



- (51) International Patent Classification:  
G06F 11/00 (2006.01)
- (21) International Application Number:  
PCT/US2015/061191
- (22) International Filing Date:  
17 November 2015 (17.11.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
62/081,510 18 November 2014 (18.11.2014) US
- (71) Applicant: VECTRA NETWORKS, INC. [US/US]; 550 S Winchester Blvd., #200, San Jose, California 95128 (US).
- (72) Inventors: BEAUCHESNE, Nicolas; 135 3rd. St., Miami Beach, Florida 33139 (US). PEGNA, David Lopes; 2388 Brittan Ave., San Carlos, California 94070 (US). LYNN, Karl; 739 Bainbridge Loop, Winter Garden, Florida 34787 (US).
- (74) Agent: MEI, Peter C.; Vista IP Law Group, LLP, 2160 Lundy Ave., Ste. 230, San Jose, California 95131 (US).

- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published: — with international search report (Art. 21(3))

(54) Title: METHOD AND SYSTEM FOR DETECTING THREATS USING METADATA VECTORS

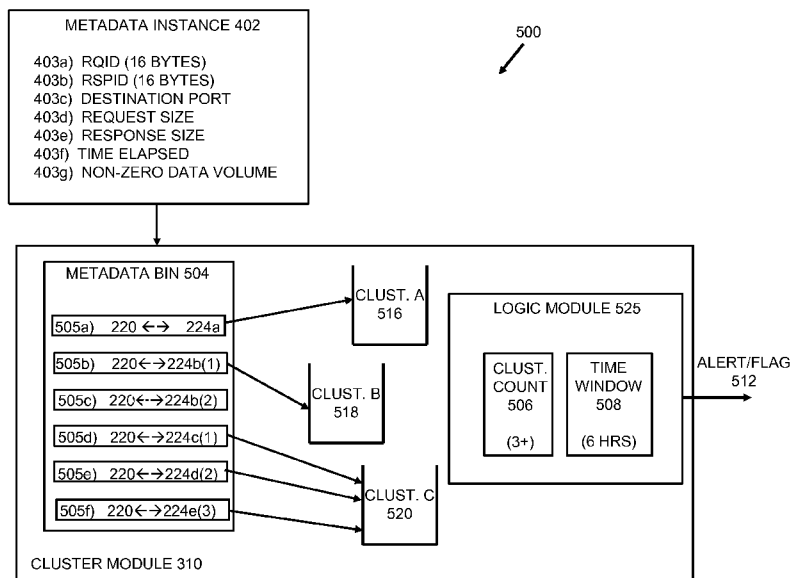


FIG. 5A

(57) Abstract: An approach for detecting network attacks using metadata vectors may initially involve receiving network communications or packets, extracting metadata items from the packets. The metadata items describe the communications without requiring deep content inspection of the data payload or contents. The communications may be clustered into groups using the metadata items. If a cluster exceeds a threshold, an alarm may be generated.

WO 2016/081520 A1

## METHOD AND SYSTEM FOR DETECTING THREATS USING METADATA VECTORS

### BACKGROUND

**[0001]** It is often an important goal for network security systems to be able to passively identify suspicious communications in a way that does not modify or disrupt the network. Typically, conventional passive network security systems, such as intrusion detection systems, rely on deep content inspection of packets to identify suspicious communications. Such approaches require deep knowledge, intricate parsing, and usually a library of signatures or heuristics for each type of suspicious activity that is to be identified. Often the library of signatures or heuristics can sprawl to 30,000 or more, which can make the use of such libraries unwieldy and/or computationally expensive. Furthermore, in some cases, deep content inspection may not be available due to restricted network permissions and/or the sensitivity of the data being transmitted. These requirements create an enormous burden on vendors seeking to create products that passively detect threats over a range of networking protocols.

**[0002]** As is evident, there is a need for an approach to passively identify suspicious communications or network threats without deep content inspection and/or sprawling signature libraries.

## SUMMARY

**[0003]** In some embodiments, a vector engine may be implemented to detect network threats using metadata vectors. In some embodiments, the vector engine may first receive network communications, extract metadata from the network communications, group the metadata into clusters, and if the size of one or more metadata clusters exceeds a given threshold within a predefined period of time, generate an alarm. The network comprises a plurality of client computers communicating with an outside network (i.e. the internet) as well as internally with each other. When a computer is infected by a malicious payload, such as malicious automated instructions or script, the computer may act as a bot and may send copies of the same or a different malicious payload to a number of other computers in the network, to infect them as well. The malicious automated script may include logic that instructs the first infected computer to send the same or a different automated script to multiple other computers. The script may further be addressed to the same destination port on each computer. Further, when the multiple other computers receive and process (e.g., execute/interpret) the script, they may respond to the first computer using the same type of response. Thus, in this case, for example, the characteristics of such an attack would be a first computer in a network sending substantially the same data (e.g., the malicious payload) of a set size, addressed to the same destination port on multiple other computers, each of which may respond to the first computer using substantially the same response message. Further, the time between each request and its matching response will typically be short (e.g., less than five seconds).

**[0004]** In some embodiments, the vector engine can collect the above network communications by tapping a network switch and making a copy of the network communications for analysis, without disrupting or slowing down the network. Next, the vector engine may extract metadata from the network communications that can be used to characterize the network communications without inspecting the contents. The vector engine may group the metadata into matching clusters that allow the vector engine to track how many times a given network communication has occurred. If a certain network communication, such as one corresponding to a malicious attack, occurs too many times, the cluster may grow in size and exceed a given threshold. If the threshold is exceeded or surpassed, the vector engine may generate an alarm or may flag the data for further monitoring.

**[0005]** As an example, if the threshold for a given cluster's size is ten, and a first computer is infected by a malicious payload and tries to spread copies of the same or different malicious payload to eleven other computers, then eleven instances of metadata will be generated by the vector engine. Further, because the eleven instances of metadata were all generated from the same source, the eleven instances of metadata will all have similar data items and can be grouped into the same cluster. As the threshold is ten, the size of the cluster containing the eleven instances exceeds the threshold and the vector engine may generate an alert to signal that the first computer is potentially attacking other computers on the network.

**[0006]** Described below are general non-limiting example definitions that may be used to understand or enhance understanding of how certain embodiments may be implemented. In some embodiments, a network may be collection of computing entities

that communicate with one another. In some embodiments, an internal network is a network of computers that are isolated from other (e.g., external) networks by networking components and/or devices as are known in the art, such as a firewall. In some embodiments, a client computer is a computer or node within a network. In some embodiments, a computer may act as a client and/or host depending on the type of connection with another computer. As used here, clients are generally uninfected but can become infected by malicious payloads, in which case they may act as bots for the outside or attacking entity.

**[0007]** In some embodiments, an attacking computer is a computer or entity (e.g., malicious user) that initiates a malicious attack on a computer, client computer, or network. In some instances, an attacking computer may use a malicious payload to attack other computers in the network, whereby the malicious payload may infect a first computer inside the network and direct the first computer to distribute copies of the same or another malicious payload to other computers, who then may further replicate and distribute the malicious payload. A bot may be a client computer that is under the effective control of the attacking computer. For example, the attacking computer can use the bot to achieve further malicious goals, and/or perform tasks. A malicious payload may be a collection of instructions, which may be in the form of code (e.g., JavaScript) that is used to perform malicious tasks. The payload may in some form be compilable code, scripted code, or other types of instructions as is known to those of ordinary skill in the art. In some instances, malicious payloads may correspond to large toolkits that are 500 megabytes or more in size. In other instances, a malicious payload may be a small automated script, such as a script of 1 megabyte (or less) in size.

**[0008]** As according to some embodiments, vector metadata is metadata that is used to characterize and/or categorize network communications without inspecting the contents of the network communications (e.g., deep content inspection). For example, in some embodiments the metadata may comprise data from packet headers, and in some embodiments the metadata may be information that describes the overall session between two hosts. A metadata instance is a collection of one or more metadata items (e.g., packet information, session information) specific to a given network communication. A metadata item as according to some embodiments, is a parameter or item, such as a response size, that describes a certain feature of the network communication. For example, in some embodiments the metadata may comprise data from packet headers, and in some embodiments the metadata may be information that describes the overall session between two hosts. In some embodiments, a metadata item is used to characterize a network communication without inspecting the contents of the network communication. In some embodiments, the metadata item does not contain any reference or indicator of what is inside (e.g., the contents) of a network communication. A metadata cluster may be a grouping of metadata instances that may generally share similar metadata items. In some embodiments, a metadata cluster can be used to track how many network communications of a certain type have occurred.

**[0009]** In some embodiments, a first communication is a communication initiated by one computer directed to another computer. In some embodiments, a first communication may take the form of a request or a “upstream” communication emanating from a client computer. In some embodiments, a first communication may correspond to a transmission that includes a file in a direction from a first computer to a

second computer. A second communication may be a communication initiated by the second computer directed to the first computer. In some embodiments, a second communication may take the form of a response or a “downstream” communication received from the host or responding computer. In some embodiments, a second communication may correspond to a response message initiated in response to receiving and/or processing the first communication. A network communication is a communication or session between two or more computers in a network using various networking protocols, such as TCP/IP. In some embodiments, a network communication comprises the first communication and a second communication as a pair, session, or session portion. In some embodiments, a network communication comprises an initiating message (e.g., request) in one direction that begins the communication, and a response message in the opposite direction, which may be a response to the initiating message. In some embodiments, a network communication may be a unilateral communication attempt by one computer where no response occurs (e.g., a failed request).

**[0010]** In some embodiments, the approaches here may be implemented in an example networked environment in which an outside computer is attacking an internal network. The internal network comprises a plurality of client computers. The attacking computer may coax the client computer to retrieve a malicious payload, such as an automated script, by sending a malicious email with an embedded web link. When the client computer processes the malicious payload, it becomes a "bot". Under direction of the malicious payload, the client computer or bot may then attempt to send copies of the same malicious payload or a different malicious payload to a plurality of client

computers. When the plurality of client computers receive and process the replicated malicious payloads, they may respond to the bot computer with a response message (not depicted). The plurality client computers, may also send out more replicas of the malicious payload to potentially more client computers inside or outside the internal network. In this way, an outside attacking computer may potentially use a small automated script to attack and/or establish control over many clients in an internal network. Though in this example only client five client computers are sent the payload, in some embodiments the payload may be sent to only two computers, in other embodiments, the payload may be sent to fifteen clients to be infected, in a network comprising clients for example.

**[0011]** The below discloses an example as according to some embodiments for how a vector engine may be implemented to detect the network threats. The vector engine receives network communications for a given network. For example, a networking environment comprises an internal network, which may be isolated by networking components and devices as are known in the art, such as a firewall (not depicted). A plurality of client computers may be networked together to communicate through use of a network switch or other networking components. In some embodiments, the vector engine may receive network communications for the internal network through the network switch. The vector engine can also be implemented as software that resides upon networking equipment, nodes, or entities within a network. Finally, the vector engine can be implemented in a virtual machine and be deployed on hardware supplied by the company that owns the network.



**[0012]** An outside computer or entity is illustrated as an attacker. The attacker may interface with the internal network through the switch to send a malicious payload (e.g., a payload linked through or attached to an email) to the client computer. When the client computer processes the malicious payload, it becomes a "bot". Alternately, the client computer could become a "bot" under many other circumstances, including, but not limited to, getting infected when browsing a compromised web site. Under the direction of the malicious payload (and in some cases under the direction of the attacker), the bot sends a plurality of requests (represented by dashed lines) to the respective client computers. The requests in this example are copies of the malicious payload or of an entirely new malicious payload, and as such all requests may have the same or similar sizes (e.g., ~1 MB). Further, the requests may also be addressed to the same destination port on each client machine, respectively (e.g., request is addressed to example TCP port "135" on client A, while request is addressed to the same example TCP port "135" but on client B.).

**[0013]** When the plurality of client machines, receive and process the requests, they may reply to the bot with a response. Since the requests consisted of replicas of the same malicious payload, when each client machine processes the request, the client machines may send similar responses to the bot. Thus, the responses may all be approximately the same size and share similar characteristics with one another. As explained, in some cases, the malicious payload consists of an automated script that when processed causes actions to occur automatically or near instantaneously. And in some case, the malicious payloads, sent by the bot, is the same as the malicious payload it received. Given the nature of malicious payloads, the high speed of

networks, and the speed of computing devices the time elapsed between the requests and the responses may be small (e.g., as little as five seconds). Collectively, the communications (e.g., the requests and responses) comprise the network communications that the vector engine can receive and analyze. Next the vector engine may extract vector metadata from the network communications. In one example, vector metadata may be metadata that can be used to describe the network communications so that the vector engine can group or cluster the metadata and/or network communications for further analysis. For example, the metadata may comprise packet header information, or information about the session (e.g., time between request/responses) that is tracked as a vector comprising a plurality of features. After the metadata is extracted, the vector engine may group the metadata into clusters using a matching process. The metadata clusters may then be used to track how many instances of a given network communication has occurred. For example, the bot has sent out five requests, and received five responses thus in one cluster there may be five instances of metadata that correspond to the five actual network communication exchanges. The vector engine can set threshold size limits on the clusters to monitor them. The underlying rationale being that if a given cluster grows too large, the metadata instances in the cluster correspond to a rapidly growing network attack using malicious payloads, such as an automated script. For example, if the cluster (not depicted) that is tracking the five network communications (e.g., the five instances of requests and responses) has a threshold size limit of three, the vector engine may generate an alarm, to signal that the internal network and specifically the bot may be

infected and attacking other clients when the five instances of metadata are grouped in one cluster.

**[0014]** In some embodiments, the engine comprises one or more of the following components. A passive interface module that is designed to passively receive network communications, such as from a network switch. The passive interface module may then send the network communications to an extractor module to extract one or more instances of metadata that correspond to each network communication.

**[0015]** Once the metadata instances are extracted from the network communications, a cluster module may group the metadata instances into clusters. As explained, each cluster may have a threshold size limit such that if the amount of metadata instances in the cluster exceeds a threshold an alarm may be generated and output at.

**[0016]** The following are examples of a metadata instance for a given network communication. In the example, a bot has sent a request containing a malicious payload to a client computer. When the client computer processes the malicious payload it may send a response back to the bot. Collectively, the request and the response may comprise a single instance of a network communication. For example, network communication may correspond to a complete session or portion of a session. The extractor module may receive the network communication and extract one or more metadata items that describe the network communication. The one or more metadata items correspond to a metadata instance that describes the network communication without the need for inspecting the contents of the request or the response. The contents of the network communication may comprise the “data” payload for a given

network packet. The metadata items illustrated serve merely as an example of a possible set of metadata that may be extracted from a network communication, though one of ordinary skill in the art appreciates that other metadata items can be implemented as necessary.

**[0017]** In some embodiments, a communication identifier such as a request identifier (RQID) may correspond to the starting bytes of a packet. For example, request identifier may comprise the first sixteen bytes of a request of communication. Generally, the first bytes of a request contain header information or other data that is common across all requests. In this way, the first sixteen bytes can be used to identify similar requests without inspecting the contents or data payload of the communications and/or packets.

**[0018]** Similarly, the response identifier (RSPID), may be the first sixteen bytes of a response, such as communication. Because the response is generally a message created in response to the malicious payload being processed (e.g., executed), responses then generally contain the same header information. Just as with RQID, RSPID can likewise be used to identify similar responses without inspecting the contents of the response.

**[0019]** The destination port metadata data item corresponds to the protocol (e.g., TCP or UDP) and port on the receiving machine to which a request is directed. As explained above, malicious payloads in the form of automated scripts may be restricted to small payload sizes to remain inconspicuous in the network. As such, malicious payloads often direct the bot to send the replicated malicious payloads to the same destination port (as opposed to specifying multiple addresses/ports which may take up more space). Because the replicated malicious payloads (e.g., the requests) are all sent

to the same destination port on each respective client machine, the destination port metadata data item may be used to match similar network communications that are initiated by the same malicious payload.

**[0020]** The request size metadata data item describes the size of a request, such as a request. Because the requests sent to clients are generally the same, they likely share the same request size. If the requests sent to clients are also copies of the original malicious payload, the sizes of the requests sent to clients will be approximately the same as the size of payload. Thus, the request size metadata item can be used to detect matching requests. Similarly, the response size metadata data item describes the size of a response, such as response. Because the responses are typically all generated in response to the same requests, the responses may all have approximately the same sizes as well. Thus, the response size can be used to detect matching responses. The time elapsed metadata data item describes the length of time between when a request was sent out and when a response was sent back. Because the targeted hosts may be generally quick to process automated scripts and other types of network attacks, the time between request and response may be very short, e.g., less than five seconds. Thus, it may be the case that the duration of time between when each request was sent and when each corresponding response was sent is roughly the same. Thus, the time elapsed data item can be used to describe and match similar network communications.

**[0021]** The non-zero data volume metadata data item is a metadata item that ensures that all analyzed network communications are not empty. For example, if the request contains actual malicious payload data, then the non-zero data volume

metadata data item will reflect that the request is not empty. However, if was sent in error and is empty, then the non-zero metadata item will reflect that the request is empty. To save system resources, some embodiments of the invention may use the non-zero data volume metadata item to filter out all empty network communications.

**[0022]** Collectively, the metadata items may comprise a single metadata instance particular to a given network communication, such as network communication. Once the metadata items are extracted to compose one or more metadata instances, the cluster module can group the metadata instances into clusters using a cluster module.

**[0023]** Below is an example of how aspects of the cluster module may be implemented. There, a metadata instance is received from the extractor module. The cluster module may buffer one or more metadata instances in a metadata bin. As illustrated for this example, metadata bin contains six metadata instances. The cluster module may group the metadata instances into metadata clusters (e.g., Cluster-A, Cluster-B, Cluster-C) using a matching process, explained in greater detail below. However, as an example, metadata instances are all grouped into the same cluster, Cluster-C because some or all of their metadata items have been determined to approximately match.

**[0024]** As explained, each cluster may have a threshold size limit such that if the threshold is exceeded an alarm is generated. In some embodiments, the cluster module may implement a logic module to keep track of cluster counts and a time window to expire old metadata instances. The cluster count module can hold a data value that corresponds to the threshold size limit and further monitor the cluster count for each module. For example, if the cluster count module sets the threshold to be "2", then

Cluster-C, which currently holds three metadata instances would exceed the threshold, and the cluster count module and/or the logic module may generate an alarm data for output at.

**[0025]** In some embodiments, the time window module may operate as a sliding window of a discrete duration. For example, if the sliding window time duration is 6 hours then metadata instances in the clusters that are older than 6 hours are expired and removed from the cluster. The expiring and removing of old metadata instances using a sliding window can protect against false positives caused by potential build-up of unrelated metadata instances over very long periods of time. Further, though specific examples are used here for the threshold size limit (e.g., 3) and the time duration (e.g., six hours), one of ordinary skill in the art appreciates that such parameters can be modified or customized per implementation, as required.

**[0026]** Thus the clusters may be required to satisfy two conditions before an alert may be triggered. First, the cluster must exceed the threshold size limit (e.g., more than three metadata instances), and second, the threshold size must be exceeded in the given time window (e.g., more than three metadata instances within the last six hours). In some embodiments, the second condition is implicit in the checking of the first condition. For example, by constantly removing metadata instances older than six hours, the second condition is automatically enforced. In some embodiments, the alert data generated by the cluster module may generate an actual alarm automatically. That is, if the threshold size limit is exceeded within the time duration, an alarm is generated without intervention. In some embodiments, if the threshold size limit is exceeded within the time duration, the corresponding network activity is flagged as potentially malicious

activity and no alarm is directly generated. In this way, the potentially malicious activity may be combined with other unrelated anomalies to generate an automatic alarm or network administrators or security administrators can further monitor the network activity and manually generate an alarm as required.

**[0027]** Further, in some embodiments, a training window or period may be implemented, where the metadata instances are grouped into clusters to generate a baseline state of the network. The training period may be set to a time duration such as 24 hours. For example, the cluster module may group metadata instances into clusters for the first 24 hours to ascertain the appropriate value to set for each cluster's threshold size limit and may decide that special limits are required for some clients in the network. The training window could be used to identify periodic network communications that satisfy all the criteria to generate an alert, but that, as they happen relatively regularly in the network, could be associated to benign communications intentionally initiated from an uninfected client by a given user or process.

**[0028]** Below is an example for clustering metadata. First, the cluster module receives a group of metadata instances for analysis and grouping. Then, the cluster module analyzes the metadata instances in a strict match stage. In the strict match stage, in order for one or more metadata instances to be grouped together, certain metadata items in each metadata instance must exactly match one another, or must exactly match metadata items of instances already in the cluster. For example, if the request identifier (RQID) for the request exactly matches the request identifier (RQID) for a different request, then the corresponding metadata instances may be grouped together in the same cluster (assuming the other requirements are met). In some



embodiments, strict match items for Cluster-A, as an example, may include the following metadata items: RQID, RSPID, and Destination Port, Thus, to be grouped in Cluster-A at the strict match stage, a metadata instance must have the above metadata items exactly match the metadata instances already in Cluster-A or exactly matched to another metadata instance to start/form a cluster. If a metadata instance under consideration has metadata items that do not strictly match those in Cluster-A, for example, then the metadata instance is not grouped with Cluster-A at. Though continuing, if the metadata instance under consideration does in fact have metadata items that exactly match each other or those in Cluster-A, then the metadata instance may be analyzed in the approximate match stage.

**[0029]** In another stage, the approximate match stage, a metadata instance is analyzed to determine whether certain metadata items in the metadata instance approximately match each other or those in the cluster. The metadata items in the approximate match stage may include the request size and the response size (e.g., a communication size data item). Further, a range or percentage threshold may be specified to determine whether a given metadata item approximately matches the metadata items in the cluster. For example, the range is 5%. Thus, if a request size is within 5% of the average of those request sizes already in the cluster, the corresponding metadata instance may be grouped in the cluster at (assuming other conditions are met).

**[0030]** The following is an example flow for an approach for how the time window module may be implemented. After a start, two operations may occur approximately simultaneously. First, using a matching process, such as the one explained above,

matching metadata instances can be added to a given cluster. Second, as metadata instances are added to the cluster, the time in which they were added is recorded, and when a metadata item is more than six hours old it may be removed from the metadata cluster. In this way, the time window module implements a sliding window that is six hours wide to incrementally remove old metadata instances, as new metadata instances are added. In some embodiments, the metadata bin may be allowed to buffer for six hours at a time. At each six-hour increment, the metadata instances buffered in the metadata bin are then grouped into clusters. Once the metadata bin is empty, new metadata instances may be buffered in the bin for another six hours. In this way, the time window module may also implement a buffered or turn-based grouping process.

**[0031]** Below is a flow for an approach for determining whether a threshold size limit has been exceeded. First, a new matching metadata instance is added to a cluster. Next, after the new metadata instance is added to the cluster, a determination is made on whether the cluster contains more than the threshold size limit. For example, if the threshold size limit is 3, then the cluster count module determines whether there are more than 3 metadata instances in the cluster. If the cluster does not have more than three metadata instances, the threshold size limit is not exceeded and no alarm data is generated. On the other hand, if after adding the new matching metadata instance, the cluster contains more than three metadata instances, the threshold size limit is exceeded and an alarm may be generated.

**[0032]** Further details of aspects, objects, and advantages of some embodiments are described below in the detailed description, drawings, and claims. Both the foregoing

general description and the following detailed description are exemplary and explanatory, and are not intended to be limiting as to the scope of the embodiments.

### BRIEF DESCRIPTION OF THE DRAWINGS

**[0033]** The drawings illustrate the design and utility of embodiments of the present invention, in which similar elements are referred to by common reference numerals. In order to better appreciate the advantages and objects of embodiments of the invention, reference should be made to the accompanying drawings. However, the drawings depict only certain embodiments of the invention, and should not be taken as limiting the scope of the invention.

**[0034]** FIG. 1 illustrates an example network in which a vector engine may be implemented.

**[0035]** FIG. 2A shows an approach for detecting network threats using metadata vectors.

**[0036]** FIG. 2B illustrates an example network in which a vector engine has been implemented, as according to some embodiments.

**[0037]** FIG. 2C illustrates an example internal network in which a vector engine has been implemented, as according to some embodiments.

**[0038]** FIG. 3 illustrates an example block diagram showing internal features of a vector engine.

**[0039]** FIG. 4 shows example details of how metadata may be extracted from network communications.

**[0040]** FIG. 5A shows an examples of clustering metadata.

**[0041]** FIG. 5B shows an example flowchart for an approach for clustering metadata instances into similar clusters.

**[0042]** FIG. 5C shows an example flowchart on how to implement a sliding window in the vector engine.

**[0043]** FIG. 5D shows an example flowchart on how to detect whether the number of metadata instances in a cluster exceeds a threshold.

**[0044]** FIG. 6 illustrates example system architecture.

### DETAILED DESCRIPTION

**[0045]** Various embodiments of the invention are directed to a method, system, and computer program product for detecting threats using replicative metadata vectors. Other objects, features, and advantages of the invention are described in the detailed description, figures, and claims.

**[0046]** Various embodiments of the methods, systems, and articles of manufacture will now be described in detail with reference to the drawings, which are provided as illustrative examples of the invention so as to enable those skilled in the art to practice the invention. Notably, the figures and the examples below are not meant to limit the scope of the present invention. Where certain elements of the present invention can be partially or fully implemented using known components (or methods or processes), only those portions of such known components (or methods or processes) that are necessary for an understanding of the present invention will be described, and the detailed descriptions of other portions of such known components (or methods or processes) will be omitted so as not to obscure the invention. Further, the present invention encompasses present and future known equivalents to the components referred to herein by way of illustration.

**[0047]** Before describing the examples illustratively depicted in the several figures, a general introduction is provided for further understanding.

**[0048]** In some embodiments, a vector engine is implemented to detect network threats using metadata vectors. In some embodiments, the vector engine may first receive network communications, extract metadata from the network communications, group the metadata into clusters, and if the size of one or more metadata clusters

exceeds a given threshold within a predefined period of time, generate an alarm. The network comprises a plurality of client computers communicating with an outside network (i.e. the internet) as well as internally with each other. When a computer is infected by a malicious payload, such as malicious automated instructions or script, the computer may act as a bot and may send copies of the same or a different malicious payload to a number of other computers in the network, to infect them as well. The malicious automated script may include logic that instructs the first infected computer to send the same or a different automated script to multiple other computers. The script may further be addressed to the same destination port on each computer. Further, when the multiple other computers receive and process (e.g., execute/interpret) the script, they may respond to the first computer using the same type of response.

**[0049]** Thus, in this case, for example, the characteristics of such an attack would be a first computer in a network sending substantially the same data (e.g., the malicious payload) of a set size, addressed to the same destination port on multiple other computers, each of which may respond to the first computer using substantially the same response message. Further, the time between each request and its matching response will typically be short (e.g., less than five seconds).

**[0050]** In some embodiments, the vector engine can collect the above network communications by tapping a network switch and making a copy of the network communications for analysis, without disrupting or slowing down the network. Next, the vector engine may extract metadata from the network communications that can be used to characterize the network communications without inspecting the contents. The vector engine may group the metadata into matching clusters that allow the vector engine to

track how many times a given network communication has occurred. If a certain network communication, such as one corresponding to a malicious attack, occurs too many times, the cluster may grow in size and exceed a given threshold. If the threshold is exceeded or surpassed, the vector engine may generate an alarm or may flag the data for further monitoring.

**[0051]** As an example, if the threshold for a given cluster's size is ten, and a first computer is infected by a malicious payload and tries to spread copies of the same or different malicious payload to eleven other computers, then eleven instances of metadata will be generated by the vector engine. Further, because the eleven instances of metadata were all generated from the same source, the eleven instances of metadata will all have similar data items and can be grouped into the same cluster. As the threshold is ten, the size of the cluster containing the eleven instances exceeds the threshold and the vector engine may generate an alert to signal that the first computer is potentially attacking other computers on the network.

**[0052]** Described below are general non-limiting example definitions that may be used to understand or enhance understanding of how certain embodiments may be implemented.

**[0053]** A network is collection of computing entities that communicate with one another. In some embodiments, an internal network is a network of computers that are isolated from other (e.g., external) networks by networking components and/or devices as are known in the art, such as a firewall.

**[0054]** A client computer is a computer or node within a network. In some embodiments, a computer may act as a client and/or host depending on the type of



connection with another computer. As used here, clients are generally uninfected but can become infected by malicious payloads, in which case they may act as bots for the outside or attacking entity.

**[0055]** An attacking computer is a computer or entity (e.g., malicious user) that initiates a malicious attack on a computer, client computer, or network. In some instances, an attacking computer may use a malicious payload to attack other computers in the network, whereby the malicious payload may infect a first computer inside the network and direct the first computer to distribute copies of the same or another malicious payload to other computers, who then may further replicate and distribute the malicious payload.

**[0056]** A bot is a client computer that is under the effective control of the attacking computer. For example, the attacking computer can use the bot to achieve further malicious goals, and/or perform tasks.

**[0057]** A malicious payload is a collection of instructions, which may be in the form of code (e.g., JavaScript) that is used to perform malicious tasks. The payload may in some form be compilable code, scripted code, or other types of instructions as is known to those of ordinary skill in the art. In some instances, malicious payloads may correspond to large toolkits that are 500 megabytes or more in size. In other instances, a malicious payload may be a small automated script, such as a script of 1 megabyte (or less) in size.

**[0058]** Vector metadata is metadata that is used to characterize and/or categorize network communications without inspecting the contents of the network communications (e.g., deep content inspection). For example, in some embodiments

the metadata may comprise data from packet headers, and in some embodiments the metadata may be information that describes the overall session between two hosts. A metadata instance is a collection of one or more metadata items (e.g., packet information, session information) specific to a given network communication.

**[0059]** A metadata item is a parameter or item, such as a response size, that describes a certain feature of the network communication. For example, in some embodiments the metadata may comprise data from packet headers, and in some embodiments the metadata may be information that describes the overall session between two hosts. In some embodiments, a metadata item is used to characterize a network communication without inspecting the contents of the network communication. In some embodiments, the metadata item does not contain any reference or indicator of what is inside (e.g., the contents) of a network communication.

**[0060]** A metadata cluster is a grouping of metadata instances that may generally share similar metadata items. In some embodiments, a metadata cluster can be used to track how many network communications of a certain type have occurred.

**[0061]** A first communication is a communication initiated by one computer directed to another computer. In some embodiments, a first communication may take the form of a request or a “upstream” communication emanating from a client computer. In some embodiments, a first communication may correspond to a transmission that includes a file in a direction from a first computer to a second computer.

**[0062]** A second communication is a communication initiated by the second computer directed to the first computer. In some embodiments, a second communication may take the form of a response or a “downstream” communication

received from the host or responding computer. In some embodiments, a second communication may correspond to a response message initiated in response to receiving and/or processing the first communication.

**[0063]** A network communication is a communication or session between two or more computers in a network using various networking protocols, such as TCP/IP. In some embodiments, a network communication comprises the first communication and a second communication as a pair, session, or session portion. In some embodiments, a network communication comprises an initiating message (e.g., request) in one direction that begins the communication, and a response message in the opposite direction, which may be a response to the initiating message. In some embodiments, a network communication may be a unilateral communication attempt by one computer where no response occurs (e.g., a failed request).

**[0064]** FIG. 1 illustrates an example networked environment 100 in which an outside computer 104 is attacking an internal network 102. The internal network 102 comprises a plurality of client computers 106 and 108a-e. As illustrated, the attacking computer may coax the client computer 106 to retrieve a malicious payload 105, such as an automated script, by sending a malicious email with an embedded web link. When the client computer 106 processes the malicious payload 105, it becomes a "bot". Under direction of the malicious payload, the client computer or bot 106 may then attempt to send copies of the same malicious payload or a different malicious payload to a plurality of client computers 108a-e. When the plurality of client computers 108a-e receive and process the replicated malicious payloads, they may respond to the bot computer 106 with a response message (not depicted). The plurality client computers 108a-e, may

also send out more replicas of the malicious payload to potentially more client computers inside or outside the internal network 102. In this way, an outside attacking computer 104 may potentially use a small automated script to attack and/or establish control over many clients in an internal network 102. Though in this example only client five client computers are sent the payload, in some embodiments the payload may be sent to only two computers, in other embodiments, the payload may be sent to fifteen clients to be infected, in a network comprising 1000 clients for example.

**[0065]** FIG. 2A shows an example flowchart 200 of an approach for how a vector engine may be implemented to detect the network threats, such as the example attack described with reference to FIG. 1. Also, FIG. 2A is explained in conjunction with FIG. 2B and FIG. 2C, which show examples of how a vector engine may be implemented in a network setting.

**[0066]** At 202 (FIG. 2A), the vector engine receives network communications for a given network. FIG. 2B illustrates an example network in which step 202 may occur. There, an example of networking environment 210 comprises an internal network 212, which may be isolated by networking components and devices as are known in the art, such as a firewall (not depicted). A plurality of client computers 220 and 224a-e may be networked together to communicate through use of a network switch 216 or other networking components. In some embodiments, the vector engine 222 may receive network communications for the internal network 212 through the network switch 216. The vector engine 222 can also be implemented as software that resides upon networking equipment, nodes, or entities within a network (e.g., internal network 212).

Finally, the vector engine 222 can be implemented in a virtual machine and be deployed on hardware supplied by the company that owns the network.

**[0067]** In FIG. 2B, an outside computer or entity is illustrated as an attacker 214. The attacker 214 may interface with the internal network 212 through the switch 216 to send a malicious payload 227 (e.g., a payload linked through or attached to an email) to the client computer 220. When the client computer 220 processes the malicious payload 227, it becomes a "bot". Alternately, the client computer could become a "bot" under many other circumstances, including, but not limited to, getting infected when browsing a compromised web site.

**[0068]** Under the direction of the malicious payload (and in some cases under the direction of the attacker 214), the bot 220 sends a plurality of requests 225a-e (represented by dashed lines) to the respective client computers 224a-e. The requests 225a-e in this example are copies of the malicious payload 227 or of an entirely new malicious payload, and as such all requests 225a-e may have the same or similar sizes (e.g., ~1 MB). Further, the requests 225a-e may also be addressed to the same destination port on each client machine 224a-e, respectively (e.g., request 225a is addressed to example TCP port "135" on client 224a, while request 225b is addressed to the same example TCP port "135" but on client 224b.).

**[0069]** When the plurality of client machines 224a-e, receive and process the requests 225a-e, they may reply to the bot 220 with a response 229a-e, as illustrated by dotted lines in FIG. 2C. Since the requests 225a-e consisted of replicas of the same malicious payload, when each client machine 224a-e processes the request, the client

machines may send similar responses to the bot 220. Thus, the responses 229a-e may all be approximately the same size and share similar characteristics with one another.

**[0070]** As explained, in some cases, the malicious payload 227 consists of an automated script that when processed causes actions to occur automatically or near instantaneously. And in some case, the malicious payloads 225a-e, sent by the bot 220, is the same as the malicious payload 227 it received. Given the nature of malicious payloads, the high speed of networks, and the speed of computing devices, such as those shown in FIG. 2B and FIG. 2C, the time elapsed between the requests 225a-e (FIG. 2B) and the responses 229a-e (FIG. 2C) may be small (e.g., as little as five seconds).

**[0071]** Collectively, the communications (e.g., the requests and responses) comprise the network communications that the vector engine can receive and analyze at 202 (FIG. 2A). Next, at 204, the vector engine may extract vector metadata from the network communications. Though explained in greater detail below with reference to FIG. 4, vector metadata is briefly explained here as metadata that can be used to describe the network communications so that the vector engine can group or cluster the metadata and/or network communications for further analysis. For example, the metadata may comprise packet header information, or information about the session (e.g., time between request/responses).

**[0072]** At 206, after the metadata is extracted, the vector engine may group the metadata into clusters using a matching process. The metadata clusters may then be used to track how many instances of a given network communication has occurred. For example, with reference to FIG. 2B and FIG. 2C, the bot 220 has sent out five requests

(e.g., 225a-e), and received five responses (e.g., 229a-e), thus in one cluster there may be five instances of metadata that correspond to the five actual network communication exchanges.

**[0073]** The vector engine can set threshold size limits on the clusters to monitor them. The underlying rationale being that if a given cluster grows too large, the metadata instances in the cluster correspond to a rapidly growing network attack using malicious payloads, such as an automated script. For example, with reference to FIG. 2B and FIG. 2C, if the cluster (not depicted) that is tracking the five network communications (e.g., the five instances of requests and responses) has a threshold size limit of three, the vector engine may generate an alarm at 208 (FIG. 2A), to signal that the internal network 212 and specifically the bot 220 may be infected and attacking other clients when the five instances of metadata are grouped in one cluster.

**[0074]** FIG. 3 illustrates an example block diagram showing internal features of a vector engine 222. There, vector engine 222 contains a passive interface module 306 that is designed to passively receive network communications, such as from a network switch 304. The passive interface module may then send the network communications to an extractor module 308 to extract one or more instances of metadata that correspond to each network communication.

**[0075]** Once the metadata instances are extracted from the network communications, a cluster module may group the metadata instances into clusters. As explained, each cluster may have a threshold size limit such that if the amount of metadata instances in the cluster exceeds a threshold an alarm may be generated and output at 314.

**[0076]** FIG. 4 shows example details of a metadata instance for a given network communication. There, bot 220 has sent a request 225a containing a malicious payload to a client computer 224a. When the client computer 224a processes the malicious payload it may send a response 229a back to the bot 220. Collectively, the request 225a and the response 229a may comprise a single instance of a network communication 404. For example, network communication may correspond to a complete session or portion of a session. The extractor module 308 may receive the network communication 404 and extract one or more metadata items 403a-g that describe the network communication 404. The one or more metadata items correspond to a metadata instance 402 that describes the network communication 404 without the need for inspecting the contents of the request or the response. The contents of the network communication may comprise the “data” payload for a given network packet.

**[0077]** The metadata items illustrated serve merely as an example of a possible set of metadata that may be extracted from a network communication, though one of ordinary skill in the art appreciates that other metadata items can be implemented as necessary.

**[0078]** In some embodiments, a communication identifier such as a request identifier (RQID) 403a may correspond to the starting bytes of a packet. For example, request identifier 403a may comprise the first sixteen bytes of a request of communication 225a. Generally, the first bytes of a request contain header information or other data that is common across all requests 225a-e (FIG. 2B). In this way, the first sixteen bytes can be used to identify similar requests without inspecting the contents or data payload of the communications and/or packets.



**[0079]** Similarly, the response identifier (RSPID) 403, may be the first sixteen bytes of a response, such as communication 229a. Because the response is generally a message created in response to the malicious payload being processed (e.g., executed), responses then generally contain the same header information. Just as with RQID 403a, RSPID 403b can likewise be used to identify similar responses without inspecting the contents of the response.

**[0080]** The destination port metadata data item 403c corresponds to the protocol (e.g., TCP or UDP) and port on the receiving machine (e.g., 224a-e) to which a request is directed. As explained above, malicious payloads in the form of automated scripts may be restricted to small payload sizes to remain inconspicuous in the network. As such, malicious payloads often direct the bot 220 to send the replicated malicious payloads to the same destination port (as opposed to specifying multiple addresses/ports which may take up more space). Because the replicated malicious payloads (e.g., the requests) are all sent to the same destination port on each respective client machine 224a-e, the destination port metadata data item 403c may be used to match similar network communications that are initiated by the same malicious payload.

**[0081]** The request size metadata data item 403d describes the size of a request, such as a request 225a. Because the requests sent to clients 224a-e are generally the same, they likely share the same request size. If the requests sent to clients 224a-e are also copies of the original malicious payload 227, the sizes of the requests sent to clients 224a-e will be approximately the same as the size of payload 227. Thus, the request size metadata item 403d can be used to detect matching requests. Similarly,

the response size metadata data item 403e describes the size of a response, such as response 229a. Because the responses 229a-e (FIG. 2C) are typically all generated in response to the same requests 225a-e, the responses may all have approximately the same sizes as well. Thus, the response size 403e can be used to detect matching responses.

**[0082]** The time elapsed metadata data item 403f describes the length of time between when a request was sent out and when a response was sent back. Because the targeted hosts may be generally quick to process automated scripts and other types of network attacks, the time between request and response may be very short, e.g., less than five seconds. Thus, it may be the case that the duration of time between when each request 225a-e was sent and when each corresponding response was sent is roughly the same. Thus, the time elapsed data item 403f can be used to describe and match similar network communications.

**[0083]** The non-zero data volume metadata data item 403g is a metadata item that ensures that all analyzed network communications are not empty. For example, referring to FIG. 2B, if the request 225a contains actual malicious payload data, then the non-zero data volume metadata data item 403g will reflect that the request is not empty. However, if 225b was sent in error and is empty, then the non-zero metadata item 403g will reflect that the request 225b is empty. To save system resources, some embodiments of the invention may use the non-zero data volume metadata item 403g to filter out all empty network communications.

**[0084]** Collectively, the metadata items 403a-g may comprise a single metadata instance particular to a given network communication, such as network communication

404. Once the metadata items are extracted to compose one or more metadata instances, the cluster module can group the metadata instances into clusters using a cluster module.

**[0085]** FIG. 5A shows an example block diagram 500 of how aspects of the cluster module 310 may be implemented. There, a metadata instance 402 is received from the extractor module 308. The cluster module 310 may buffer one or more metadata instances in a metadata bin 504. As illustrated for this example, metadata bin 504 contains six metadata instances 505a-f. The cluster module 310 may group the metadata instances into metadata clusters (e.g., Cluster-A 516, Cluster-B 518, Cluster-C 520) using a matching process, explained in greater detail below. However, as an example, metadata instances 505d, 505e, 505f are all grouped into the same cluster, Cluster-C 520 because some or all of their metadata items have been determined to approximately match.

**[0086]** As explained, each cluster may have a threshold size limit such that if the threshold is exceeded an alarm is generated. In some embodiments, the cluster module 310 may implement a logic module 525 to keep track of cluster counts and a time window to expire old metadata instances. The cluster count module 506 can hold a data value that corresponds to the threshold size limit and further monitor the cluster count for each module. For example, if the cluster count module 506 sets the threshold to be "2", then Cluster-C 520, which currently holds three metadata instances (505d, 505e, and 505f), would exceed the threshold, and the cluster count module and/or the logic module 525 may generate an alarm data for output at 512.

**[0087]** In some embodiments, the time window module 508 may operate as a sliding window of a discrete duration. For example, if the sliding window time duration is 6 hours, as illustrated in FIG. 5A (see element 508), then metadata instances in the clusters that are older than 6 hours are expired and removed from the cluster. The expiring and removing of old metadata instances using a sliding window can protect against false positives caused by potential build-up of unrelated metadata instances over very long periods of time. Further, though specific examples are used here for the threshold size limit (e.g., 3) and the time duration (e.g., six hours), one of ordinary skill in the art appreciates that such parameters can be modified or customized per implementation, as required.

**[0088]** Thus, as FIG. 5A illustrates, the clusters may be required to satisfy two conditions before an alert may be triggered. First, the cluster must exceed the threshold size limit (e.g., more than three metadata instances), and second, the threshold size must be exceeded in the given time window (e.g., more than three metadata instances within the last six hours). In some embodiments, the second condition is implicit in the checking of the first condition. For example, by constantly removing metadata instances older than six hours, the second condition is automatically enforced.

**[0089]** In some embodiments, the alert data 512 generated by the cluster module 310 may generate an actual alarm automatically. That is, if the threshold size limit is exceeded within the time duration, an alarm is generated without intervention. In some embodiments, if the threshold size limit is exceeded within the time duration, the corresponding network activity is flagged as potentially malicious activity and no alarm is directly generated. In this way, the potentially malicious activity may be combined with

other unrelated anomalies to generate an automatic alarm or network administrators or security administrators can further monitor the network activity and manually generate an alarm as required.

**[0090]** Further, in some embodiments, a training window or period may be implemented, where the metadata instances are grouped into clusters to generate a baseline state of the network. The training period may be set to a time duration such as 24 hours. For example, the cluster module may group metadata instances into clusters for the first 24 hours to ascertain the appropriate value to set for each cluster's threshold size limit and may decide that special limits are required for some clients in the network. The training window could be used to identify periodic network communications that satisfy all the criteria to generate an alert, but that, as they happen relatively regularly in the network, could be associated to benign communications intentionally initiated from an uninfected client by a given user or process.

**[0091]** FIG. 5B shows a flowchart on an approach for how metadata instances may be grouped into clusters. At 526, the cluster module receives a group of metadata instances for analysis and grouping. At 528, the cluster module analyzes the metadata instances in a strict match stage. In the strict match stage 528, in order for one or more metadata instances to be grouped together, certain metadata items in each metadata instance must exactly match one another, or must exactly match metadata items of instances already in the cluster. For example, referring to FIG. 2B, if the request identifier (RQID) 403a for the request 225a exactly matches the request identifier (RQID) 403a for a different request 225b, then the corresponding metadata instances

may be grouped together in the same cluster (assuming the other requirements are met).

**[0092]** As illustrated in FIG. 5B, the strict match items for Cluster-A, as an example, may include the following metadata items: RQID 403a, RSPID 403b, and Destination Port 403c. Thus, to be grouped in Cluster-A at the strict match stage 528, a metadata instance must have the above metadata items exactly match the metadata instances already in Cluster-A or exactly matched to another metadata instance to start/form a cluster.

**[0093]** If a metadata instance under consideration has metadata items that do not strictly match those in Cluster-A, for example, then the metadata instance is not grouped with Cluster-A at 530. Though continuing, if the metadata instance under consideration does in fact have metadata items that exactly match each other or those in Cluster-A, then the metadata instance may be analyzed in the approximate match stage 532.

**[0094]** In the approximate match stage 532, a metadata instance is analyzed to determine whether certain metadata items in the metadata instance approximately match each other or those in the cluster. As illustrated in FIG. 5B, the metadata items in the approximate match stage 532 may include the request size 403d and the response size 403e (e.g., a communication size data item). Further, a range or percentage threshold may be specified to determine whether a given metadata item approximately matches the metadata items in the cluster. For example, in FIG. 5B the range is 5%. Thus, if a request size 403d is within 5% of the average of those request sizes already

in the cluster, the corresponding metadata instance may be grouped in the cluster at 534 (assuming other conditions illustrated in FIG. 5B are met).

**[0095]** FIG. 5C shows an example flowchart for an approach 575 for how the time window module 508 (FIG. 5A) may be implemented. After a start block at 576, two operations may occur approximately simultaneously. First, using a matching process, such as the one explained above, matching metadata instances can be added to a given cluster, at 578a. Second, as metadata instances are added to the cluster, the time in which they were added is recorded, and when a metadata item is more than six hours old it may be removed from the metadata cluster, at 578b. In this way, the time window module implements a sliding window that is six hours wide to incrementally remove old metadata instances, as new metadata instances are added. In some embodiments, the metadata bin 504 (FIG. 5A) may be allowed to buffer for six hours at a time. At each six-hour increment, the metadata instances buffered in the metadata bin 504 are then grouped into clusters. Once the metadata bin 504 is empty, new metadata instances may be buffered in the bin for another six hours. In this way, the time window module 508 may also implement a buffered or turn-based grouping process.

**[0096]** FIG. 5D illustrates a flowchart for one approach 582 for determining whether a threshold size limit has been exceeded. At 583, a new matching metadata instance is added to a cluster. At 584, after the new metadata instance is added to the cluster, a determination is made on whether the cluster contains more than the threshold size limit. For example, if the threshold size limit is 3, then at 584 the cluster count module 506 determines whether there are more than 3 metadata instances in the cluster. If the cluster does not have more than three metadata instances, the threshold size limit is not

exceeded and no alarm data is generated, at 585. On the other hand, if after adding the new matching metadata instance, the cluster contains more than three metadata instances, the threshold size limit is exceeded and an alarm may be generated, at 586.

## SYSTEM ARCHITECTURE OVERVIEW

**[0097]** FIG. 6 is a block diagram of an illustrative computing system 1400 suitable for implementing an embodiment of the present invention for performing intrusion detection. Computer system 1400 includes a bus 1406 or other communication mechanism for communicating information, which interconnects subsystems and devices, such as processor 1407, system memory 1408 (e.g., RAM), static storage device 1409 (e.g., ROM), disk drive 1410 (e.g., magnetic or optical), communication interface 1414 (e.g., modem or Ethernet card), display 1411 (e.g., CRT or LCD), input device 1412 (e.g., keyboard), and cursor control.

**[0098]** According to one embodiment of the invention, computer system 1400 performs specific operations by processor 1407 executing one or more sequences of one or more instructions contained in system memory 1408. Such instructions may be read into system memory 1408 from another computer readable/usable medium, such as static storage device 1409 or disk drive 1410. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and/or software. In one embodiment, the term “logic” shall mean any combination of software or hardware that is used to implement all or part of the invention.



**[0099]** The term “computer readable medium” or “computer usable medium” as used herein refers to any medium that participates in providing instructions to processor 1407 for execution. Such a medium may take many forms, including but not limited to, non-volatile media and volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as disk drive 1410. Volatile media includes dynamic memory, such as system memory 1408.

**[00100]** Common forms of computer readable media include, for example, floppy disk, flexible disk, hard disk, magnetic tape, any other magnetic medium, CD-ROM, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, RAM, PROM, EPROM, FLASH-EPROM, any other memory chip or cartridge, or any other medium from which a computer can read.

**[00101]** In an embodiment of the invention, execution of the sequences of instructions to practice the invention is performed by a single computer system 1400. According to other embodiments of the invention, two or more computer systems 1400 coupled by communication link 1415 (e.g., LAN, PTSN, or wireless network) may perform the sequence of instructions required to practice the invention in coordination with one another.

**[00102]** Computer system 1400 may transmit and receive messages, data, and instructions, including program, i.e., application code, through communication link 1415 and communication interface 1414. Received program code may be executed by processor 1407 as it is received, and/or stored in disk drive 1410, or other non-volatile storage for later execution.

**[00103]** In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. For example, the above-described process flows are described with reference to a particular ordering of process actions. However, the ordering of many of the described process actions may be changed without affecting the scope or operation of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than restrictive sense.

CLAIMS

What is claimed:

1. A method for detecting threats in networks using metadata, comprising:  
receiving one or more network communications, wherein a network communication includes at least a first communication in a first direction from a first host to a second host and a second communication in a second direction from a second host to a first host, the second direction being opposite of the first direction;  
extracting one or more metadata instances from the network communications that describe the network communications without inspecting contents of the network communications,  
grouping the metadata instances into one or more clusters based on whether the metadata instances match; and  
generating alert data for a cluster when an amount of metadata instances in the cluster exceeds a threshold.
2. The method of claim 1, wherein the clusters correspond to a sliding window in which metadata instances older than a time duration are removed.
3. The method of claim 1, wherein network communications are grouped into clusters by matching a plurality of metadata items in a first metadata instance to a plurality of metadata items in a second metadata instance.
4. The method of claim 3, wherein the metadata items includes one or more of the following group: a first communication identifier, a second communication identifier, a destination port data item, a first communication size data item, a second communication size data item, a time elapsed data item, and/or a non-zero data volume data item.

5. The method of claim 4, wherein two metadata instances are grouped into a same cluster if one or more metadata items belonging to each of the two metadata instances strictly match.
6. The method of claim 4, wherein two metadata instances are grouped into a same cluster if one or more metadata items belonging to each of the two metadata instances approximately match.
7. The method of claim 1, further comprising a training window in which network communications are grouped into clusters.
8. The method of claim 4, wherein the first communication identifier corresponds to a number of starting bytes of the first communication, and the second communication identifier corresponds to a number of starting bytes of the second communication.
9. A system for detecting threats in networks using metadata, comprising:
  - a computer processor to execute a set of program code instructions;
  - a memory to hold the set of program code instructions, in which the set of program code instructions comprises program code to perform any of the methods of claims 1-8.
10. A computer program product embodied on a non-transitory computer readable medium, the non-transitory computer readable medium having stored thereon a sequence of instructions which, when executed by a processor causes the processor to execute any of the methods of claims 1-9.





FIG. 1



200 →

FIG. 2A




FIG. 2B

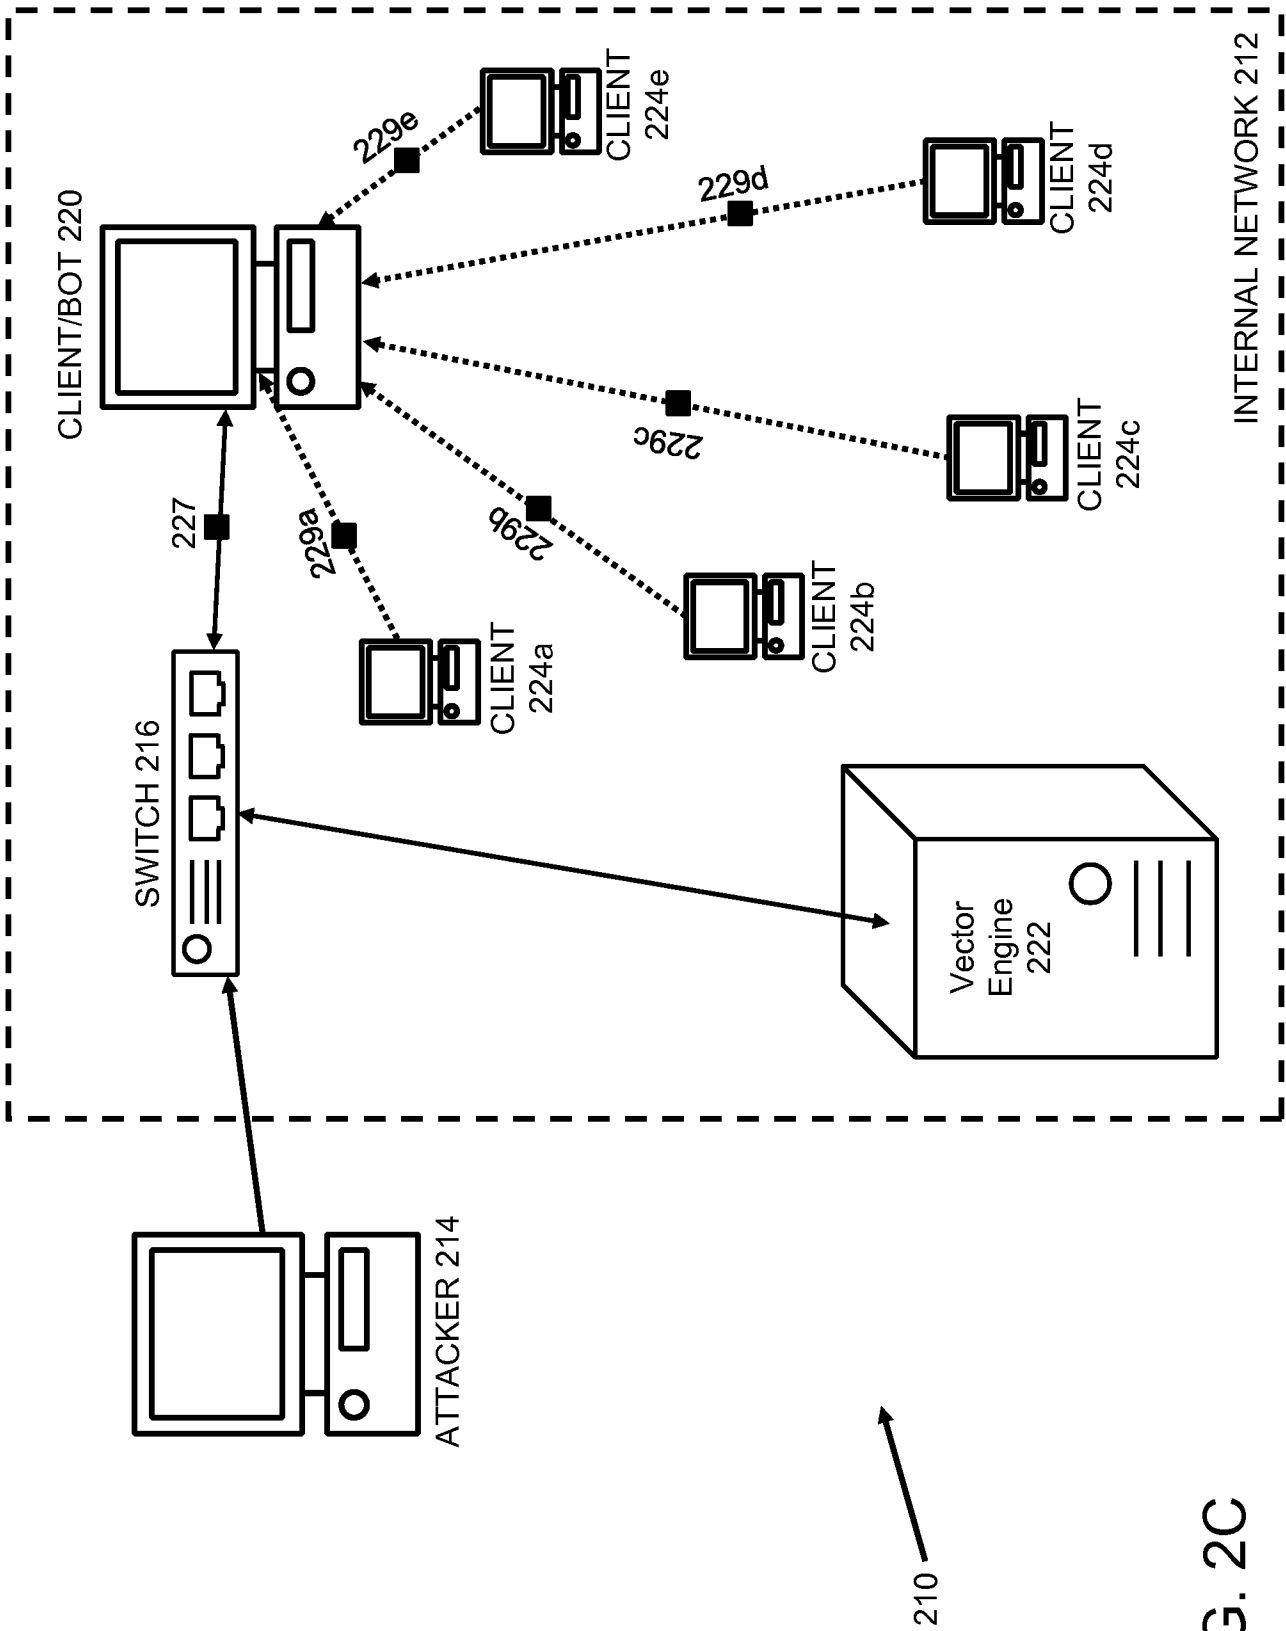


FIG. 2C






FIG. 3




FIG. 4

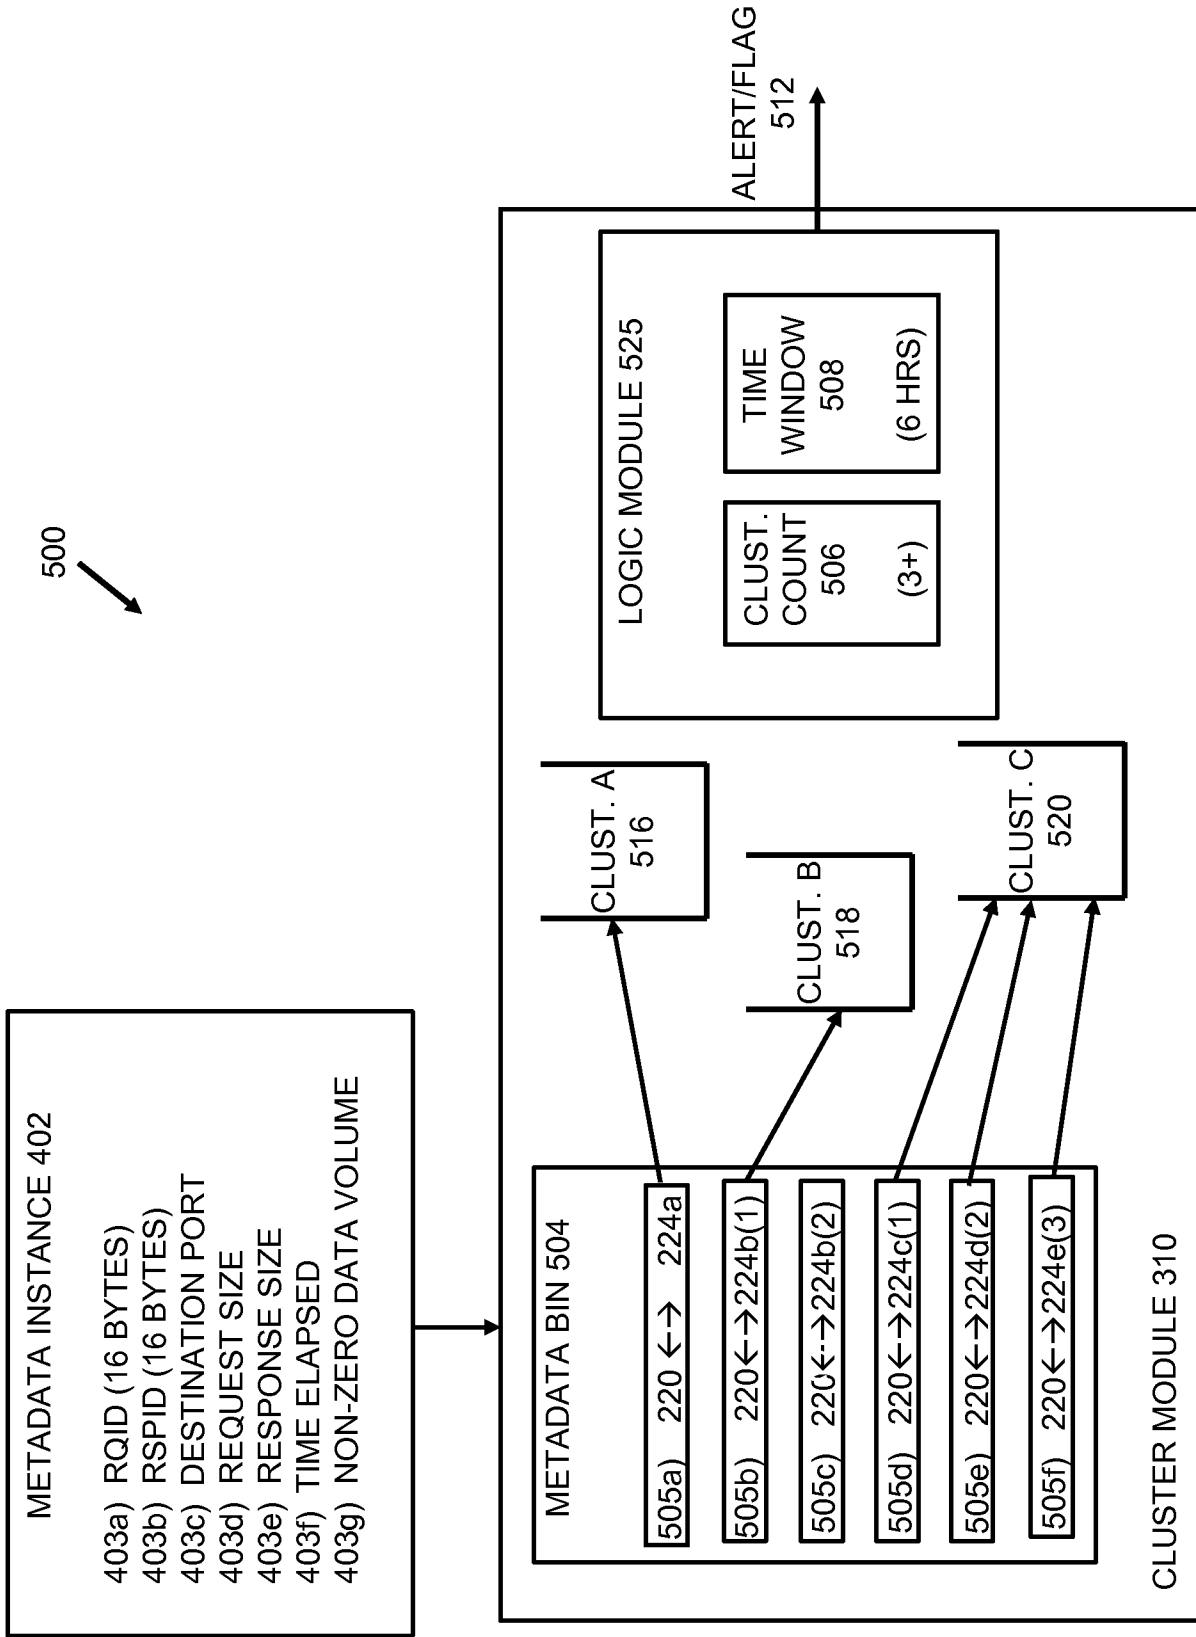


FIG. 5A

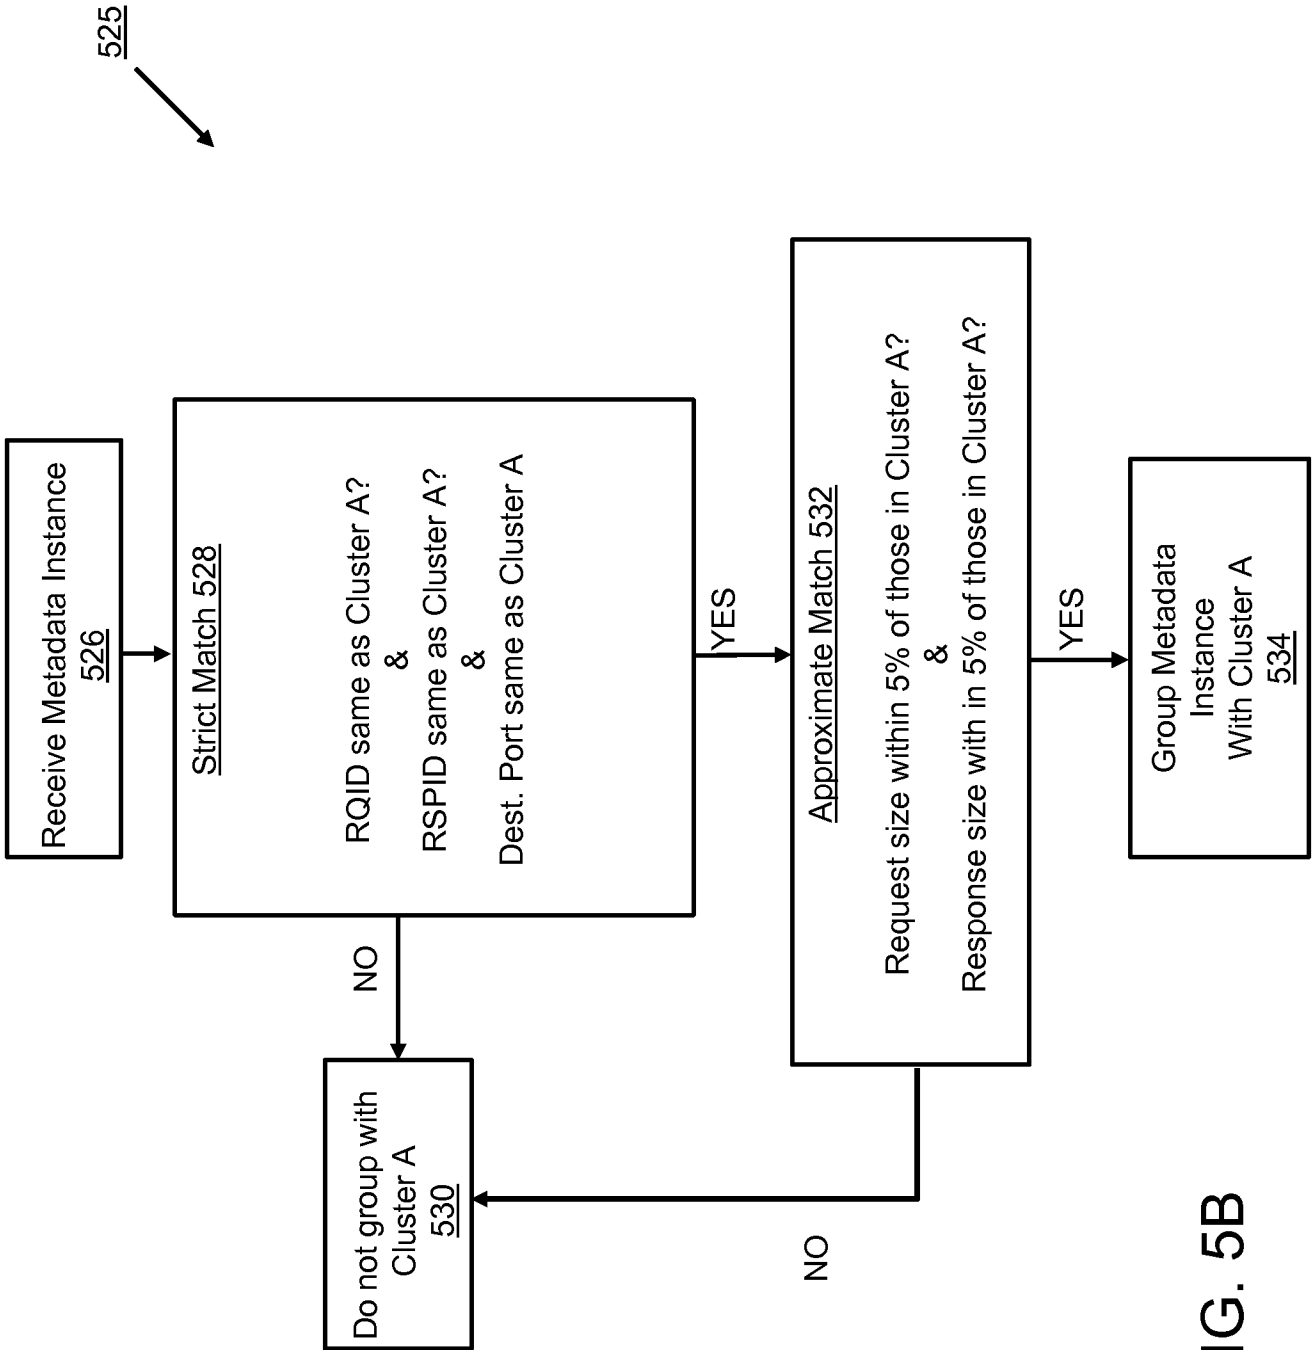


FIG. 5B

575 →




FIG. 5C





FIG. 5D

Figure 6



INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US 15/61191

A. CLASSIFICATION OF SUBJECT MATTER  
IPC(8) - G06F 11/00 (2015.01)  
CPC - H04L63/1416  
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
USPC: 726/22; CPC: H04L63/1416; IPC(8): G06F 11/00 (2015.01)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
USPC: 726/1, 726/22-26; CPC: H04L63/1416, H04L63/1408, H04L63/145, H04L63/1441, H04L63/1458; IPC(8): G06F 11/00 (2015.01)  
(keyword limited, terms below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
PatBase, Google Patents, IEEE; Search Terms: DDOS, distributed denial of service; attack; histogram; bucket; bin; standard deviation, sigma; outliers; filtering; time window; log; monitor; inspect; statistics, statistical analysis; suspicious; packet; request; identifier, IP address, IP protocol, IP TTL

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2003/0145232 A1 (Poletto et al.) 31 July 2003 (31.07.2003), entire document especially paras [0006], [0035], [0037], [0038], [0051], [0057], [0064], [0065], [0096]	1 - 9
A	US 2011/0138463 A1 (Kim et al.) 09 June 2011 (09.06.2011), entire document	1 - 9
A	US 2012/0124666 A1 (Kim et al.) 17 May 2012 (17.05.2012), entire document	1 - 9
A	US 2013/0219502 A1 (Danford et al.) 22 August 2013 (22.08.2013), entire document	1 - 9
A	US 2013/0219502 A1 (Danford et al.) 01 April 2013 (01.04.2013), entire document	1 - 9

Further documents are listed in the continuation of Box C.

\* 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 11 January 2016 (11.01.2016)	Date of mailing of the international search report <b>04 FEB 2016</b>
---	--

Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-8300	Authorized officer: Lee W. Young  PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774
---	--



## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 15/61191

**Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1.  Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
  
2.  Claims Nos.:  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
  
3.  Claims Nos.: 10  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

**Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)**

This International Searching Authority found multiple inventions in this international application, as follows:

1.  As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2.  As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3.  As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
  
4.  No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

**Remark on Protest**

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.