



- (51) **International Patent Classification:**  
G06F 15/173 (2006.01) H04L 12/28 (2006.01)  
H04L 12/26 (2006.01)
- (21) **International Application Number:**  
PCT/US2016/036025
- (22) **International Filing Date:**  
6 June 2016 (06.06.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
14/738,380 12 June 2015 (12.06.2015) US
- (71) **Applicant:** LEVEL 3 COMMUNICATIONS, LLC [US/US]; 1025 Eldorado Blvd., Broomfield, Colorado 80021 (US).
- (72) **Inventors:** CAPUTO II, Pete J.; 10705 Addison Ct., Highlands Ranch, Colorado 80126 (US). YERMAKOV, Sergey; 10475 Park Meadows Drive, Littleton, Colorado 80124 (US).
- (74) **Agents:** MUTSCHELKNAUS, Joseph E. et al.; Sterne, Kessler, Goldstein & Fox, PLLC, 1100 New York Ave., NW, Washington, District of Columbia 20005 (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:** NETWORK OPERATIONAL FLAW DETECTION USING METRICS

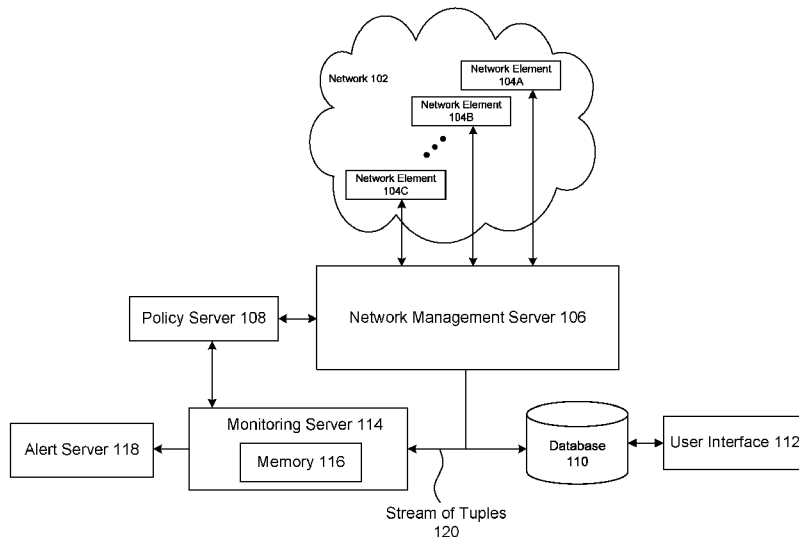


FIG. 1

(57) **Abstract:** In an embodiment, a computer-implemented method receives and monitors performance metrics from network element. The method also includes receiving periodic control metrics corresponding to object instances. Performance metrics and control metrics provide information about operation of object instances. By monitoring the metrics, a network server is able to detect an operational flaw in the network. Monitoring the performance and control metrics in real time increases the speed of detecting any operational flaw in the network.

WO 2016/200731 A1

## NETWORK OPERATIONAL FLAW DETECTION USING METRICS

## BACKGROUND

## Technical Field

**[0001]** Embodiments generally relate to monitoring network operation.

## Background

**[0002]** A communication network may, for example, provide a network connection that allows data to be transferred between two geographically remote locations. A network may include network elements connected by links. The network elements may be any type of managed device on the network, including routers, access servers, switches, bridges, hubs, IP telephones, IP video cameras, computer hosts, and printers. Network elements can be physical or logical and can communicate with one another via interconnected links.

**[0003]** Network operation may be impaired for different reasons. For example, component failure, including link failures or network element failure, may cause operational flaws in the network. Network operational flaws may also be caused by misconfiguration or malfunction of network elements or congestion in the network.

**[0004]** Networks may provide clients with statistics, reports, and other information related to their elements and their performance. For example, clients may wish to see how much their traffic is delayed by the network, whether the service is meeting service level agreements, whether the network elements are functioning as intended, etc. Such performance information helps to discover operational flaws in the network.

**[0005]** Network management servers may use performance metrics that provide information about network performance. To collect metrics, a standard protocol, such as Simple Network Management Protocol (SNMP), may be used. SNMP is part of the Internet Protocol Suite as defined by the Internet Engineering Task Force (IETF). It includes of a set of standards for network management, including an application layer protocol, a database schema, and a set of data objects.

**[0006]** The database schema SNMP uses a management information base (MIB). The MIB describes network objects. The objects can be specified by object identifiers (OID). An object can include one or more object instances.

[0007] SNMP may support a query providing for discovery of the instances available for an object. The instances may be identified by suffixes to the object identifiers. The instance identifier may be used to retrieve metric values for the instance. Performance metric values stored in a database may be used to determine if there is an operational flaw with the corresponding object instance.

## SUMMARY

[0008] In an embodiment, a computer-implemented method detects operational flaws in a network. The method includes receiving, by a network management server, metrics corresponding to object instances of objects associated with network elements. The network management server generates a stream of tuples, wherein a tuple includes a network element identifier, a metric, and an object instance identifier. A metric is either a performance metric or a control metric corresponding to the object instance, and wherein the control metric appears periodically in the stream. A monitoring server monitors the stream of tuples and sets a flag when a first performance metric corresponding to the object instance is received. When detecting a control metric corresponding to the object instance the monitoring server determines whether the first flag is set. The monitoring server creates an alert if the first flag is not set.

[0009] Method and computer-readable medium embodiments are also disclosed.

[0010] Further embodiments and features, as well as the structure and operation of the various embodiments, are described in detail below with reference to accompanying drawings.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The accompanying drawings are incorporated herein and form a part of the specification.

[0012] FIG. 1 illustrates a network environment where a monitoring server monitors performance of object instances.

[0013] FIG. 2A illustrates a method of generating a stream of tuples corresponding to object instances.

- [0014] FIG. 2B illustrates a method of generating a stream of tuples corresponding to object instances.
- [0015] FIG. 3A illustrates a method of monitoring a stream of tuples corresponding to object instances for detecting operational flaws.
- [0016] FIG. 3B illustrates a method of monitoring a stream of tuples object instance operational flaw similar to the method in FIG. 3A, but with a .
- [0017] FIG. 4 illustrates an example memory a monitoring server uses to monitor performance metrics of object instances.
- [0018] FIG. 5 illustrates an example of a method for monitoring a stream of tuples object instance operational flaw.
- [0019] FIG. 6 illustrates another example memory a monitoring server uses to monitor performance metrics of an object instance.
- [0020] In the drawings, like reference numbers generally indicate identical or similar elements. Generally, the left-most digit(s) of a reference number identifies the drawing in which the reference number first appears.

## DETAILED DESCRIPTION

- [0021] A network management server may receive performance metrics corresponding to instances of network objects associated with network elements. Network objects may include network services, applications or processes for a client. A network management server may collect metrics from network elements by sending request messages. Alternatively, network elements may send the metrics to the network management server. After receiving the metrics, network management server generates a stream of tuples. Each tuple corresponds to an object instance and may include an identification of the corresponding network element, a performance metric value from the object instance, and an identification of the object instance.
- [0022] The performance metrics may be used to determine operation status of their corresponding object instances. A performance metric value may for example describe operational status of its corresponding object instance. The performance metrics may be stored in a database.
- [0023] According to embodiments, a monitoring server monitors the stream of tuples that the network management server generates. The monitoring server may receive the stream

in parallel with a database or before the stream is sent to the database. The monitoring server may therefore monitor the performance metrics in the stream in real time to detect any possible operational flaws without a need to access the database.

**[0024]** In embodiments, by monitoring the frequency or value of specific performance metrics in the stream of tuples, the monitoring server is able to detect any flaws in operation of the instances of network objects associated with the performance metrics. By monitoring of the tuple stream in real time using the monitoring server, as opposed to using a database, embodiments can reduce the time to detect any flaws in operation of object instances.

**[0025]** In the description that follows, a system where the monitoring server monitors the stream is first described with respect to FIG. 1. The network management server's creating a tuple stream is described with respect to FIGs. 2A and 2B. How the monitoring server uses the tuple stream to detect network flaws is described with respect to FIGs. 3A and 3B. An example memory used by monitoring server is described with respect to FIG. 4. FIG. 5 walks through an example of the monitoring server's operation.. FIG. 6 provides another example memory used by monitoring server .

**[0026]** FIG. 1 illustrates a network environment 100 that includes a plurality of network elements 140A-C in network 102. Network elements 140 may be connected to a network management server 106. Network management server 106 receives metrics corresponding to network object instances on the network elements.

**[0027]** Network management server 106, using performance metrics received from network elements in network 102, creates a stream of tuples. Each tuple in the stream corresponds to an instance of an object associated with a network element and may include an identifier of the network element, a metric corresponding to the object instance, and an identifier of the object instance.

**[0028]** Network element identifier may for example be an IP or MAC address of a network device. As mentioned above, the objects may be specified using OIDs. Each OID may be a sequence of integers separated by decimal points. An example of an OID is 1.3.6.1.2.1.4.6. Each OID may have a textual description as well. For example the textual description of 1.3.6.1.2.1.4.6 may be iso.org.dod.internet.mgmt.mib-2.ip.ipForwDatagrams. In that example, the ipForwDatagrams object may be an integer counter that stores the number of forwarded datagrams at a router. As mentioned above,

the specification may include a list of specific objects, for example, 1.3.6.1.2.1.4.4, 1.3.6.1.2.1.4.5, and 1.3.6.1.2.1.4.6. Or the specification may designate an entire portion of the object hierarchy, for example, 1.3.6.1.2.1.4.\*. Objects can be have instances. When the objects are stored in a treelike hierarchy, the instances may be all of the object's children. They may, for example, be an entire subtree with the object specified by an object identifier as its root. The requests and response may be formatted according to SNMP. The object instances may also have an object ID. For example, the object instances may have the object's ID, appended with a new suffix.

**[0029]** Examples for performance metrics may include passive metrics such as InOctets or OutOctets provided by a network switch. Other examples of performance metrics may include active metrics such as round trip time (RTT), latency, or jitter of traffic on a service provided by the network. For example, a tuple may be {Switch Port MAC address, Performance Metric, Instance Identifier}. In an embodiment, network management server 106 uses the method illustrated in FIGs. 2A or 2B to create the stream of tuples.

**[0030]** Network management server 106 may be connected to a database 110. Performance metrics generated by network management server 106 may be stored in database 110 and accessible to users through user interface 112.

**[0031]** In an embodiment, monitoring server 114 receives the stream of tuples from network management server 106. Monitoring server 114 monitors performance metrics received in the stream of tuples 120 to detect any flaw with operation of object instances corresponding to the performance metrics. Monitoring server 114 monitors the stream of tuples 120 as it is generated by the network management server 106 without a need to query database 110. Monitoring server 114 may therefore monitor the performance stream of tuples in real time to detect any operational flaws in the network 102. In this way, monitoring server 114 detects an operational flaw in the network much quicker than a user accessing database 110 may realize the same operational flaw.

**[0032]** In an embodiment, monitoring server 114 uses the method illustrated in FIGs. 3A or 3B to detect an operational flaw of an object instance. Monitoring server 114 may temporarily store flags corresponding to the received performance metrics in a memory 116. Monitoring server 114 may detect flaws when a flag shows missing performance metric for an object instances when the metric is expected. Monitoring server 114 may

also detects operational flaw when the value of the performance metric is not as it is expected to be for the corresponding object instance.

**[0033]** When monitoring server 114 detects an operational flaw in an object instance, it may inform alert server 118 of operational flaw in the object instance, the corresponding performance metric indicating the flaw, and identification of associated network element and object instance. Alerts server 118, using this information, alerts a network operator of the operational flaw. Alert server 118 may send the network operator information in the tuple that contains the metric indicating of the flaw. Alert server 118 may inform the network operator of the operational flaw through email, or any other messaging system using wired or wireless networks, for example cellular network.

**[0034]** FIG. 2A illustrates a method of generating a stream of tuples, which may be used in operation of network management server 106. At step 210A, network management server 106 receives performance metrics corresponding to object instances in network 102. Network management server 106 may also receive identification of the corresponding object instance and element when receiving a performance metric in step 210A.

**[0035]** At step 220A, network management server 106 generates control metrics for each object instance. Control metrics are generated periodically. The frequency of the control metrics may be determined according to a business rule set in policy module 108 and communicated to the network management server 106. Policy module 108 may set the frequency of the control metric for an object instance. In embodiments, policy module 108 may set the frequency of the control metric for an object instance less than the frequency of any performance metric associated with the object instance.

**[0036]** The control metrics act as a heartbeat for a network element. They set a pace for arrival of the performance metrics. Control metrics may notify the monitoring server of arrival expectancy of performance metrics for a corresponding network instance. Control metrics may not have a value describing performance or functionality of a network instance, but their arrival frequency sets a pace for receiving and evaluating performance metrics. Control metrics of a network instance arrive less frequently than any performance metrics for the network instance. Therefore when a control metric corresponding to a network instance is arrived, network monitoring server expects that all the performance metrics for the network instance to have been arrived.

- [0037]** At step 230A, network management server 106 may generate the stream of tuples using performance metrics and control metric for each object instance. In an embodiment, network management server 106 generates and sends a tuple for each performance metric that it receives in step 210A. Network management server 106 may generate tuples including the control metric of an object instance according to the frequency of the control metric for the object instance as specified by policy server 108.
- [0038]** After receiving a performance metric, network management server 106 may generate a tuple that includes the performance metric, identification of the corresponding object instance, and identification of the corresponding network elements. When a control metric is generated according to a business rule set by policy server 108 for an object instance, network management server 106 generates and sends a tuple that includes the control metric, identification of the corresponding object instance, and identification of the corresponding network element. Therefore tuples that include control metrics appear periodically for each object instance, according to the business rule set by policy server 108 for each object instance. Network management server 106 may also format the tuple according to an instruction set by policy module 108.
- [0039]** FIG. 2B illustrates a method of generating a stream of tuples by network management server 106 according to another embodiment. At step 210B, network management server 106 receives performance metrics and periodic control metrics corresponding to object instances in network 102.
- [0040]** At step 230B, network management server 106 may generate the stream of tuples using performance metrics and control metrics for each object instance. In an embodiment, network management server 106 generates and sends a tuple for each performance or control metric that it receives. After receiving a performance metric, network management server 106 may generate a tuple that includes the performance metric, identification of the corresponding object instance, and identification of the corresponding network element. After receiving a control metric, network management server 106 generates and sends a tuple that includes the control metric, identification of the corresponding object instance, and identification of the corresponding network element. Therefore tuples that include control metrics appear periodically for each object instance, according to the frequency they are received by network management server 106.

- [0041] In embodiments, network management server 106 receives performance metrics from network 102 and monitoring server 114 uses a least frequent performance metric corresponding to a network instance as a control metric for the instance. Policy server 108 may configure network management server 106 and monitoring server 114 to use the least frequent metric of a network instance as the control metric. The least frequent performance metric may function as both a performance and the control metric for the network instance. Thus, control metric role can be assumed by a regular "performance" metric, for example ifAlias (1.3.6.1.2.1.31.1.1.1.18). In that case, the control metric injection/generation described in step 220A may be unnecessary.
- [0042] FIG. 3A illustrates a method of monitoring streams of tuples 120 by monitoring server 114 for detecting an operational flaw in network 102. At step 310, monitoring server 114 may receive a tuple on stream of tuples 120. At step 320 monitoring server 114 may detect the object instance associated with the received tuple, using the field in the tuple that contains the identification of the object instance.
- [0043] At step 330, monitoring server 114 may detect a metric in the received tuple. Monitoring server 114 may gain information about formatting of the tuple, including location of the metric or other information in the tuple, from policy server 116.
- [0044] At step 340, monitoring server 114 determines whether the metric in the received tuple is a performance metric or a control metric. If it is a performance metric, at step 360, monitoring server 114 will set a flag in memory 116 corresponding to the performance metric. In an embodiment, monitoring server 114 sets the flag in a portion of memory 116 allocated to the object instance. An example embodiment of allocating portions of memory 116 to object instances is shown in FIG. 4.
- [0045] In normal operation of an object instance, all tuples including performance metrics for the object instance are generated between two consecutive tuples including control metrics for the object instance. Monitoring server 114 therefore expects certain performance metrics for an object instance before the control metric for the object instance arrives. Monitoring server 114 may track what performance metrics to expect for an object instance from policy server 116. If monitoring server 114 does not receive all the performance metrics expected for the object instance by the time it receives the corresponding control metric, there is a flaw in operation of the object instance.

- [0046] If at step 340, network management server 106 determines that the metric in the received tuple is a control metric, at step 350 monitoring server 114 determines whether all the flags indicating receipt of all the performance metrics expected for the object instance are set. If not all the flags are set, at step 370, monitoring server 114 informs alert server 118 to generate an alert. Monitoring server 114, at step 380, erases all the flags in the portion of memory 116 corresponding to the object instance and then continues to monitor the incoming tuples in method 300A.
- [0047] If at step 350, monitoring server 114 determines that the flags corresponding to all the performance metrics are set, monitoring server 114, at step 380, erases all the flags in the portion of memory 116 corresponding to the object instance. Monitoring server 114 then continues to monitor the incoming tuples in method 300A.
- [0048] In another embodiment, shown in FIG. 3B, after setting a flag corresponding to a received performance metric at step 360, monitoring server, at step 362, determines whether the performance metric has a correct value. Monitoring server 114 may gain information on the accepted values for performance metric of an object instance from policy server 108.
- [0049] If monitoring server 114, at step 362 determined that the performance metric does not have a correct value, it will inform alert server 118 to generate an alert at step 364. The alert may include information indicating the received value of the performance metric, and the corresponding identification of the object instance and the network element. Other steps of method 330B in FIG. 3B is the same as method 300A in FIG. 3A.
- [0050] FIG. 4 shows an embodiment for a portion of memory 116 used to store metric flags 440 corresponding to performance metrics received for object instances 420. In the example embodiment of FIG. 4, each columns is assigned to an object instance. Each row is assigned to a performance metric flag. A flag may be one bit in the memory. Value "1" may indicates that a flag is set and value "0" in a memory location may indicate that the flag is reset.
- [0051] After a control metric arrives for an object instance, all the memory locations in the columns of the object instance are reset to "0." Any time a performance metric corresponding to a flag is arrived, the corresponding row in the object instance columns is set to "1." When next control metric arrives, monitoring server 114 by examining the

column assigned to the object instance determines if all the locations corresponding to the performance metrics that should have arrived for the object instance have value “1.”

**[0052]** FIG. 5 shows an example of method 300A and FIG. 6 shows an example of memory 400. Steps of method 500 are described below with reference to memory 600 in FIG. 6. In this example, the object instances may have two performance metrics iso.org.dod.internet.mgmt.mib-2.interfaces.ifTable.ifEntry.ifInOctets (ifInOctet for short, OID: 1.3.6.1.2.1.2.2.1.10) ( $\text{metric}_{\text{IN}}$ ) and iso.org.dod.internet.mgmt.mib-2.interfaces.ifTable.ifEntry.ifOutOctets (ifOutOctet for short, OID: 1.3.6.1.2.1.2.2.1.16) ( $\text{metric}_{\text{OUT}}$ ), and a control metric ( $\text{metric}_{\text{CTL}}$ ). Monitoring server 114, at step 510 receives a tuple  $\{\text{element}_{\text{ID}}, \text{metric}_{[\text{IN}, \text{OUT}, \text{CTL}]}, \text{instance}_{\text{ID}}\}$  from stream of tuples 120. First field in the tuple,  $\text{element}_{\text{ID}}$ , includes an identification of the corresponding network element, for example IP or MAC address of a router. Second field may be ifInOctets metric ( $\text{metric}_{\text{IN}}$ ), ifOutOctets metric ( $\text{metric}_{\text{OUT}}$ ), or control metric ( $\text{metric}_{\text{CTL}}$ ). The third field in the tuple  $\text{instance}_{\text{ID}}$  includes an indication of the corresponding object instance. In normal operation of the object instance, tuples including  $\text{metric}_{\text{CTL}}$  arrives less frequently than  $\text{metric}_{\text{IN}}$  or  $\text{metric}_{\text{OUT}}$  in the stream of tuples 120. Therefore when network management server 106 receives  $\text{metric}_{\text{CTL}}$ , it expects tuples that include  $\text{metric}_{\text{IN}}$  and  $\text{metric}_{\text{OUT}}$  to have already been received.

**[0053]** Monitoring server 114, at step 520, reads the third field  $\text{instance}_{\text{ID}}$  in the received tuple. Monitoring server 114, at step 530 reads the second field, which is the metric, from the tuple.

**[0054]** If monitoring server 114 at step 540 determines that the metric in the received tuple is a  $\text{metric}_{\text{IN}}$  or  $\text{metric}_{\text{OUT}}$ , it will write value “1” in the corresponding memory location to the metric and the  $\text{instance}_{\text{ID}}$ . For example, referring to FIG. 6, if  $\text{metric}_{\text{IN}}$  is arrived for  $\text{instance}_{\text{ID}}=2$ , monitoring server 114 writes value 1 in the first row and the third column. Monitoring server 114 then continues to receive other tuples on the tuple stream. Because control metric for the  $\text{instance}_{\text{ID}}$  arrives less frequently than all other metrics for  $\text{instance}_{\text{ID}}$ , if there is no flaws in operation of the object instance, both memory bits for  $\text{instance}_{\text{ID}}$  will be “1” when  $\text{metric}_{\text{CTL}}$  arrives.

**[0055]** If monitoring server 114 at step 540 determines that the metric in the received tuple is a control metric, at step 550 it verifies whether  $\text{metric}_{\text{IN}}$  and  $\text{metric}_{\text{OUT}}$  bits in the memory column corresponding to  $\text{instance}_{\text{ID}}$  have value “1.” If they both have value “1,”

monitoring server 114 writes “0” on metric<sub>IN</sub> and metric<sub>OUT</sub> bits of instance<sub>ID</sub> column and continues to receive tuples on the tuple stream. If one or both of the values corresponding to instance<sub>ID</sub> in memory 600 is “0,” monitoring server 114 informs alert server 118 to create an alert. Monitoring server 114 then resets both bits corresponding to instance<sub>ID</sub> and continues to receive tuples on the tuple stream.

### *Conclusion*

- [0056] Each of the blocks and modules in FIGs. 1, 4, or 6 may be implemented in hardware, software, firmware, or any combination thereof.
- [0057] Each of the blocks and modules in FIGs. 1, 4, or 6 may be implemented on the same or different computing devices. Such computing devices can include, but are not limited to, a personal computer, a mobile device such as a mobile phone, workstation, embedded system, game console, television, set-top box, or any other computing device. Further, a computing device can include, but is not limited to, a device having a processor and memory, including a nontransitory memory, for executing and storing instructions. The memory may tangibly embody the data and program instructions. Software may include one or more applications and an operating system. Hardware can include, but is not limited to, a processor, memory, and graphical user interface display. The computing device may also have multiple processors and multiple shared or separate memory components. For example, the computing device may be a part of or the entirety of a clustered computing environment or server farm.
- [0058] Identifiers, such as “(a),” “(b),” “(i),” “(ii),” etc., are sometimes used for different elements or steps. These identifiers are used for clarity and do not necessarily designate an order for the elements or steps.
- [0059] The present invention has been described above with the aid of functional building blocks illustrating the implementation of specified functions and relationships thereof. The boundaries of these functional building blocks have been arbitrarily defined herein for the convenience of the description. Alternate boundaries can be defined so long as the specified functions and relationships thereof are appropriately performed.
- [0060] The foregoing description of the specific embodiments will so fully reveal the general nature of the invention that others can, by applying knowledge within the skill of the art, readily modify and/or adapt for various applications such specific embodiments, without undue experimentation, without departing from the general concept of the present

invention. Therefore, such adaptations and modifications are intended to be within the meaning and range of equivalents of the disclosed embodiments, based on the teaching and guidance presented herein. It is to be understood that the phraseology or terminology herein is for the purpose of description and not of limitation, such that the terminology or phraseology of the present specification is to be interpreted by the skilled artisan in light of the teachings and guidance.

**[0061]** The breadth and scope of the present embodiments should not be limited by any of the above-described examples, but should be defined only in accordance with the following claims and their equivalents.

## WHAT IS CLAIMED IS:

1. A computer implemented method for detecting operation flaws in a network comprising of network elements, the method comprising:
  - (a) receiving, by a network management server, metrics corresponding to network instances of objects associated with network elements;
  - (b) generating, by the network management server, a stream of tuples, wherein a tuple includes a network element identifier, a metric, and a network instance identifier, and wherein the metric is either one of one or more performance metrics or a control metric corresponding to the network instance, and wherein the control metric appears periodically in the stream;
  - (c) monitoring, by a monitoring server, the stream of tuples;
  - (d) setting, by the monitoring server, a first flag when a first performance metric corresponding to the network instance is received;
  - (e) detecting, by the monitoring server, the control metric corresponding to the network instance;
  - (f) determining, by the monitoring server, after detecting the control metric in step (e), whether the first flag is set;
  - (g) creating an alert, by an alert server, when step (f) determines that the first flag is not set, wherein the alert indicates an operational flaw; and
  - (h) resetting the first flag, by the monitoring server after it detects the control metric in step (e).
  
2. The method of claim 1, further comprising allocating memory to each network instance, and wherein setting the first flag in step (d) includes storing a value corresponding to the first performance metric in the allocated memory for the

instance, and resetting the first flag in step (h) includes erasing the stored value corresponding to the first performance metric of the instance.

3. The method of claim 2, wherein steps (d), (f), and (h) are performed using the memory allocated to the network instance.

4. The method of claim 1, further comprising:

(i) setting, by the monitoring server, a second flag when a second performance metric corresponding to the object instance is received;

wherein determining in step (f) further comprises determining, after detecting the control metric in step (e), whether the second flag is set;

wherein creating the alert in step (g) further comprises creating an alert when either of the first flag or the second flag are not set in step (f), and wherein the alert includes information in the tuple corresponding to the first performance metric when the first flag is not set in step (f) or to the second performance metric when the second flag is not set in step (f), respectively; and

wherein resetting in step (h) further comprising resetting the second flag by the status monitoring module after it detects the control metric in step (e).

5. The method of claim 4, wherein setting the second flag in step (i) includes storing a value corresponding to the second metric in the allocated memory for the instance, and resetting the second flag in step (h) includes erasing the stored value corresponding to the second flag stored in the allocated memory.

6. The method of claim 1, wherein the steps (a)-(h) are performed in real-time.

7. The method of claim 1, further comprising:

(j) monitoring, by the monitoring server, a value of the first performance metric, wherein the value indicates operation status of the object instance;

and wherein creating an alert in step (g) further comprises creating an alert when the value indicates operational flaw for the object instance.

8. A network monitoring system for detecting operational flaws in a network comprising of network elements, the system comprising:

a network management server configured to:

receive metrics corresponding to object instances of objects associated with network elements;

generate a stream of tuples, wherein a tuple includes a network element identifier, a metric, and an object instance identifier, and wherein the metric is either one of one or more performance metrics or a control metric corresponding to the object instance, and wherein the control metric appears periodically in the stream;

a monitoring server configured to:

monitor the stream of tuples;

set a first flag when a first performance metric corresponding to the object instance is received;

detect the control metric corresponding to the object instance;

determine, after detecting the control metric, whether the first flag is set; and

reset the first flag after it detects the control metric;

an alert server configured to:

create an alert when monitoring server determines that the first flag is not set, wherein the alert includes information in the tuple corresponding to the first performance metric.

9. The system of claim 8, wherein the monitoring server further comprises a memory, and the monitoring server further configured to allocate a memory portion to each object instance, and wherein setting the first flag includes storing a value corresponding to the first performance metric in the allocated memory portion for the instance, and resetting the first flag includes erasing the stored value corresponding to the first performance metric of the instance.

10. The system of claim 8, wherein

the monitoring server is further configured to:

set a second flag when a second performance metric corresponding to the object instance is received;

determine, after detecting the control metric, whether the second flag is set; and

reset the second flag after it detects the control metric;

the alert server further configured to:

create an alert when either of the first flag or the second flag are not set, after detecting the control metric, and wherein the alert includes information in the tuple corresponding to the first performance metric or to the second performance metric when the first flag or the second flag is not set, respectively.

11. The system of claim 10, wherein the monitoring server is further configured to store a value corresponding to the second metric in the allocated memory portion for the instance to set the second flag, and erasing the stored value corresponding to the second flag stored in the allocated memory portion to reset the second flag.
12. The system of claim wherein the monitoring server is further configured to monitor a value of the first performance metric, wherein the value indicates operation status of the object instance, and wherein the alert server is further configured to create an alert when the value indicates operational flaw for the object instance.
13. A non-transitory computer-readable medium having instructions stored thereon that, when executed by at least one computing device, causes the at least one computing device to perform a method for detecting operational flaws in a network comprising of network elements, the method comprising:
  - (a) receiving, by a network management server, metrics corresponding to object instances of objects associated with network elements;
  - (b) generating, by the network management server, a stream of tuples, wherein a tuple includes a network element identifier, a metric, and an object instance identifier, and wherein the metric is either one of one or more performance metrics or a control

metric corresponding to the object instance, and wherein the control metric appears periodically in the stream;

- (c) monitoring, by a monitoring server, the stream of tuples;
- (d) setting, by the monitoring server, a first flag when a first performance metric corresponding to the object instance is received;
- (e) detecting, by the monitoring server, the control metric corresponding to the object instance;
- (f) determining, by the monitoring server, after detecting the control metric in step (e), whether the first flag is set;
- (g) creating an alert, by an alert server, when step (f) determines that the first flag is not set, wherein the alert includes information in the tuple corresponding to the first performance metric; and
- (h) resetting the first flag, by the monitoring server after it detects the control metric in step (e).

14. The computer-readable medium claim 13, wherein the method further comprising allocating memory to each object instance, and wherein setting the first flag in step (d) includes storing a value corresponding to the first performance metric in the allocated memory for the instance, and resetting the first flag in step (h) includes erasing the stored value corresponding to the first performance metric of the instance.

15. The computer-readable medium claim 14, wherein steps (d), (f), and (h) are performed using the memory allocated to the object instance.

16. The computer-readable medium claim 13, wherein the method further comprising:

- (i) setting, by the monitoring server, a second flag when a second performance metric corresponding to the object instance is received;

wherein determining in step (f) further comprises determining, after detecting the control metric in step (e), whether the second flag is set;

wherein creating the alert in step (g) further comprises creating an alert when either of the first flag or the second flag are not set in step (f), and wherein the alert includes information in the tuple corresponding to the first performance metric when the first flag is not set in step (f) or to the second performance metric when the second flag is not set in step (f), respectively; and

wherein resetting in step (h) further comprising resetting the second flag by the status monitoring module after it detects the control metric in step (e).

- 17. The computer-readable medium claim 16, wherein setting the second flag in step (i) includes storing a value corresponding to the second metric in the allocated memory for the instance, and resetting the second flag in step (h) includes erasing the stored value corresponding to the second flag stored in the allocated memory.

- 18. The computer-readable medium claim 13, wherein the steps (a)-(h) are performed in real-time.

- 19. The computer-readable medium claim 13, wherein the method further comprising:

- (j) monitoring, by the monitoring server, a value of the first performance metric, wherein the value indicates operation status of the object instance;

and wherein creating an alert in step (g) further comprises creating an alert when the value indicates operational flaw for the object instance.

- 20. A computer implemented method for detecting operational flaws in a network comprising of network elements, the method comprising:

- (a) receiving a stream of metrics corresponding to object instances of objects associated with a network element, the metrics including (i) a plurality of performance metrics each indicating a value for at least one of the object instances and (ii) a control metric for the network element;
- (b) tracking, by a monitoring server, performance metrics for a plurality of monitored object instances;
- (c) on receipt of the control metric for the network element, determining, by the monitoring server, whether a performance metric for each of the plurality of monitored object instances has been received since the last receipt a prior control metric for the object instance; and
- (d) when a performance metric for a monitored object instances is determined not to have been received in (c), creating an alert indicating a potential flaw with the network element.

21. The method of claim 1, wherein the tracking (b) comprises, on receipt of respective performance metrics, setting a flag corresponding to an object instance of the performance metric,

wherein the determining (c) comprises determining whether flags for each of the plurality of monitored object instances have been raised, and further comprising:

- (e) on receipt of the control metric for the network element, clearing flags for each of the plurality of monitored object instances.

100

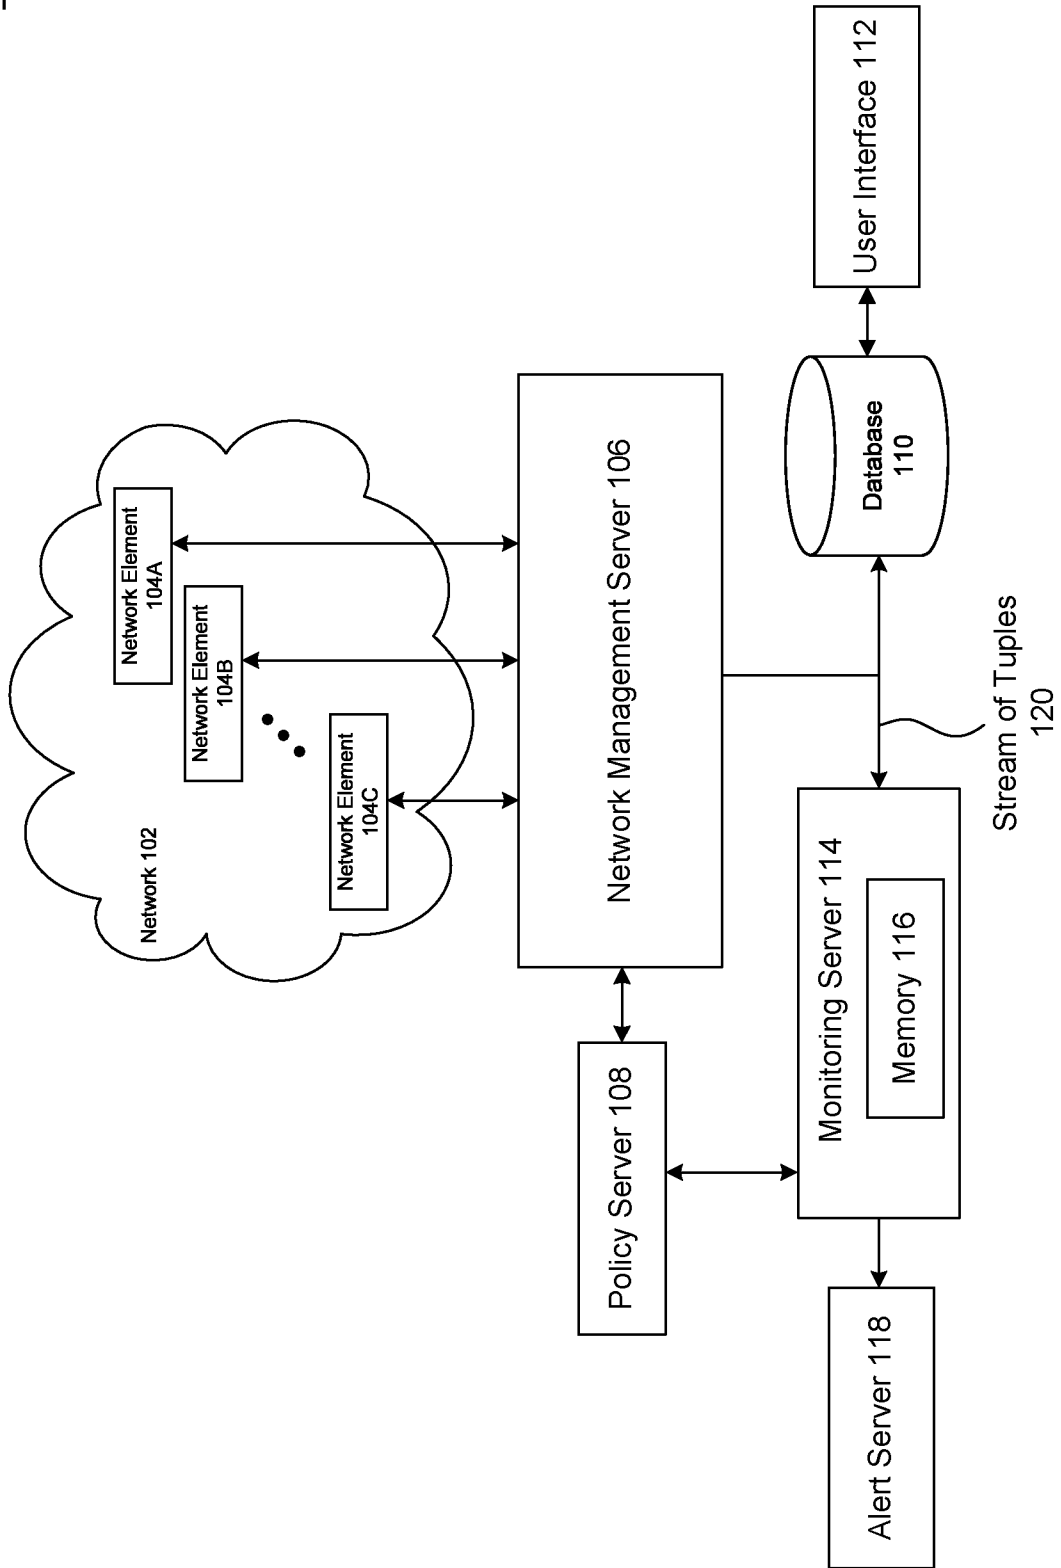


FIG. 1

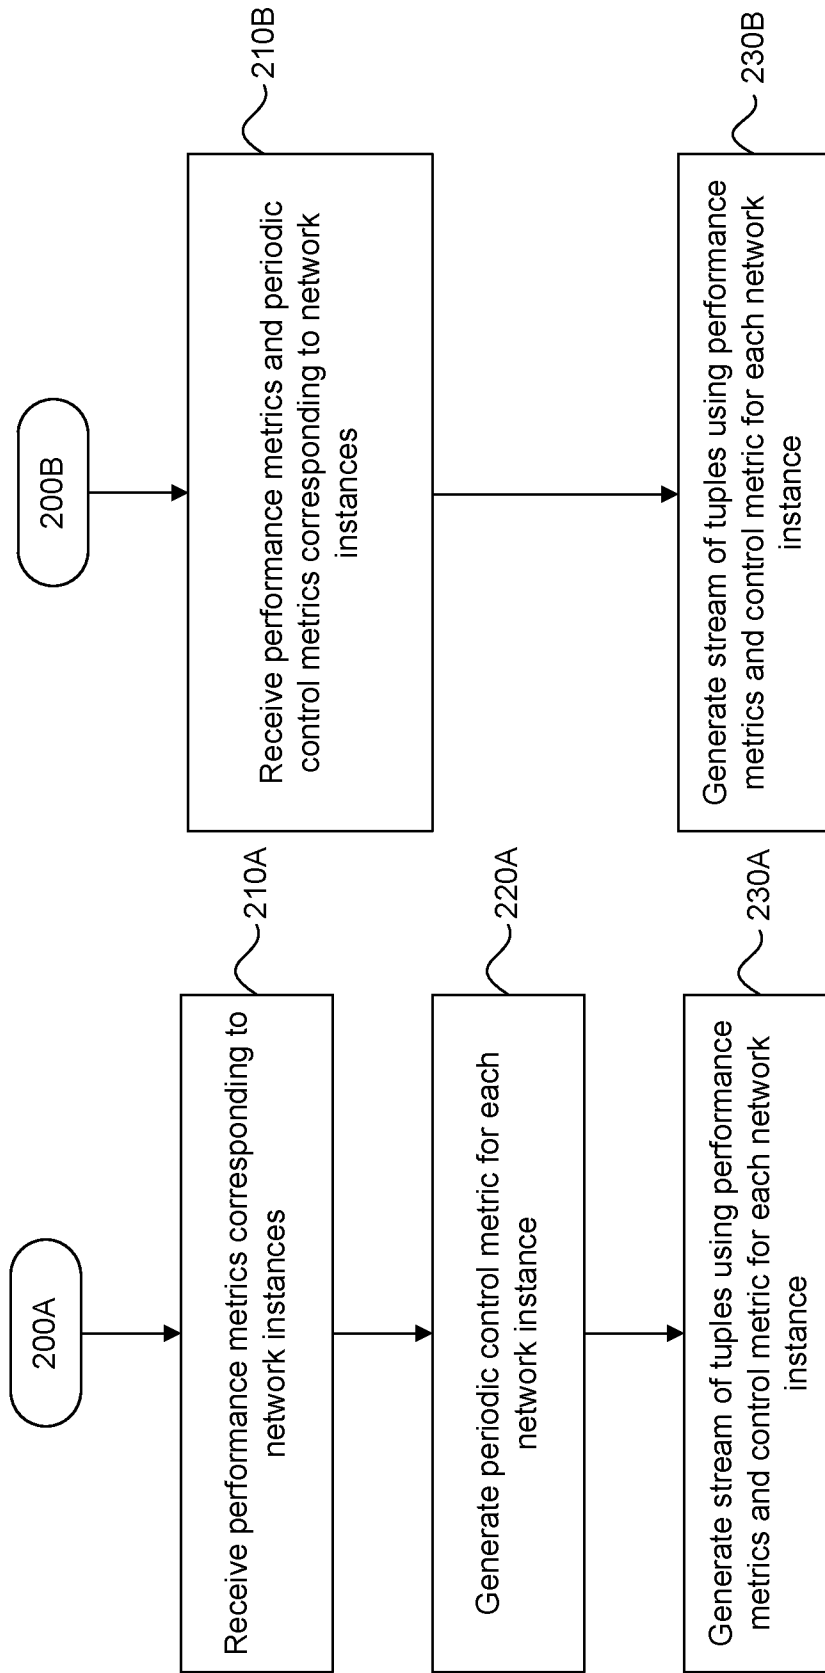


FIG. 2A

FIG. 2B

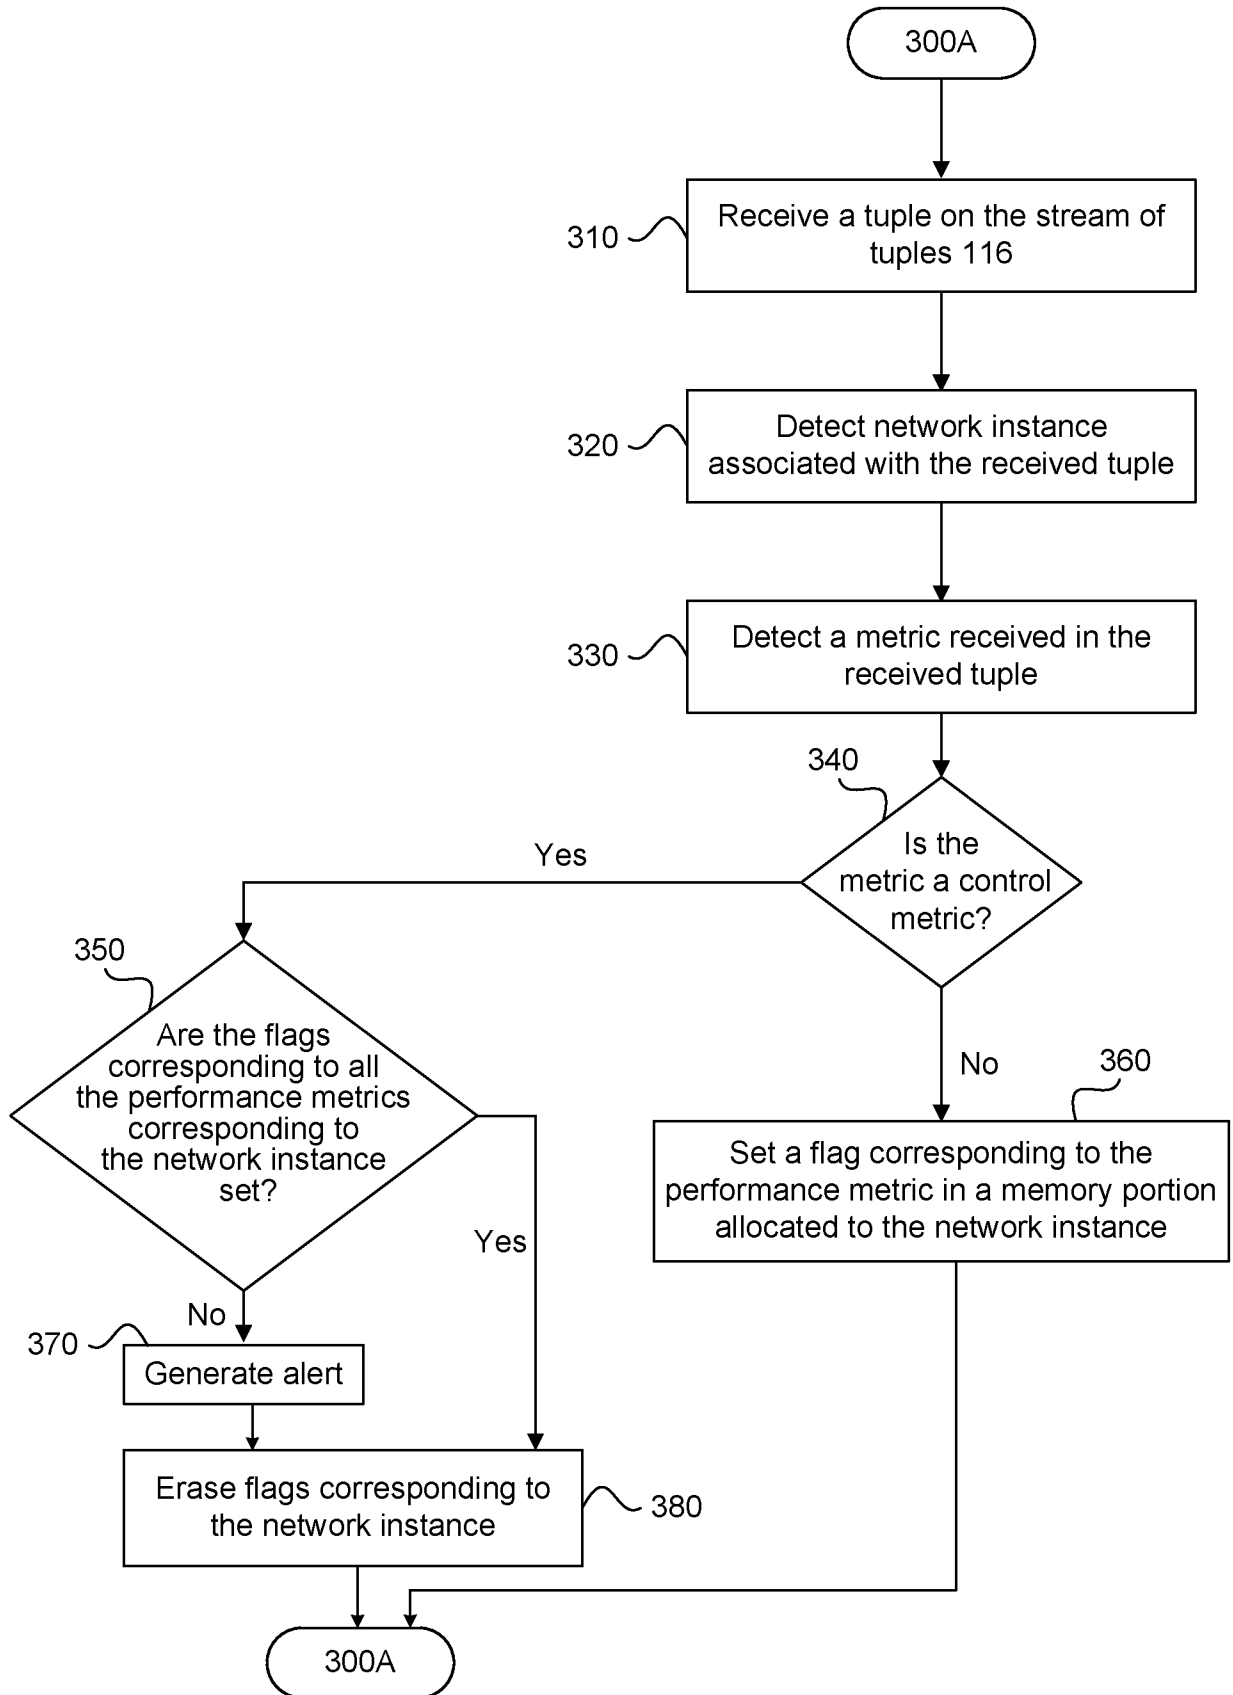


FIG. 3A

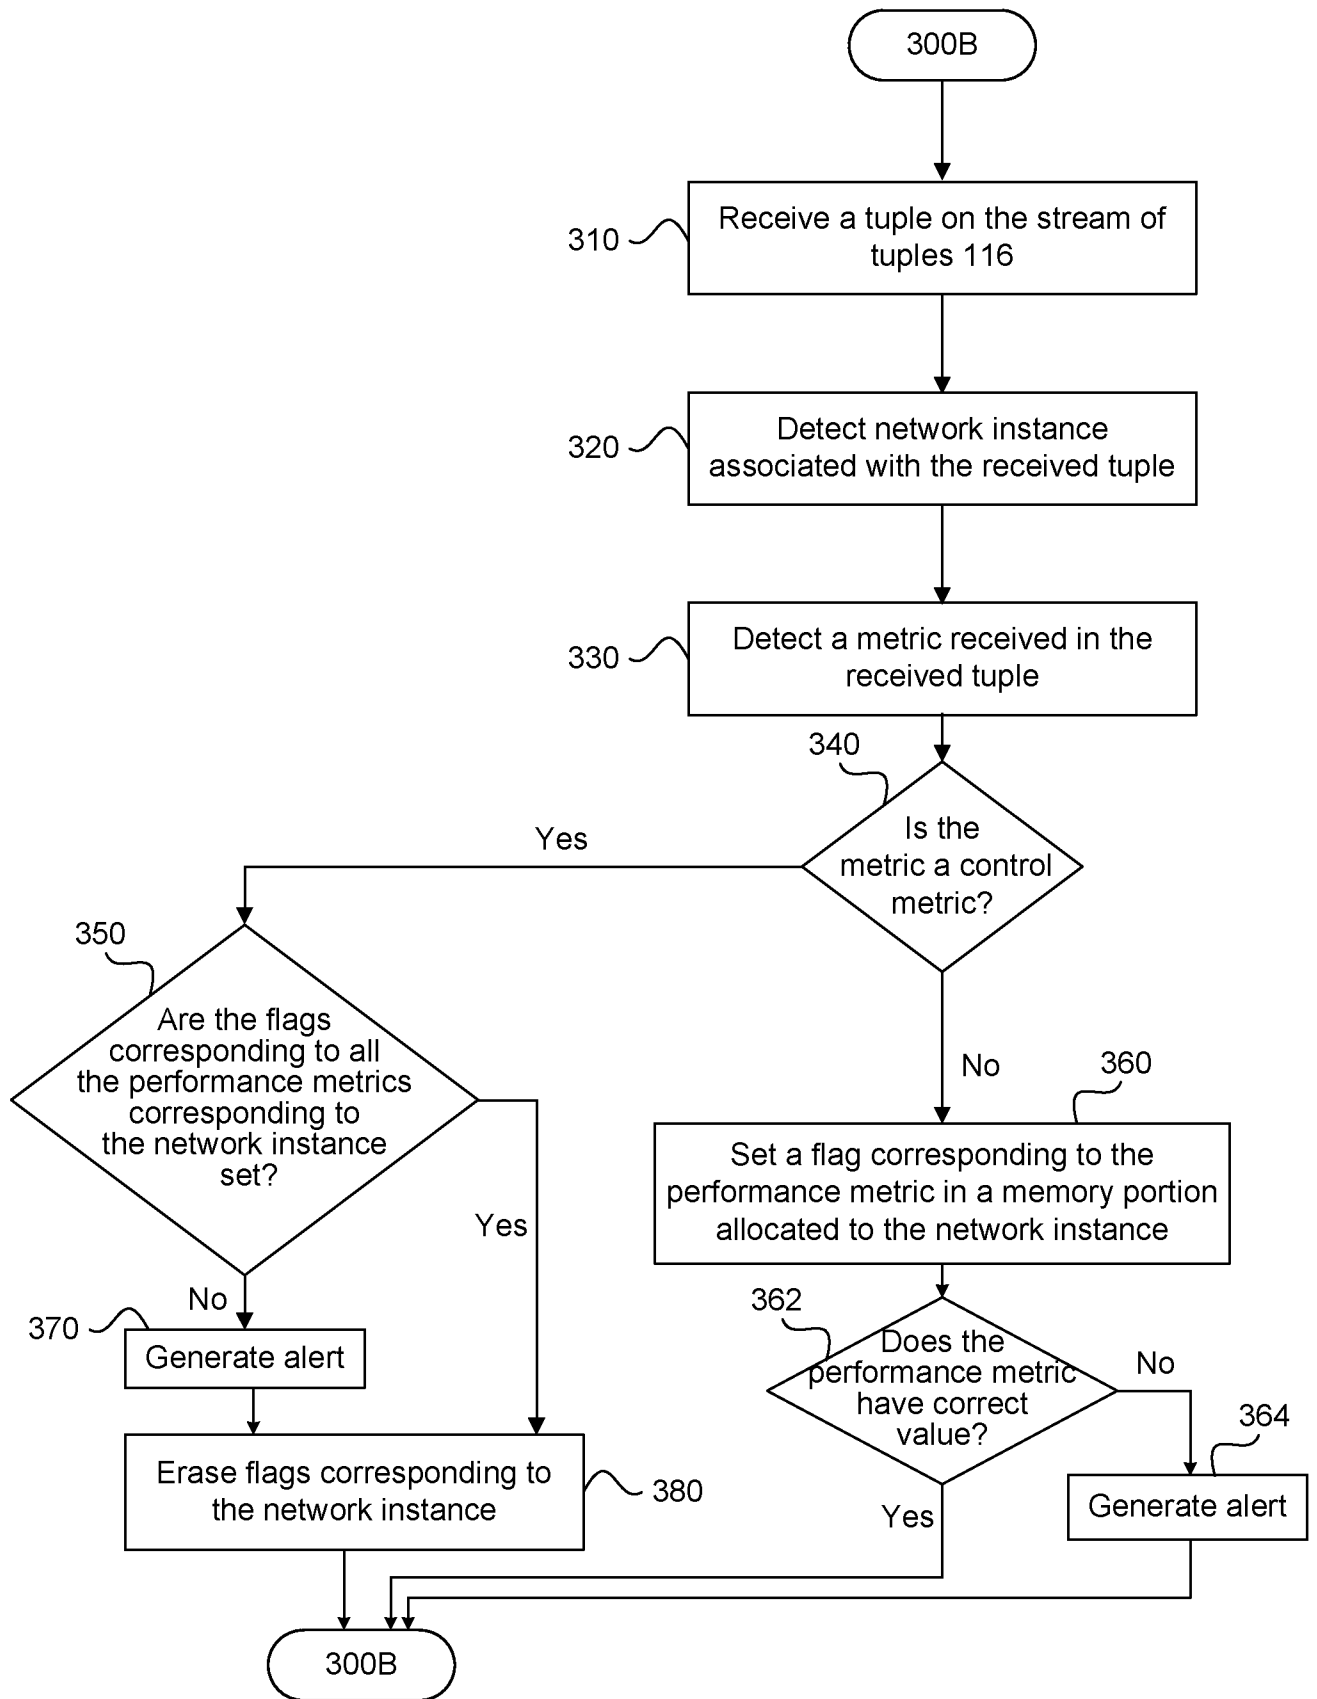



FIG. 3B

400



**FIG. 4**

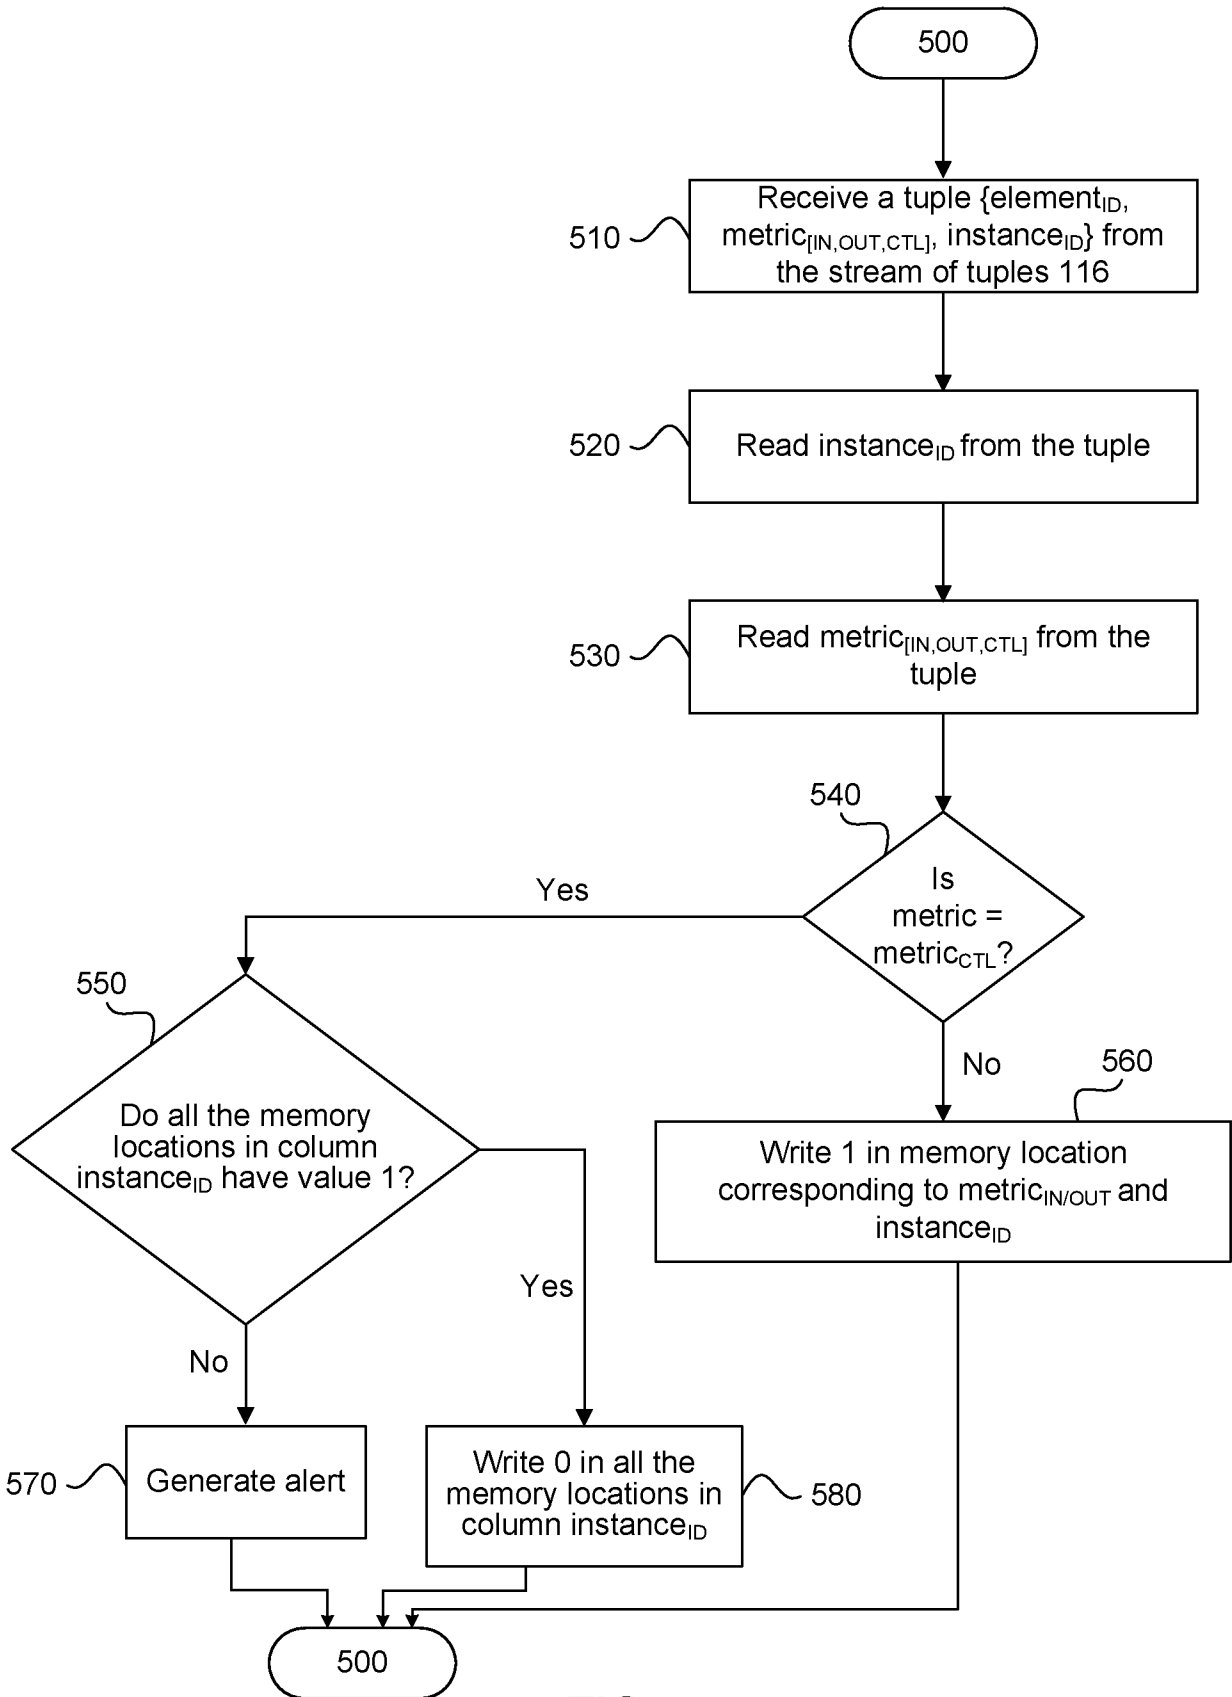


FIG. 5

500



**FIG. 6**

**INTERNATIONAL SEARCH REPORT**

International application No.

PCT/US2016/036025

**A. CLASSIFICATION OF SUBJECT MATTER**  
 IPC(8) - G06F 15/173; H04L 12/26; H04L 12/28 (2016.01)  
 CPC - H04L 29/06; H04L 29/08072; H04L 43/00 (2016.05)  
 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)

IPC - G06F 15/173; H04L 12/26; H04L 12/28  
 CPC - H04L 29/06; H04L 29/08072; H04L 43/00

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
 USPC - 370/400; 709/223; 709/224 (keyword delimited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Orbit, Google Patents, Google  
 Search terms used: network, performance, metric, tuple, instance, server, flag, alert

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2013/0117847 A1 (FRIEDMAN et al) 09 May 2013 (09.05.2013) entire document	1-21
Y	US 2003/0105976 A1 (COPELAND III) 05 June 2003 (05.06.2003) entire document	1-7, 9-11, 13-19, 21
Y	US 2009/0138945 A1 (SAVCHUK) 28 May 2009 (28.05.2009) entire document	8-12
Y	US 6,708,137 B2 (CARLEY) 16 March 2004 (16.03.2004) entire document	20

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

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"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
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search  
 03 August 2016

Date of mailing of the international search report  
**26 AUG 2016**

Name and mailing address of the ISA/  
 Mail Stop PCT, Attn: ISA/US, Commissioner for Patents  
 P.O. Box 1450, Alexandria, VA 22313-1450  
 Facsimile No. 571-273-8300

Authorized officer  
 Blaine R. Copenheaver  
 PCT Helpdesk: 571-272-4300  
 PCT OSP: 571-272-7774