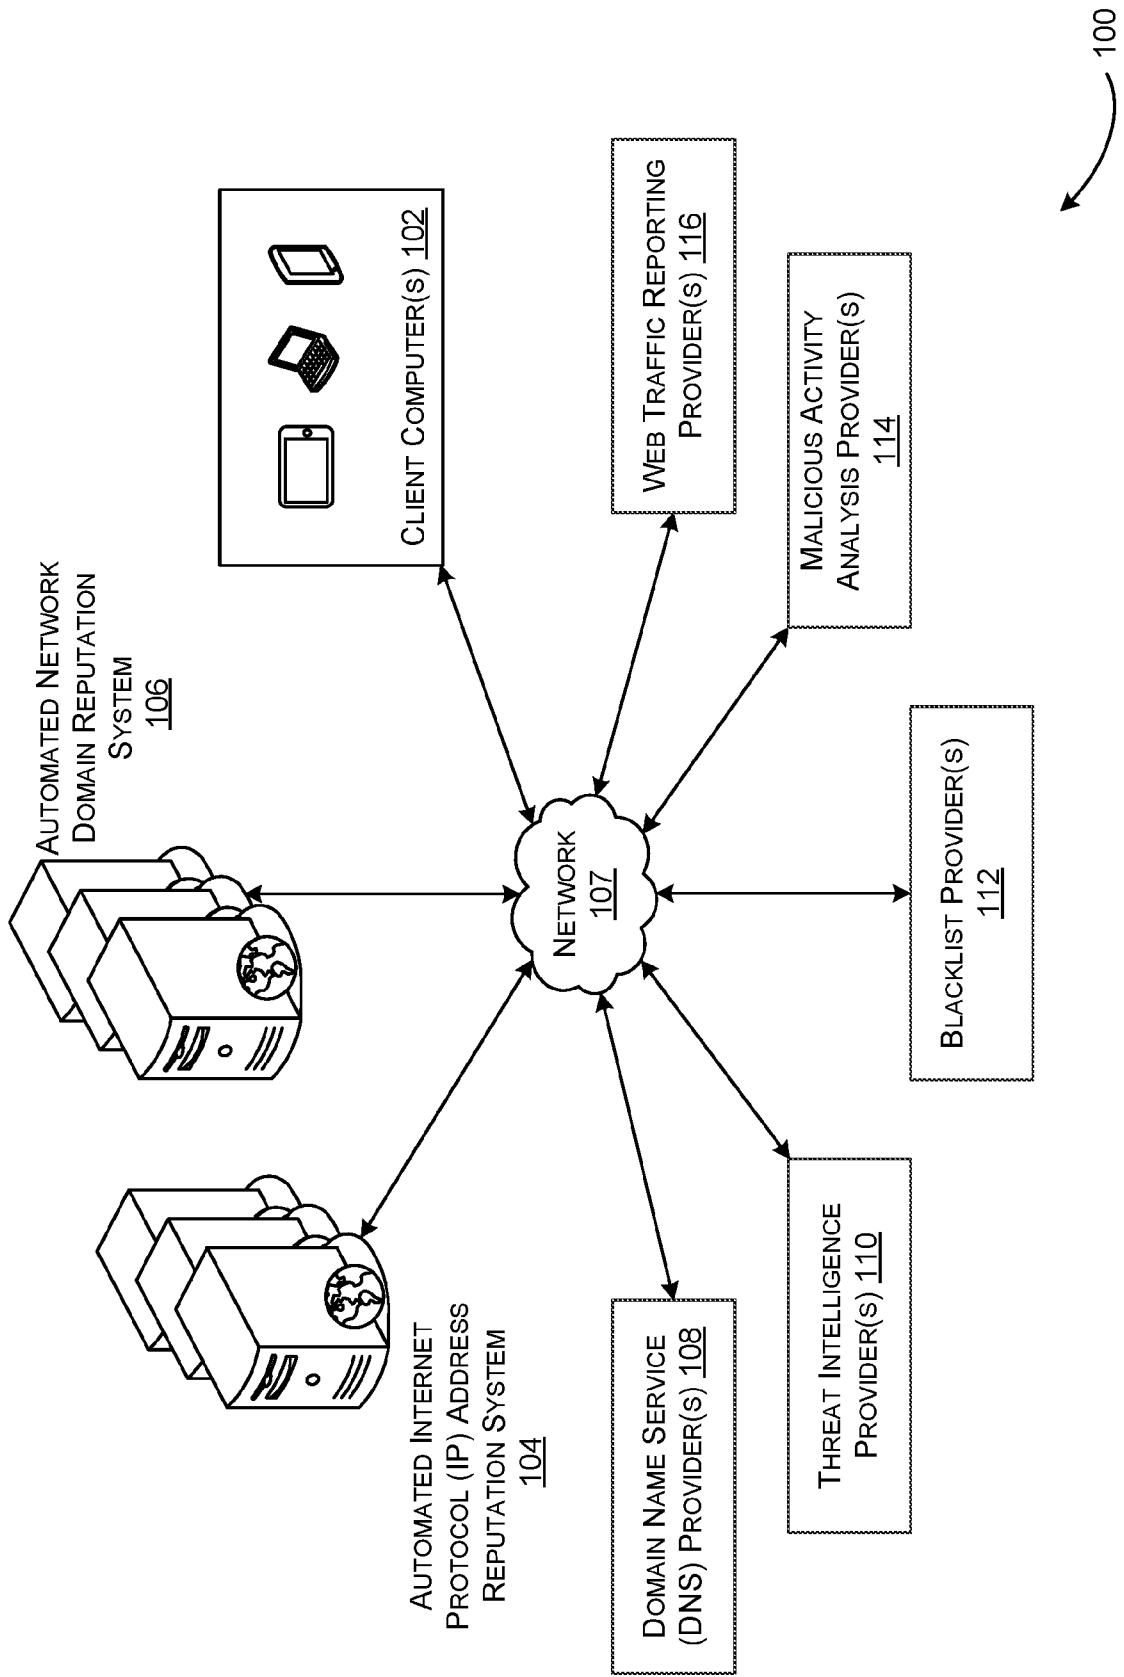


FIG. 1




FIG. 2

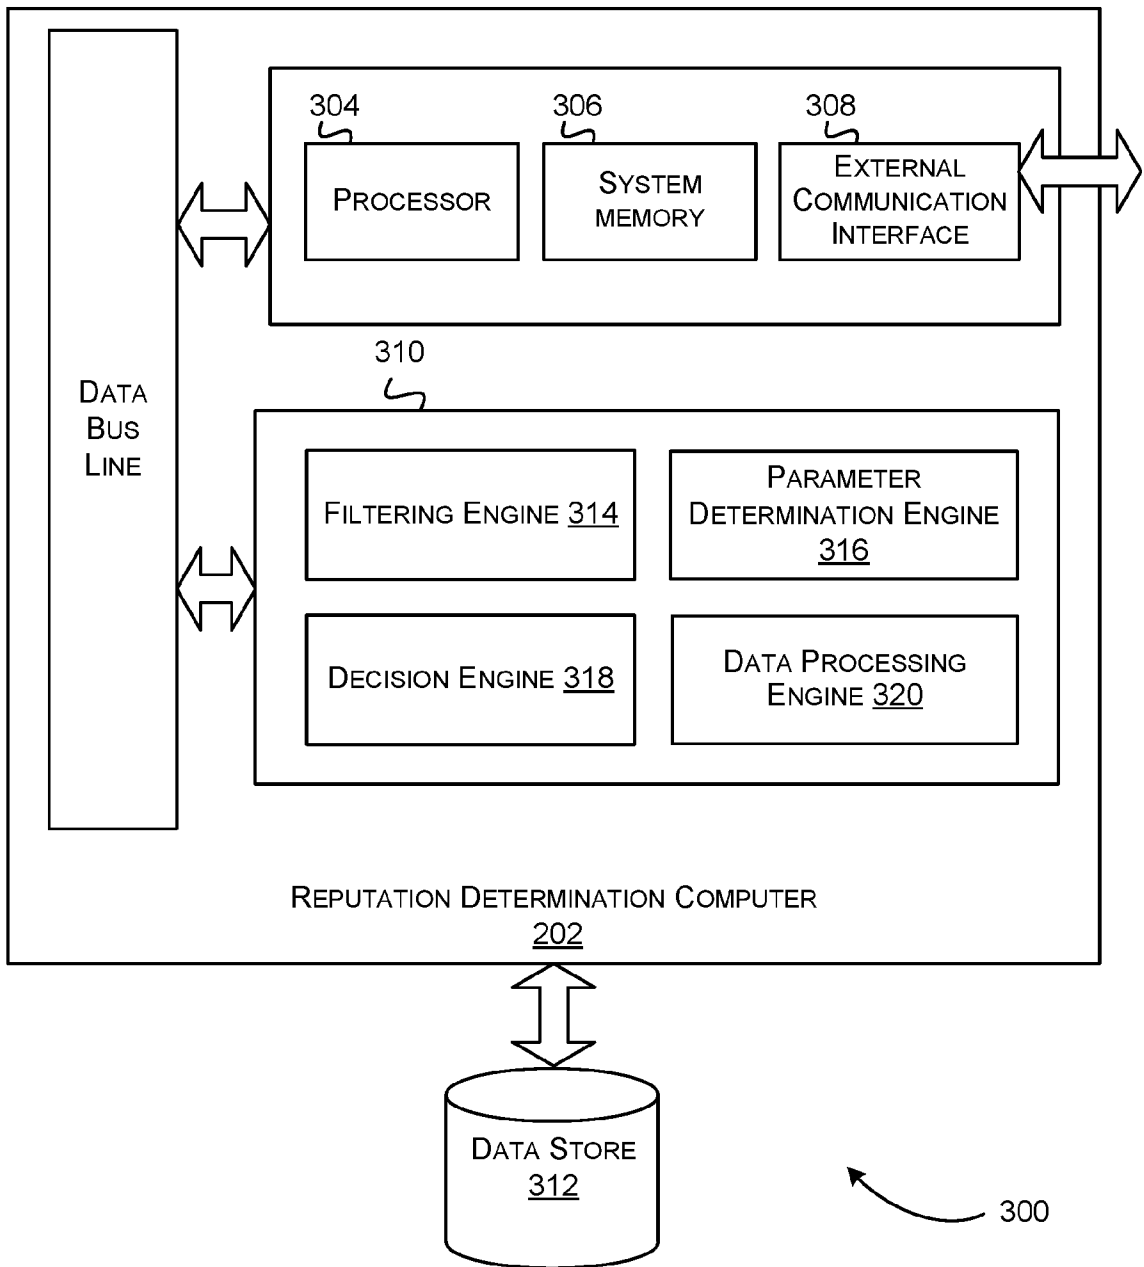


FIG. 3

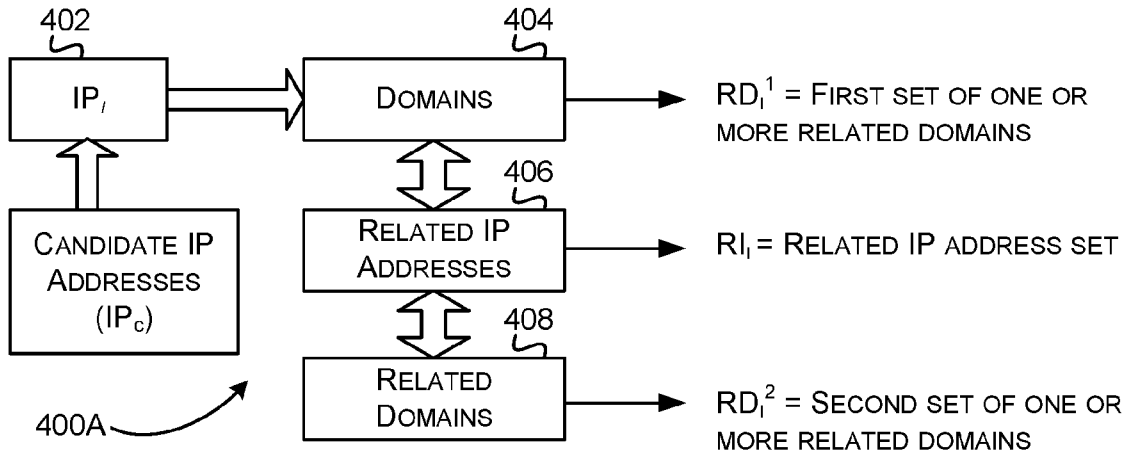


FIG. 4A

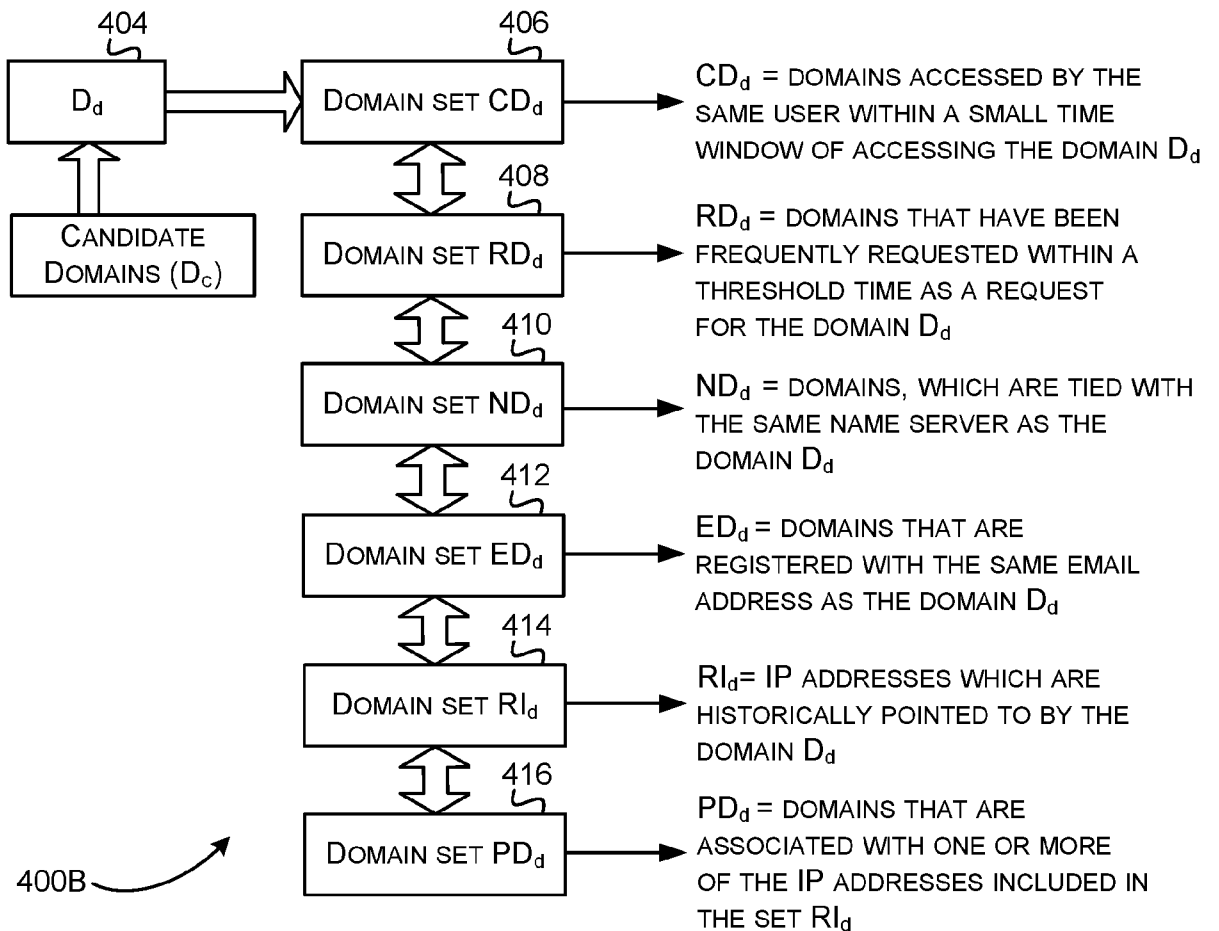
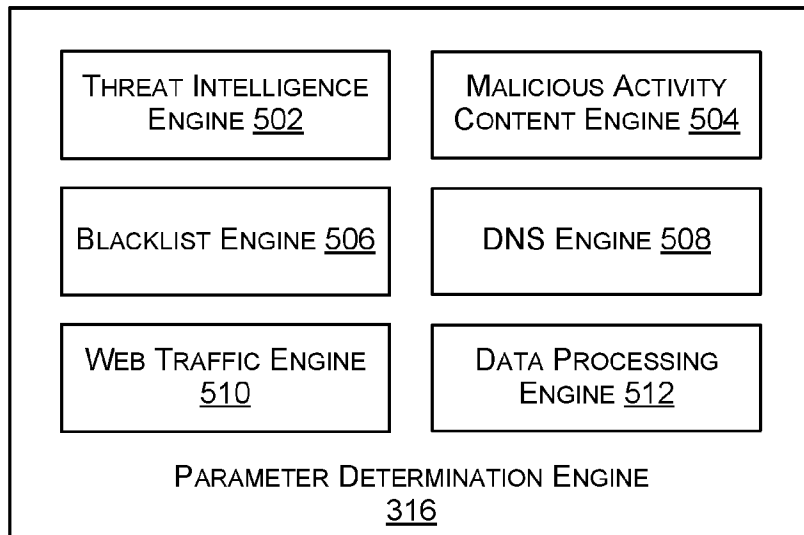
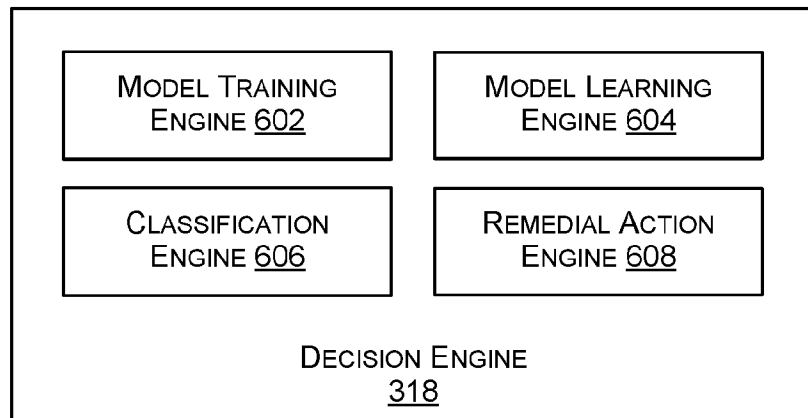


FIG. 4B



500

FIG. 5



600

FIG. 6




FIG. 7

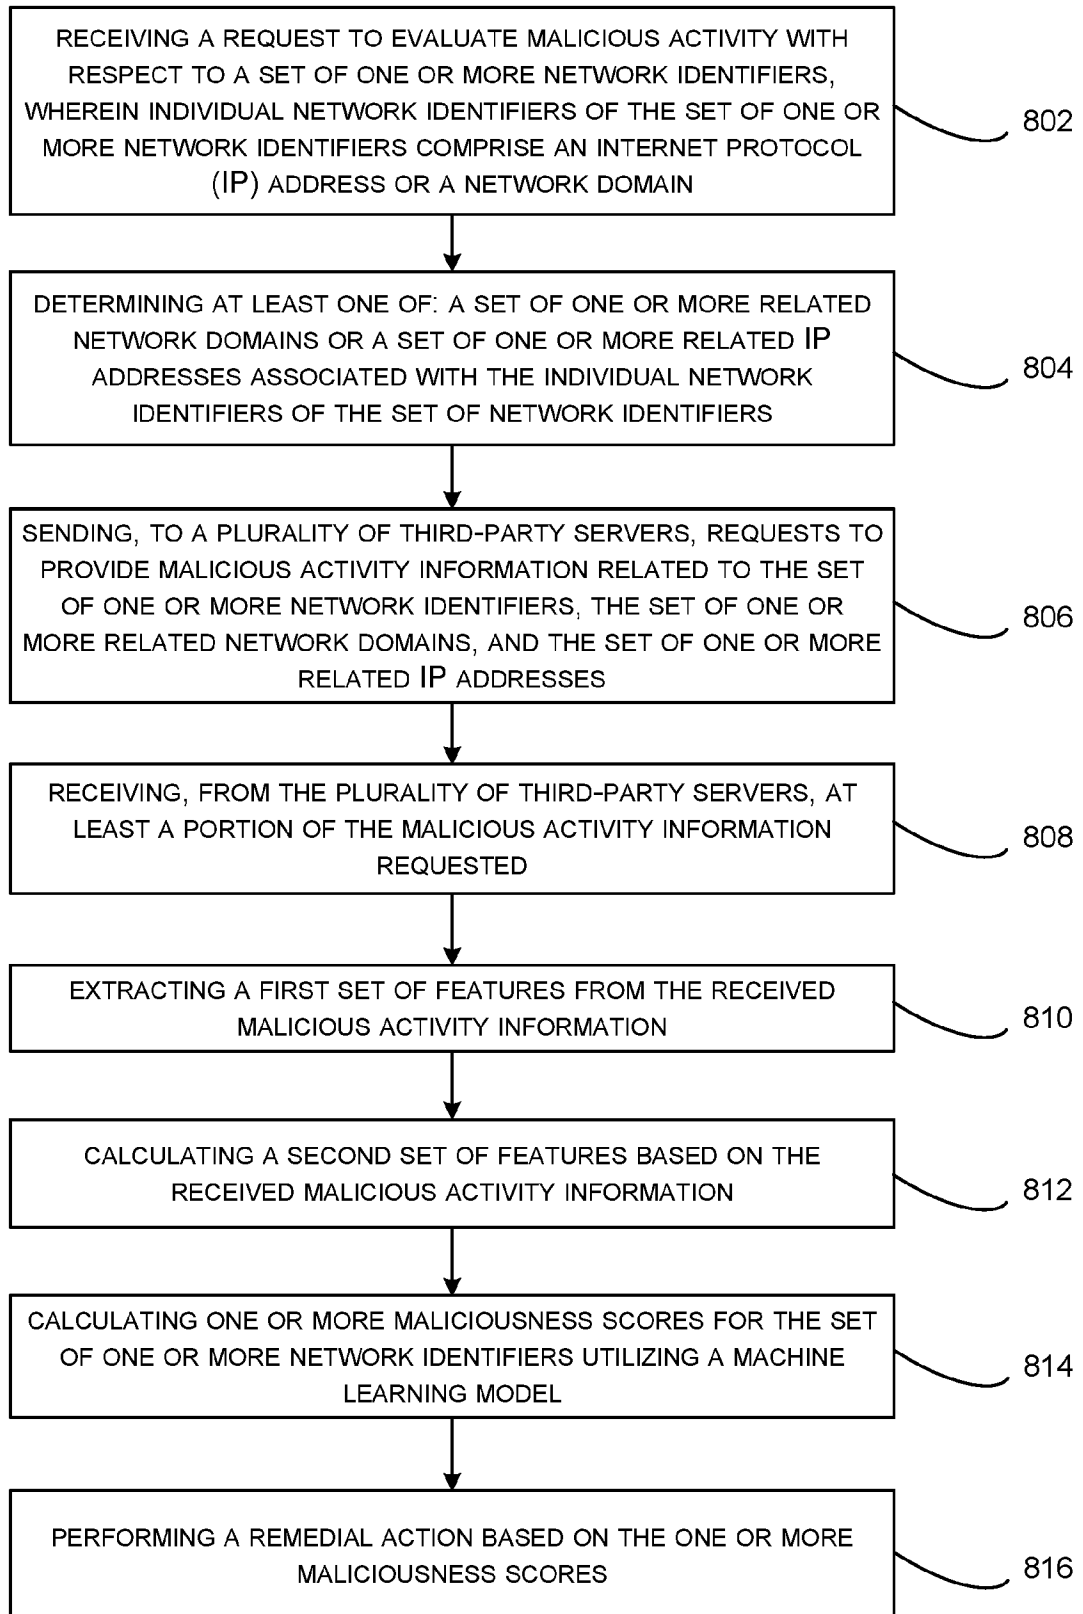


FIG. 8

A. CLASSIFICATION OF SUBJECT MATTER**H04L 29/06(2006.01)**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
H04L 29/06; G06F 15/16; G06F 15/00; H04L 9/28; H04L 29/12; H04L 12/22Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: third-party, server, malicious activity, learning, model, score, network, domain, identifier**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2016-0065597 A1 (NOMINUM, INC.) 03 March 2016 See paragraphs [0007]-[0009], [0022]-[0028], [0055]-[0059], [0070], [0088]; and figures 1-2.	1-22
A	US 2014-0096251 A1 (LEVEL 3 COMMUNICATIONS, LLC) 03 April 2014 See paragraphs [0012]-[0020]; and figure 1.	1-22
A	US 2014-0075558 A1 (DAMBALLA, INC.) 13 March 2014 See paragraphs [0011]-[0033]; and figures 1-2.	1-22
A	KR 10-0877911 B1 (INDUSTRY FOUNDATION OF CHONNAM NATIONAL UNIVERSITY) 12 January 2009 See paragraphs [0016]-[0025]; and figure 1.	1-22
A	KR 10-0615080 B1 (INFOSEC TECHNOLOGIES CO., LTD et al.) 25 August 2006 See paragraphs [0036]-[0038], [0050]; and figures 1, 3.	1-22

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

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

08 December 2017 (08.12.2017)

Date of mailing of the international search report

08 December 2017 (08.12.2017)

Name and mailing address of the ISA/KR


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

Facsimile No. +82-42-481-8578

Authorized officer

LEE, Chang Ho

Telephone No. +82-42-481-8288



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/021961

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016-0065597 A1	03/03/2016	US 2013-0014253 A1 US 2016-0065534 A1 US 2016-0065535 A1 US 2016-0065611 A1 US 9185127 B2	10/01/2013 03/03/2016 03/03/2016 03/03/2016 10/11/2015
US 2014-0096251 A1	03/04/2014	CA 2886058 A1 EP 2901612 A2 EP 2901612 A4 HK 1213107 A1 US 2015-0215334 A1 WO 2014-052756 A2 WO 2014-052756 A3 WO 2016-164403 A1	03/04/2014 05/08/2015 15/06/2016 24/06/2016 30/07/2015 03/04/2014 30/05/2014 13/10/2016
US 2014-0075558 A1	13/03/2014	US 9166994 B2	20/10/2015
KR 10-0877911 B1	12/01/2009	None	
KR 10-0615080 B1	25/08/2006	None	