



- (51) International Patent Classification: Not classified
- (21) International Application Number: PCT/CN2014/081910
- (22) International Filing Date: 9 July 2014 (09.07.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant: **SPLUNK INC.** [US/US]; 250 Brannan Street 1st Floor, San Francisco, California 94107 (US).
- (72) Inventors; and
- (71) Applicants (for US only): **ZHONG, Qianjie** [CN/CN]; Baocheng Road, No. 155, Building 12, Room 801, Minhang District, Shanghai 201100 (CN). **WANG, Ting** [CN/CN]; Rm 502, No. 78, Lane 1389, Dong Qiu Rd., Pudong District, Shanghai 200126 (CN). **LEE, Margaret** [US/US]; 5136 Newington Court, Dublin, California 94568 (US). **LI, Dawei** [CN/CN]; Rm 201, Bldg A7, No 1 Longdong Rd., Shanghai 201203 (CN). **FILIPPI, Nick** [US/US]; 63 Stockbridge Avenue, Atherton, California 94027 (US). **NI, Yue** [CN/CN]; Rm 302, No 10, Lane 415, Pingtang Road, Shanghai 200335 (CN). **YUAN, Shiming** [CN/CN]; Rm 305, Building 3, Shilansancun, No 250 Langao Rd., Putuo District, Shanghai 200061 (CN).
- (74) Agent: **CHINA SINDA INTELLECTUAL PROPERTY LTD.**; B11th Floor, Focus Place, 19 Financial Street, Xicheng District, Beijing 100033 (CN).

- (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, LT, 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, 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:**

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: MANAGING DATASETS PRODUCED BY ALERT-TRIGGERING SEARCH QUERIES

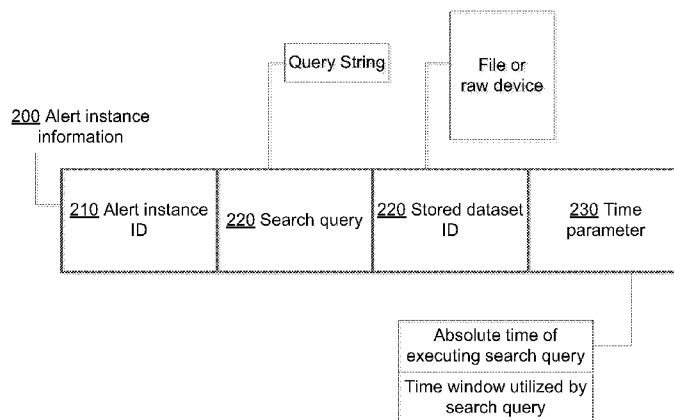


FIG. 2

(57) Abstract: Systems and methods for managing datasets produced by alert-triggering search queries in data aggregation and analysis systems. An example method may comprise: executing, by one or more processing devices, a search query on a portion of searchable data associated with a time window to produce a dataset comprising one or more results; responsive to determining that at least a portion of the dataset satisfies a triggering condition defining an alert associated with the search query, generating an instance of the alert; associating, by a memory data structure, the instance of the alert with an identifier of the search query and a time parameter specifying the time window; receiving, from a client computing device, a request for the portion of the dataset; and responsive to determining that the portion of the dataset is not stored in the memory in a manner associating it with the instance of the alert, re-producing the portion of the dataset by re-executing the search query in view of the time parameter.

WO 2016/004594 A2

**MANAGING DATASETS PRODUCED BY ALERT-TRIGGERING SEARCH QUERIES****TECHNICAL FIELD**

**[0001]** The present disclosure is generally related to data aggregation and analysis systems, and is more specifically related to alerting mechanisms employed by data aggregation and analysis systems.

**BACKGROUND**

**[0002]** Modern data centers often comprise thousands of hosts that operate collectively to service requests from even larger numbers of remote clients. During operation, components of these data centers can produce significant volumes of machine-generated data. The unstructured nature of much of this data has made it challenging to perform indexing and searching operations because of the difficulty of applying semantic meaning to unstructured data. As the number of hosts and clients associated with a data center continues to grow, processing large volumes of machine-generated data in an intelligent manner and effectively presenting the results of such processing continues to be a priority.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The present disclosure is illustrated by way of examples, and not by way of limitation, and may be more fully understood with references to the following detailed description when considered in connection with the figures, in which:

[0004] FIG. 1 schematically illustrates an example data structure that may be utilized by the systems and methods described herein for defining alerts in accordance with one or more aspects of the present disclosure.

[0005] FIG. 2 schematically illustrates an example data structure that may be utilized by the systems and methods described herein for storing the search query, dataset, and a time parameter associated with an alert instance, in accordance with one or more aspects of the present disclosure.

[0006] FIGs. 3A-3C depict flow diagrams of example methods 300A-300C for managing datasets produced by alert-triggering search queries, in accordance with one or more aspects of the present disclosure.

[0007] FIG. 4 presents a block diagram of an event-processing system in accordance with the disclosed embodiments.

[0008] FIG. 5 presents a flowchart illustrating how indexers process, index, and store data received from forwarders in accordance with the disclosed embodiments.

[0009] FIG. 6 presents a flowchart illustrating how a search head and indexers perform a search query in accordance with the disclosed embodiments.

**[00010]** FIG. 7 presents a block diagram of a system for processing search requests that uses extraction rules for field values in accordance with the disclosed embodiments.

**[00011]** FIG. 8 illustrates an exemplary search query received from a client and is executed by search peers in accordance with the disclosed embodiments.

**[00012]** FIG. 9 depicts a block diagram of an example computing device operating in accordance with one or more aspects of the present disclosure.

## DETAILED DESCRIPTION

[00013] Disclosed herein are systems and methods for managing datasets produced by alert-triggering search queries in data aggregation and analysis systems.

[00014] An example data aggregation and analysis system may aggregate heterogeneous machine-generated data received from various sources, including servers, databases, applications, networks, etc. The aggregated source data may comprise a plurality of events. An event may be represented by a data structure that is associated with a certain point in time and comprises a portion of raw machine-generated data. The system may be configured to perform real-time indexing of the source data and to execute real-time, scheduled, or historic searches on the source data.

[00015] The data aggregation and analysis system may provide an alerting mechanism to trigger certain actions in response to an occurrence of certain conditions. An alert may be defined by a triggering condition applied to a data set produced by a search query that is executed by the system either in real time or according to a certain schedule. An alert instance may be triggered when at least a portion of the dataset returned by the search satisfies the triggering condition, as described in more details herein below.

[00016] Responsive to triggering an alert instance, the system may perform one or more actions associated with the alert, such as sending an email message to certain recipients, creating a Really Simple Syndication (RSS) feed, executing a certain script, and/or presenting the alert instance via a graphical user interface (GUI) of a client computing device. In various illustrative

examples, the system may present alert instances on a client computing device (*e.g.*, a desktop computing device or a mobile computing device).

**[00017]** In certain implementations, responsive to triggering an alert instance, the system may store, in a storage memory, a portion of the dataset that has triggered the alert instance. The system may associate (*e.g.*, by a memory data structure) the instance of the alert with an identifier of the portion of the dataset stored in the memory.

**[00018]** The system may allow a user to “drill down” to the underlying data that has triggered a particular alert instance: a user may request (*e.g.*, via a GUI associated with the client computing device) the portion of the dataset that has triggered the alert instance. Responsive to receiving the request, the system may retrieve, from the storage memory, the portion of the dataset associated with the alert instance and present it to the requestor.

**[00019]** As the disk space allocated for the files storing the datasets produced by alert-triggering search queries may be limited, the system may implement a file retention policy with respect to such files. In an illustrative example, a file may be deleted upon expiration of a certain time since the time of the file creation. Alternatively, one or more least recently created, accessed or modified files may be deleted responsive to determining that the free disk space falls below a certain threshold value.

**[00020]** Thus, requesting the portion of the dataset associated with a particular alert instance may trigger an irrecoverable error condition if the portion of the dataset has been deleted in accordance with a file retention policy. To provide a recovery mechanism in the above described scenario, an example data aggregation and analysis system may, at the time when an alert is

triggered, store, in a memory, the search query associated with the alert and a time parameter associated with the instance of the alert. The time parameter may define a point in time (*e.g.*, using POSIX time represented by number of seconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970) of executing a search query that has triggered the alert instance, and/or an absolute time range associated with the search query that has triggered the alert instance. Alternatively, the time parameter may define the time range of data searched using the query that resulted in meeting the trigger condition, where the data is time-series data.

**[00021]** Later, when a user requests the alert triggering dataset, the system may determine whether the portion of the dataset is stored in the storage memory. Responsive to successfully retrieving, from the storage memory, the portion of the dataset identified by the dataset identifier associated with the alert instance, the system may return the dataset to the requestor. Otherwise, responsive to determining that the portion of the dataset is not stored in the storage memory, the system may, transparently for the requestor, re-execute the search query using the stored alert time or the time window, and return the relevant portion of the dataset to the requestor.

**[00022]** In certain implementations, a dataset that has been recreated by the above described procedure responsive to a client request, may be stored in the storage memory subject to the general or a shortened file retention policy.

**[00023]** Various aspects of the methods and systems are described herein by way of examples, rather than by way of limitation. The methods described herein may be implemented by hardware (*e.g.*, general purpose and/or specialized processing devices, and/or other devices and

associated circuitry), software (*e.g.*, instructions executable by a processing device), or a combination thereof.

**[00024]** As noted herein above, an example data aggregation and analysis system may be configured to perform real-time indexing of the source data and to execute real-time, scheduled, or historic searches on the source data. A search query may comprise one or more search terms. Search terms may include keywords, phrases, Boolean expressions, regular expressions, field name-value pairs, *etc.* In certain implementations, the data aggregation and analysis system may produce a plurality of data items from the source data in a reverse chronological order, in order to optimize search operations on more recent data. The data aggregation and analysis system may apply a filter specifying relative (with respect to a certain point in time) or absolute (*e.g.*, POSIX) time values, to filter the dataset produced by the search by a specific time value or a specific time window.

**[00025]** The data aggregation and analysis system may provide an alerting mechanism to trigger certain actions in response to occurrences of certain conditions. An alert may be associated with a search query that is executed by the system either in real time or according to a certain schedule. An alert instance may be triggered when at least a portion of the resulting dataset produced by the search meet a triggering condition defining the alert, as described in more details herein below.

**[00026]** In certain implementations, alerts may be classified into several categories based on the search execution mode and/or the type of the triggering condition associate with an alert. In an illustrative example, alerts may be categorized as “per result” alerts, scheduled alerts, or rolling window alerts.

**[00027]** A “per result” alert may be based on a real-time search (*e.g.*, repeatedly executed by a certain process or thread in an indefinite loop which may be interrupted by occurrences of certain terminating conditions). A “per result” alert instance may be triggered whenever the search associated with the alert returns a data item (*e.g.*, an event, or a data item of another type). In other words, an alert instance is triggered for each result returned by the search query. Thus, evaluating a triggering condition for a “per result” alert may comprise determining that the dataset produced by executing the search query associated with the alert comprises at least one data item. “Per result” alert type may be utilized for triggering one or more actions to be performed within a short period of time of a matching search result being detected. In an illustrative example, a “per result” alert instance may be triggered responsive to detecting every failed login attempt. In another illustrative example, a “per result” alert instance may be triggered responsive to detecting a “file system full” error on any host.

**[00028]** A “scheduled” alert may be based on a search that runs according to a certain configurable schedule. A configurable schedule may comprise one or more values of minute, hour, day, month, and/or day-of-week. A “scheduled” alert instance may be triggered whenever the dataset produced by executing the search query associated with the alert satisfies a certain triggering condition. The search schedule may include criteria for searching data only within a specified time period relative to when the search is performed. Thus, the same search query may produce different results when run (according to the schedule) at different times.

**[00029]** In an illustrative example, evaluating a triggering condition may comprise comparing the number of data items in the dataset produced by executing the search query associated with the alert to a certain configurable integer value. In an illustrative example, an alert instance may

be triggered whenever the number of data items in the dataset produced by executing the search query associated with the alert is greater than, less than, or equal to a certain configurable integer value.

**[00030]** In another illustrative example, evaluating a triggering condition may comprise performing a secondary search on the dataset produced by executing the search query associated with the alert. If a triggering condition is utilized for defining a “scheduled” alert, an alert instance is triggered whenever the secondary search returns any results.

**[00031]** In an illustrative example, a “scheduled” alert instance may be triggered responsive to determining that the number of “404 – Object not found” errors in any given one hour time interval exceeds 100. In another illustrative example, a “scheduled” alert instance may be triggered responsive to determining that the number of items sold in the previous day is less than 500. “Scheduled” alert instances may be throttled as described in more details herein below.

**[00032]** A “rolling window” alert may be based on a real-time search (*e.g.*, repeatedly executed by a certain process or thread in an indefinite loop which may be interrupted by occurrences of certain terminating conditions) that monitors data (such as events) within a rolling time window. The rolling time window may be defined relative to the “current time” moment, *e.g.*, between “current time” and 15 minutes to “current time.” A “rolling window” alert instance may be triggered whenever the search criteria and the triggering condition are met when applied to the data within a rolling window.

**[00033]** In an illustrative example, evaluating a triggering condition may comprise comparing the number of data items in the dataset produced by executing the search query associated with

the alert to a certain configurable integer value. In an illustrative example, an alert instance may be triggered whenever the number of data items in the dataset produced by executing the real-time search query that monitors events within a rolling time window associated with the alert is greater than, less than, or equal to a certain configurable integer value.

**[00034]** In another illustrative example, evaluating a triggering condition may comprise performing a secondary search on the dataset produced by executing the real-time search query that monitors events within a rolling time window associated with the alert. An alert instance may be triggered whenever the secondary search returns any results.

**[00035]** The “rolling window” alert type may be utilized for triggering one or more actions responsive to events occurring within a rolling time window of a configurable width. In an illustrative example, a “rolling window” alert instance may be triggered responsive to detecting three consecutive login attempts by a user within ten minutes immediately before the current time. In another illustrative example, a “rolling window” alert instance may be triggered responsive to determining that a host could not complete a file transfer operation within the hour immediately before the current time. “Rolling window” alert instances may be throttled as described in more details herein below.

**[00036]** An alert definition may comprise one or more actions to be performed by the system responsive to triggering an instance of the alert. In an illustrative example, an action associated with an alert may comprise sending an e-mail message to certain recipients. The system may insert into the e-mail message at least a portion of the dataset that has triggered the alert instance. In another illustrative example, an action associated with an alert may comprise creating an RSS feed including at least a portion of the dataset that has triggered the alert instance. In another

illustrative example, an action associated with an alert may comprise executing a shell script that can perform various actions (*e.g.*, sending a Simple Network Management Protocol (SNMP) trap notification or calling an Application Programming Interface (API) function). In another illustrative example, an action associated with an alert may comprise presenting the alert via a GUI of a client computing device.

**[00037]** In certain implementations, an alert definition may comprise a rule specifying whether the actions associated with the alert should be performed for each data item in the dataset produced by executing the search query associated with the alert, or once for all data items comprised by the dataset. In an illustrative example, a search query may be configured to return user names identifying the users that had over three consecutive failed login attempts within a specified period of time. A definition of an alert associated with the search query may specify that the action associated with the alert (*e.g.*, executing a script to disable the affected user account) should be performed for each data item in the dataset produced by executing the search query. Alternatively, the alert definition may specify that the action associated with the alert (*e.g.*, sending an e-mail message to a system administrator) should be performed for all data items in the dataset produced by executing the search query.

**[00038]** In certain implementations, triggering an alert instance may be suppressed by a throttling condition. In an illustrative example, an alert definition may comprise a throttling condition to suppress triggering alert instances for a certain period of time following a previous instance of the alert. Alternatively, an alert definition may comprise a throttling condition to suppress, for a certain period of time, triggering alert instances for certain data items (*e.g.*, defined by field name-value pairs) in the dataset produced by executing the search query. For

example, a throttling condition may be defined to suppress for one hour triggering alert instances for failed login attempts by the same user.

[00039] FIG. 1 schematically illustrates an example data structure that may be utilized by the systems and methods described herein for defining alerts in accordance with one or more aspects of the present disclosure. As schematically illustrated by FIG. 1, an alert definition 100 may comprise an alert type 110, a search query 120, a schedule 130, a triggering condition 140, a throttling condition 150, and one or more actions 160. It should be noted that an alert definition 100 may include more, less or different alert characteristics than those shown in FIG. 1.

[00040] In various illustrative examples, alert type 110 may be represented by “per result,” “scheduled,” or “rolling window” alerts, as described in more details herein above.

[00041] Search query 120 may be represented by a search query string comprising one or more search terms. The search query string may satisfy to the syntax of a certain query language supported by the data aggregation and retrieval systems.

[00042] Schedule 130 may be represented by a data structure comprising values of one or more scheduling parameters (*e.g.*, minute, hour, day, month, and/or day-of-week).

[00043] Triggering condition 140 may comprise comparing the number of data items in the dataset produced by executing the search query associated with the alert to a certain configurable integer value. In another illustrative example, triggering condition 140 may comprise a secondary search on the dataset produced by executing the search query associated with the alert, as described in more details herein above.

**[00044]** Throttling condition 150 may be utilized to suppress triggering certain alert instances. In an illustrative example, a throttling condition may be used to suppress triggering alert instances for a certain period of time following a previous instance of the alert. Alternatively, an alert definition may comprise a throttling condition to suppress, for a certain period of time, triggering alert instances for certain data items in the dataset produced by executing the search query, as described in more details herein above.

**[00045]** Action 160 may specify one or more actions to be performed by the system responsive to triggering an instance of the alert. The actions may include, for example, sending an e-mail message to certain recipients, creating an RSS feed including at least part of the dataset that has triggered the alert instance, executing a shell script, and/or presenting the alert via a GUI of a client computing device, as described in more details herein above.

**[00046]** In certain implementations, responsive to triggering an alert instance, the system may store, in a storage memory, a portion of the dataset that has triggered the alert instance, the search query associated with the alert, and a time parameter associated with the instance of the alert (e.g., the time parameter specifying what data was searched that resulted in triggering of the alert instance, where the data is time-stamped). The time parameter may define an absolute (e.g., POSIX) time of executing a search query associated with a real-time or a scheduled alert, or a time window for a rolling window alert. For a given alert instance, the system may also store the trigger condition, where this is needed to reproduce the results that resulted in the alert instance. In various illustrative examples, the storage memory may be provided by a file system or by a raw device-based storage.

**[00047]** The system may associate, *e.g.*, by a memory data structure, an identifier of the alert instance with the portion of the dataset stored in the memory, the search query, and the time parameter. FIG. 2 schematically illustrates an example data structure that may be utilized by the systems and methods described herein for storing the search query, dataset, and a time parameter associated with an alert instance, in accordance with one or more aspects of the present disclosure. As schematically illustrated by FIG. 2, memory data structure 200 may map an alert instance identifier 210 to an identifier 220 of the portion of the dataset stored in the memory, the search query 230 utilized to produce the dataset, and the time parameter 240 associated with the alert instance.

**[00048]** In certain implementations, the system may implement a file retention policy with respect to various files, including the files containing datasets associated with alert instances. In an illustrative example, a file may be deleted upon expiration of a certain time (*e.g.*, 24 hours) since the time of the file creation. Alternatively, one or more least recently created, accessed or modified files may be deleted responsive to determining that the free disk space falls below a certain threshold value.

**[00049]** At any time after an alert instance has been triggered, a user may request (*e.g.*, via a GUI associated with the client computing device), the portion of the dataset that has triggered the alert instance. Responsive to receiving the request, the system may determine whether the portion of the dataset is stored in the storage memory. In an illustrative example, the system may determine an identifier of the portion of the dataset by looking up the identifier of the alert instance in the memory data structure 200. Responsive to determining that the portion of the dataset is stored in the storage memory, the system may retrieve the portion of the dataset from

the storage memory and transmit a message comprising the portion of the dataset to the client computing device. Otherwise, responsive to determining that the portion of the dataset is not stored in the storage memory, the system may identify the search query and the time parameter associated with the alert instance, by looking up the identifier of the alert instance in the memory data structure 200. The system may then re-create the portion of the dataset by executing the search query using the time parameter associated with the alert instance. In certain implementations, the time parameter may define a point in time (*e.g.*, using POSIX time represented by number of seconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970) of executing a search query that has triggered the alert instance. The time parameter may also define a time range associated with the search query that has triggered the alert instance.

**[00050]** In an illustrative example, for a rolling window alert, the system may substitute, in the definition of the time range utilized by the search query, the “current time” moment with the value of the time parameter associated with the alert instance, where the time parameter represents the point in time of executing the search query that has triggered the alert instance. For example, if a query associated with a rolling window alert defined on a one-hour rolling window was executed at 09:00 UTC, January 1, 2014, the system may re-create the portion of the dataset that has triggered the alert instance, by re-executing the search query with a filtering condition to only select data items having timestamps falling within the time range between 08:00 and 09:00 UTC, January 1, 2014.

**[00051]** In another illustrative example, for a scheduled alert, the system may substitute the definition of the time range utilized by the search query with the value of the time parameter

associated with the alert instance, where the time parameter represents the time range associated with the alert instance, as described in more details herein above. For example, the system may re-create the portion of the dataset that has triggered the alert instance, by re-executing the search query with a filtering condition to only select data items having timestamps falling within the time range between 12:00 and 13:00 UTC, January 1, 2014.

**[00052]** Responsive to re-executing the search query, the system may identify a portion of the resulting dataset that satisfies the triggering condition and transmit a message comprising the portion of the dataset to the client computing device.

**[00053]** In certain implementations, a dataset that has been recreated by the above described procedure responsive to a client request, may be stored in the storage memory subject to the general or a shortened file retention policy.

**[00054]** FIGs. 3A-3C depict flow diagrams of example methods 300A-300C for managing datasets produced by alert-triggering search queries. Methods 300A-300C and/or each of their respective individual functions, routines, subroutines, or operations may be performed by one or more general purpose and/or specialized processing devices. Two or more functions, routines, subroutines, or operations of methods 300A-300C may be performed in parallel or in an order that may differ from the order described above. In certain implementations, one or more of methods 300A-300C may be performed by a single processing thread. Alternatively, methods 300A-300C may be performed by two or more processing threads, each thread executing one or more individual functions, routines, subroutines, or operations of the respective method. In an illustrative example, the processing threads implementing methods 300A-300C may be synchronized (*e.g.*, using semaphores, critical sections, and/or other thread synchronization

mechanisms). Alternatively, the processing threads implementing methods 300A-300C may be executed asynchronously with respect to each other. In an illustrative example, methods 300A-300C may be performed by an example computing device 1000 described herein below with references to FIG. 9. In another illustrative example, methods 300A-300C may be performed by a distributed computer system comprising two or more example computing devices 1000.

**[00055]** FIG. 3A depict a flow diagram of an example method 300A for storing search queries and time parameters associated with alert instances, in accordance with one or more aspects of the present disclosure.

**[00056]** Referring to FIG. 3A, at block 310, the computer system implementing the method may execute a search query. In an illustrative example, the search query may represent a real-time search (*e.g.*, may repeatedly be executed by a certain process or thread in an indefinite loop which may be interrupted by occurrences of certain terminating conditions). In an illustrative example, the search query may represent a scheduled search (*e.g.*, may be executed according to a certain schedule), as described in more details herein above.

**[00057]** Responsive to determining, at block 315, that a portion of the dataset produced by the search query satisfies a triggering condition defining an alert associated with the search query, the processing may continue at block 320; otherwise, the processing associated with the current search query instance may terminate.

**[00058]** At block 320, the computer system may generate an alert instance based on at least a portion of the dataset produced by executing the search query.

[00059] At block 325, the computer system may store, in a memory associated with the computer system, the portion of the dataset. In various illustrative examples, the storage memory may be provided by a file system or by a raw device-based storage.

[00060] At block 330, the computer system may associate, by a memory data structure, the instance of the alert with an identifier of the portion of the dataset, an identifier of the search query, and a time parameter. The time parameter may define a point in time (*e.g.*, using POSIX time) of executing a search query that has triggered the alert instance, and/or a time range associated with the search query that has triggered the alert instance, as described in more details herein above.

[00061] The stored portion of the dataset may be subject to a file retention policy implemented by the computer system, *e.g.*, as described herein below with references to FIG. 3B.

[00062] FIG. 3B depicts a flow diagram of an example method 300B for implementing a file retention policy with respect to stored datasets, in accordance with one or more aspects of the present disclosure.

[00063] Referring to FIG. 3B, at block 335, the computer system implementing the method may identify a file storing a dataset (*e.g.*, by cycling through a plurality of stored datasets).

[00064] At block 340, the computer system may evaluate a file retention condition applied to the identified dataset.

[00065] Responsive to determining, at block 340, that the file retention condition for the identified file is satisfied, the method may loop back to block 335. Otherwise, the computer system may, at block 345, delete the identified dataset. In an illustrative example, a file retention

policy may require deleting a file upon expiration of a certain time period (*e.g.*, 24 hours) since the time of the file creation, and the computer system can check at block 340 whether the time of the file creation is still within the time period defined by the file retention policy. Alternatively, the file retention policy may require deleting one or more (*e.g.*, a predefined number of) least recently created, accessed or modified files responsive to determining that the free disk space falls below a certain threshold value, and the computer system can check at block 340 whether the file is not one of the least recently created, accessed or modified files.

**[00066]** As noted herein above, in certain implementations, the file retention policy may be implemented by a dedicated processing device and/or dedicated processing thread operating asynchronously with respect to other operations comprised by method 300.

**[00067]** FIG. 3C depicts a flow diagram of an example method 300C for processing client requests for datasets associated with alert instances, in accordance with one or more aspects of the present disclosure.

**[00068]** Referring to FIG. 3C, at block 350, the computer system implementing the method may receive, from a client computing device, a request for the portion of the dataset associated with a certain alert instance identified by the request.

**[00069]** Responsive to determining, at block 355, that the portion of the dataset is stored in the storage memory, the processing may continue at block 365; otherwise, the method may branch to block 360.

**[00070]** At block 360, the computer system may identify, using the above referenced memory data structure, the search query utilized for producing the portion of the dataset that has triggered the alert instance identified by the client request, and the time parameter associated with the

search query. The computer system may produce a copy of the portion of the dataset by re-executing the search query in view of the time parameter.

**[00071]** In an illustrative example, for a rolling window alert, the system may substitute, in the definition of the time range utilized by the search query, the “current time” moment with the value of the time parameter associated with the alert instance, where the time parameter represents the point in time of executing the search query that has triggered the alert instance. In another illustrative example, for a scheduled alert, the system may substitute the definition of the time range utilized by the search query with the value of the time parameter associated with the alert instance, where the time parameter represents the time range associated with the alert instance, as described in more details herein above.

**[00072]** At block 365, the computer system may transmit the copy of the portion of the dataset to the client computing device. Responsive to completing operations referenced by the block 365, processing associated with the current search query instance may terminate.

**[00073]** The systems and methods described herein above may be employed by various data processing systems, *e.g.*, data aggregation and analysis systems. In various illustrative examples, the data processing system may be represented by the SPLUNK® ENTERPRISE system produced by Splunk Inc. of San Francisco, California, to store and process performance data.

**[00074]** Modern data centers often comprise thousands of host computer systems that operate collectively to service requests from even larger numbers of remote clients. During operation, these data centers generate significant volumes of performance data and diagnostic information that can be analyzed to quickly diagnose performance problems. In order to reduce the size of this performance data, the data is typically pre-processed prior to being stored based on

anticipated data-analysis needs. For example, pre-specified data items can be extracted from the performance data and stored in a database to facilitate efficient retrieval and analysis at search time. However, the rest of the performance data is not saved and is essentially discarded during pre-processing. As storage capacity becomes progressively cheaper and more plentiful, there are fewer incentives to discard this performance data and many reasons to keep it.

**[00075]** This plentiful storage capacity is presently making it feasible to store massive quantities of minimally processed performance data at “ingestion time” for later retrieval and analysis at “search time.” Note that performing the analysis operations at search time provides greater flexibility because it enables an analyst to search all of the performance data, instead of searching pre-specified data items that were stored at ingestion time. This enables the analyst to investigate different aspects of the performance data instead of being confined to the pre-specified set of data items that were selected at ingestion time.

**[00076]** However, analyzing massive quantities of heterogeneous performance data at search time can be a challenging task. A data center may generate heterogeneous performance data from thousands of different components, which can collectively generate tremendous volumes of performance data that can be time-consuming to analyze. For example, this performance data can include data from system logs, network packet data, sensor data, and data generated by various applications. Also, the unstructured nature of much of this performance data can pose additional challenges because of the difficulty of applying semantic meaning to unstructured data, and the difficulty of indexing and querying unstructured data using traditional database systems.

**[00077]** These challenges can be addressed by using an event-based system, such as the SPLUNK® ENTERPRISE system produced by Splunk Inc. of San Francisco, California, to store and process performance data. The SPLUNK® ENTERPRISE system is the leading platform for providing real-time operational intelligence that enables organizations to collect, index, and harness machine-generated data from various websites, applications, servers, networks, and mobile devices that power their businesses. The SPLUNK® ENTERPRISE system is particularly useful for analyzing unstructured performance data, which is commonly found in system log files. Although many of the techniques described herein are explained with reference to the SPLUNK® ENTERPRISE system, the techniques are also applicable to other types of data server systems.

**[00078]** In the SPLUNK® ENTERPRISE system, performance data is stored as “events,” wherein each event comprises a collection of performance data and/or diagnostic information that is generated by a computer system and is correlated with a specific point in time. Events can be derived from “time series data,” wherein time series data comprises a sequence of data points (e.g., performance measurements from a computer system) that are associated with successive points in time and are typically spaced at uniform time intervals. Events can also be derived from “structured” or “unstructured” data. Structured data has a predefined format, wherein specific data items with specific data formats reside at predefined locations in the data. For example, structured data can include data items stored in database fields or data items stored in fields in a data structure defined by a computer program. In contrast, unstructured data does not have a predefined format. This means that unstructured data can comprise various data items having different data types that can reside at different locations. For example, when the data source is an operating system log, an event can include one or more lines from the operating

system log containing raw data that includes different types of performance and diagnostic information associated with a specific point in time. Examples of data sources from which an event may be derived include, but are not limited to: web servers; application servers; databases; firewalls; routers; operating systems; and software applications that execute on computer systems, mobile devices, and sensors. The data generated by such data sources can be produced in various forms including, for example and without limitation, server log files, activity log files, configuration files, messages, network packet data, performance measurements and sensor measurements. An event typically includes a timestamp that may be derived from the raw data in the event, or may be determined through interpolation between temporally proximate events having known timestamps.

**[00079]** The SPLUNK® ENTERPRISE system also facilitates using a flexible schema to specify how to extract information from the event data, wherein the flexible schema may be developed and redefined as needed. Note that a flexible schema may be applied to event data “on the fly,” when it is needed (e.g., at search time), rather than at ingestion time of the data as in traditional database systems. Because the schema is not applied to event data until it is needed (e.g., at search time), it is referred to as a “late-binding schema.”

**[00080]** During operation, the SPLUNK® ENTERPRISE system starts with raw data, which can include unstructured data, machine data, performance measurements or other time-series data, such as data obtained from web logs, syslogs, or sensor readings. It divides this raw data into “portions,” and optionally transforms the data to produce timestamped events. The system stores the timestamped events in a data store, and enables a user to run queries against the data store to retrieve events that meet specified criteria, such as containing certain keywords or

having specific values in defined fields. Note that the term “field” refers to a location in the event data containing a value for a specific data item.

**[00081]** As noted above, the SPLUNK® ENTERPRISE system facilitates using a late-binding schema while performing queries on events. A late-binding schema specifies “extraction rules” that are applied to data in the events to extract values for specific fields. More specifically, the extraction rules for a field can include one or more instructions that specify how to extract a value for the field from the event data. An extraction rule can generally include any type of instruction for extracting values from data in events. In some cases, an extraction rule comprises a regular expression, in which case the rule is referred to as a “regex rule.”

**[00082]** In contrast to a conventional schema for a database system, a late-binding schema is not defined at data ingestion time. Instead, the late-binding schema can be developed on an ongoing basis until the time a query is actually executed. This means that extraction rules for the fields in a query may be provided in the query itself, or may be located during execution of the query. Hence, as an analyst learns more about the data in the events, the analyst can continue to refine the late-binding schema by adding new fields, deleting fields, or changing the field extraction rules until the next time the schema is used by a query. Because the SPLUNK® ENTERPRISE system maintains the underlying raw data and provides a late-binding schema for searching the raw data, it enables an analyst to investigate questions that arise as the analyst learns more about the events.

**[00083]** In the SPLUNK® ENTERPRISE system, a field extractor may be configured to automatically generate extraction rules for certain fields in the events when the events are being

created, indexed, or stored, or possibly at a later time. Alternatively, a user may manually define extraction rules for fields using a variety of techniques.

**[00084]** Also, a number of “default fields” that specify metadata about the events rather than data in the events themselves can be created automatically. For example, such default fields can specify: a timestamp for the event data; a host from which the event data originated; a source of the event data; and a source type for the event data. These default fields may be determined automatically when the events are created, indexed or stored.

**[00085]** In some embodiments, a “tag” may be assigned to two or more fields that contain equivalent data items, even though the fields are associated with different events and possibly different extraction rules. By enabling a single tag (e.g., a field name) to be used to identify equivalent fields from different types of events generated by different data sources, the system facilitates use of a “common information model” (CIM) across the different data sources.

**[00086]** FIG. 4 presents a block diagram of an exemplary event-processing system 100, similar to the SPLUNK® ENTERPRISE system. System 100 includes one or more forwarders 101 that collect data obtained from a variety of different data sources 105, and one or more indexers 102 that store, process, and/or perform operations on this data, wherein each indexer operates on data contained in a specific data store 103. These forwarders and indexers can comprise separate computer systems in a data center, or may alternatively comprise separate processes executing on various computer systems in a data center.

**[00087]** During operation, the forwarders 101 identify which indexers 102 will receive the collected data and then forward the data to the identified indexers. Also, certain forwarders 101, referred to as “heavy forwarders,” can perform operations to strip extraneous data and detect

timestamps for the data. (In contrast, normal forwarders do not perform such operations as stripping extraneous data and detecting timestamps.) Based on the detected timestamps, the heavy forwarders can index and group the data into buckets that fall within a common time range. The heavy forwarders next determine which indexers 102 will receive each bucket of data and then forward the data to these indexers 102.

**[00088]** Note that distributing buckets for different time ranges across different indexers facilitates parallel processing. This parallel processing can take place at data ingestion time, because multiple indexers can process incoming data for different time ranges in parallel. The parallel processing can also take place at search time, because multiple indexers can perform searches for different time ranges in parallel while processing a query.

**[00089]** System 100 and the processes described below with respect to FIGs. 1-5 are further described in “Exploring Splunk Search Processing Language (SPL) Primer and Cookbook” by David Carasso, CITO Research, 2012, and in “Optimizing Data Analysis With a Semi-Structured Time Series Database” by Ledion Bitincka, Archana Ganapathi, Stephen Sorkin, and Steve Zhang, SLAML, 2010, each of which is hereby incorporated herein by reference in its entirety for all purposes.

**[00090]** FIG. 5 presents a flowchart illustrating how an indexer processes, indexes, and stores data received from forwarders in accordance with the disclosed embodiments. At block 201, the indexer receives the data from the forwarder. Next, at block 202, the indexer apportions the data into events. Note that the data can include lines of text that are separated by carriage returns or line breaks and an event may include one or more of these lines. During the apportioning process, the indexer can use heuristic rules to automatically determine the boundaries of the

events, which for example coincide with line boundaries. These heuristic rules may be determined based on the source of the data, wherein the indexer can be explicitly informed about the source of the data or can infer the source of the data by examining the data. These heuristic rules can include regular expression-based rules or delimiter-based rules for determining event boundaries, wherein the event boundaries may be indicated by predefined characters or character strings. These predefined characters may include punctuation marks or other special characters including, for example, carriage returns, tabs, spaces or line breaks. In some cases, a user can fine-tune or configure the rules that the indexers use to determine event boundaries in order to adapt the rules to the user's specific requirements.

**[00091]** Next, the indexer determines a timestamp for each event at block 203. As mentioned above, these timestamps can be determined by extracting the time directly from data in the event, or by interpolating the time based on timestamps from temporally proximate events. In some cases, a timestamp can be determined based on the time the data was received or generated. The indexer subsequently associates the determined timestamp with each event at block 204, for example by storing the timestamp as metadata for each event.

**[00092]** Then, the system can apply transformations to data to be included in events at block 205. For log data, such transformations can include removing a portion of an event (e.g., a portion used to define event boundaries, extraneous text, characters, etc.) or removing redundant portions of an event. Note that a user can specify portions to be removed using a regular expression or any other possible technique.

**[00093]** Next, a keyword index can optionally be generated to facilitate fast keyword searching for events. To build a keyword index, the indexer first identifies a set of keywords in

block 206. Then, at block 207 the indexer includes the identified keywords in an index, which associates each stored keyword with references to events containing that keyword (or to locations within events where that keyword is located). When an indexer subsequently receives a keyword-based query, the indexer can access the keyword index to quickly identify events containing the keyword.

**[00094]** In some embodiments, the keyword index may include entries for name-value pairs found in events, wherein a name-value pair can include a pair of keywords connected by a symbol, such as an equals sign or colon. In this way, events containing these name-value pairs can be quickly located. In some embodiments, fields can automatically be generated for some or all of the name-value pairs at the time of indexing. For example, if the string “dest=10.0.1.2” is found in an event, a field named “dest” may be created for the event, and assigned a value of “10.0.1.2.”

**[00095]** Finally, the indexer stores the events in a data store at block 208, wherein a timestamp can be stored with each event to facilitate searching for events based on a time range. In some cases, the stored events are organized into a plurality of buckets, wherein each bucket stores events associated with a specific time range. This not only improves time-based searches, but it also allows events with recent timestamps that may have a higher likelihood of being accessed to be stored in faster memory to facilitate faster retrieval. For example, a bucket containing the most recent events can be stored as flash memory instead of on hard disk.

**[00096]** Each indexer 102 is responsible for storing and searching a subset of the events contained in a corresponding data store 103. By distributing events among the indexers and data stores, the indexers can analyze events for a query in parallel, for example using map-reduce

techniques, wherein each indexer returns partial responses for a subset of events to a search head that combines the results to produce an answer for the query. By storing events in buckets for specific time ranges, an indexer may further optimize searching by looking only in buckets for time ranges that are relevant to a query.

**[00097]** Moreover, events and buckets can also be replicated across different indexers and data stores to facilitate high availability and disaster recovery as is described in U.S. Patent Application No. 14/266,812 filed on 30 April 2014, and in U.S. Application Patent No. 14/266,817 also filed on 30 April 2014.

**[00098]** FIG. 6 presents a flowchart illustrating how a search head and indexers perform a search query in accordance with the disclosed embodiments. At the start of this process, a search head receives a search query from a client at block 301. Next, at block 302, the search head analyzes the search query to determine what portions can be delegated to indexers and what portions should be executed locally by the search head. At block 303, the search head distributes the determined portions of the query to the indexers. Note that commands that operate on single events can be trivially delegated to the indexers, while commands that involve events from multiple indexers are harder to delegate.

**[00099]** Then, at block 304, the indexers to which the query was distributed search their data stores for events that are responsive to the query. To determine which events are responsive to the query, the indexer searches for events that match the criteria specified in the query. These criteria can include matching keywords or specific values for certain fields. In a query that uses a late-binding schema, the searching operations in block 304 may involve using the late-binding scheme to extract values for specified fields from events at the time the query is processed.

Next, the indexers can either send the relevant events back to the search head, or can use the events to calculate a partial result, and send the partial result back to the search head.

**[000100]** Finally, at block 305, the search head combines the partial results and/or events received from the indexers to produce a final result for the query. This final result can comprise different types of data depending upon what the query is asking for. For example, the final results can include a listing of matching events returned by the query, or some type of visualization of data from the returned events. In another example, the final result can include one or more calculated values derived from the matching events.

**[000101]** Moreover, the results generated by system 100 can be returned to a client using different techniques. For example, one technique streams results back to a client in real-time as they are identified. Another technique waits to report results to the client until a complete set of results is ready to return to the client. Yet another technique streams interim results back to the client in real-time until a complete set of results is ready, and then returns the complete set of results to the client. In another technique, certain results are stored as “search jobs,” and the client may subsequently retrieve the results by referencing the search jobs.

**[000102]** The search head can also perform various operations to make the search more efficient. For example, before the search head starts executing a query, the search head can determine a time range for the query and a set of common keywords that all matching events should include. Next, the search head can use these parameters to query the indexers to obtain a superset of the eventual results. Then, during a filtering stage, the search head can perform field-extraction operations on the superset to produce a reduced set of search results.

**[000103]** FIG. 7 presents a block diagram illustrating how fields can be extracted during query processing in accordance with the disclosed embodiments. At the start of this process, a search query 402 is received at a query processor 404. Query processor 404 includes various mechanisms for processing a query, wherein these mechanisms can reside in a search head 104 and/or an indexer 102. Note that the exemplary search query 402 illustrated in FIG. 7 is expressed in Search Processing Language (SPL), which is used in conjunction with the SPLUNK® ENTERPRISE system. SPL is a pipelined search language in which a set of inputs is operated on by a first command in a command line, and then a subsequent command following the pipe symbol “|” operates on the results produced by the first command, and so on for additional commands. Search query 402 can also be expressed in other query languages, such as the Structured Query Language (“SQL”) or any suitable query language.

**[000104]** Upon receiving search query 402, query processor 404 sees that search query 402 includes two fields “IP” and “target.” Query processor 404 also determines that the values for the “IP” and “target” fields have not already been extracted from events in data store 414, and consequently determines that query processor 404 needs to use extraction rules to extract values for the fields. Hence, query processor 404 performs a lookup for the extraction rules in a rule base 406, wherein rule base 406 maps field names to corresponding extraction rules and obtains extraction rules 408-409, wherein extraction rule 408 specifies how to extract a value for the “IP” field from an event, and extraction rule 409 specifies how to extract a value for the “target” field from an event. As is illustrated in FIG. 7, extraction rules 408-409 can comprise regular expressions that specify how to extract values for the relevant fields. Such regular-expression-based extraction rules are also referred to as “regex rules.” In addition to specifying how to extract field values, the extraction rules may also include instructions for deriving a field value

by performing a function on a character string or value retrieved by the extraction rule. For example, a transformation rule may truncate a character string, or convert the character string into a different data format. In some cases, the query itself can specify one or more extraction rules.

**[000105]** Next, query processor 404 sends extraction rules 408-409 to a field extractor 412, which applies extraction rules 408-409 to events 416-418 in a data store 414. Note that data store 414 can include one or more data stores, and extraction rules 408-409 can be applied to large numbers of events in data store 414, and are not meant to be limited to the three events 416-418 illustrated in FIG. 7. Moreover, the query processor 404 can instruct field extractor 412 to apply the extraction rules to all the events in a data store 414, or to a subset of the events that have been filtered based on some criteria.

**[000106]** Next, field extractor 412 applies extraction rule 408 for the first command “Search IP= “10\*” to events in data store 414 including events 416-418. Extraction rule 408 is used to extract values for the IP address field from events in data store 414 by looking for a pattern of one or more digits, followed by a period, followed again by one or more digits, followed by another period, followed again by one or more digits, followed by another period, and followed again by one or more digits. Next, field extractor 412 returns field values 420 to query processor 404, which uses the criterion IP=“10\*” to look for IP addresses that start with “10”. Note that events 416 and 417 match this criterion, but event 418 does not, so the result set for the first command is events 416-417.

**[000107]** Query processor 404 then sends events 416-417 to the next command “stats count target.” To process this command, query processor 404 causes field extractor 412 to apply

extraction rule 409 to events 416-417. Extraction rule 409 is used to extract values for the target field for events 416-417 by skipping the first four commas in events 416-417, and then extracting all of the following characters until a comma or period is reached. Next, field extractor 412 returns field values 421 to query processor 404, which executes the command “stats count target” to count the number of target fields that were read, which in this example produces the value “2” that is returned as a final result 422 for the query.

**[000108]** Note that query results can be returned to a client, or to a search head, or to any other system component for further processing. In general, query results may include a set of one or more values obtained from events in data store 414; a subset of the values; statistics calculated based on the values; a report containing the values; or a visualization, such as a graph or chart, generated from the values.

**[000109]** The above-described system provides significant flexibility by enabling a user to analyze massive quantities of minimally processed performance data “on the fly” at search time instead of storing pre-specified portions of the performance data in a database at ingestion time. This flexibility enables a user to see correlations in the performance data and perform subsequent queries to examine interesting aspects of the performance data that may not have been apparent at ingestion time.

**[000110]** However, performing analysis operations at search time can involve a large amount of data and require a large number of computational operations, which can cause considerable delays while processing the queries. Fortunately, a number of acceleration techniques have been developed to speed up analysis operations performed at search time. These techniques include: (1) using a keyword index; (2) performing search operations in parallel by formulating a search

as a map-reduce computation; (3) using a high performance analytics store; and (4) accelerating the process of generating reports. These techniques are described in more detail below.

**[000111]** As described above with reference to the flow charts in FIGs. 5 and 6, event-processing system 100 can construct and maintain one or more keyword indices to facilitate rapidly identifying events containing specific keywords. This can greatly speed up the processing of queries involving specific keywords.

**[000112]** To facilitate faster query processing, a query can be structured as a map-reduce computation, wherein the “map” operations are delegated to the indexers, while the corresponding “reduce” operations are performed locally at the search head. For example, FIG. 8 illustrates how a search query 501 received from a client at search head 104 can split into two phases, including: (1) a “map phase” comprising subtasks 502 (e.g., data retrieval or simple filtering) that may be performed in parallel and are “mapped” to indexers 102 for execution, and (2) a “reduce phase” comprising a merging operation 503 to be executed by the search head when the results are ultimately collected from the indexers.

**[000113]** During operation, upon receiving search query 501, search head 104 modifies search query 501 by substituting “stats” with “prestats” to produce search query 502, and then distributes search query 502 to one or more distributed indexers, which are also referred to as “search peers.” Note that search queries may generally specify search criteria or operations to be performed on events that meet the search criteria. Search queries may also specify field names, as well as search criteria for the values in the fields or operations to be performed on the values in the fields. Moreover, the search head may distribute the full search query to the search peers as is illustrated in FIG. 8, or may alternatively distribute a modified version (e.g., a more

restricted version) of the search query to the search peers. In this example, the indexers are responsible for counting the results and sending them to the search head. After the indexers return the results to the search head, the search head performs the merging operations 503 on the results. Note that by executing the computation in this way, the system effectively distributes the computational operations while minimizing data transfers.

**[000114]** To speed up certain types of queries, some embodiments of system 100 make use of a high performance analytics store, which is referred to as a “summarization table,” that contains entries for specific field-value pairs. Each of these entries keeps track of instances of a specific value in a specific field in the event data and includes references to events containing the specific value in the specific field. For example, an exemplary entry in a summarization table can keep track of occurrences of the value “94107” in a “ZIP code” field of a set of events, wherein the entry includes references to all of the events that contain the value “94107” in the ZIP code field. This enables the system to quickly process queries that seek to determine how many events have a particular value for a particular field, because the system can examine the entry in the summarization table to count instances of the specific value in the field without having to go through the individual events or do extractions at search time. Also, if the system needs to process all events that have a specific field-value combination, the system can use the references in the summarization table entry to directly access the events to extract further information without having to search all of the events to find the specific field-value combination at search time.

**[000115]** In some embodiments, the system maintains a separate summarization table for each of the above-described time-specific buckets that stores events for a specific time range, wherein a bucket-specific summarization table includes entries for specific field-value combinations that

occur in events in the specific bucket. Alternatively, the system can maintain a separate summarization table for each indexer, wherein the indexer-specific summarization table only includes entries for the events in a data store that is managed by the specific indexer.

**[000116]** The summarization table can be populated by running a “collection query” that scans a set of events to find instances of a specific field-value combination, or alternatively instances of all field-value combinations for a specific field. A collection query can be initiated by a user, or can be scheduled to occur automatically at specific time intervals. A collection query can also be automatically launched in response to a query that asks for a specific field-value combination.

**[000117]** In some cases, the summarization tables may not cover all of the events that are relevant to a query. In this situation, the system can use the summarization tables to obtain partial results for the events that are covered by summarization tables, but may also have to search other events that are not covered by the summarization tables. These partial results and results obtained by searching the other events can then be combined to produce a final set of results for the query. This summarization table and associated techniques are described in more detail in U.S. Patent No. 8,682,925, issued on March 25, 2014.

**[000118]** In some embodiments, a data server system such as the SPLUNK® ENTERPRISE system can accelerate the process of periodically generating updated reports based on query results. To accelerate this process, a summarization engine automatically examines the query to determine whether generation of updated reports can be accelerated by creating intermediate summaries. (This is possible if results from preceding time periods can be computed separately and combined to generate an updated report. In some cases, it is not possible to combine such incremental results, for example where a value in the report depends on relationships between

events from different time periods.) If reports can be accelerated, the summarization engine periodically generates a summary covering data obtained during a latest non-overlapping time period. For example, where the query seeks events meeting specified criteria, a summary for the time period includes only events within the time period that meet the specified criteria.

Similarly, if the query seeks statistics calculated from the events, such as the number of events that match the specified criteria, then the summary for the time period includes the number of events in the period that match the specified criteria.

**[000119]** In parallel with the creation of the summaries, the summarization engine schedules the periodic updating of the report associated with the query. During each scheduled report update, the query engine determines whether intermediate summaries have been generated covering portions of the time period covered by the report update. If so, then the report is generated based on the information contained in the summaries. Also, if additional event data has been received and has not yet been summarized, and is required to generate the complete report, the query can be run on this additional event data. Then, the results returned by this query on the additional event data, along with the partial results obtained from the intermediate summaries, can be combined to generate the updated report. This process is repeated each time the report is updated. Alternatively, if the system stores events in buckets covering specific time ranges, then the summaries can be generated on a bucket-by-bucket basis. Note that producing intermediate summaries can save the work involved in re-running the query for previous time periods, so only the newer event data needs to be processed while generating an updated report. These report acceleration techniques are described in more detail in U.S. Patent No. 8,589,403, issued on November 19, 2013, and U.S. Patent No. 8,412,696, issued on April 2, 2011.

**[000120]** FIG. 9 illustrates a diagrammatic representation of a computing device 1000 within which a set of instructions for causing the computing device to perform the methods discussed herein may be executed. The computing device 1000 may be connected to other computing devices in a LAN, an intranet, an extranet, and/or the Internet. The computing device 1000 may operate in the capacity of a server machine in client-server network environment. The computing device 1000 may be provided by a personal computer (PC), a set-top box (STB), a server, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single computing device is illustrated, the term “computing device” shall also be taken to include any collection of computing devices that individually or jointly execute a set (or multiple sets) of instructions to perform the methods discussed herein. In illustrative examples, the computing device 1000 may implement the above described methods 300A-300C for managing datasets produced by alert-triggering search queries.

**[000121]** The example computing device 1000 may include a processing device (*e.g.*, a general purpose processor) 1002, a main memory 1004 (*e.g.*, synchronous dynamic random access memory (DRAM), read-only memory (ROM)), a static memory 1006 (*e.g.*, flash memory and a data storage device 1018), which may communicate with each other via a bus 1030.

**[000122]** The processing device 1002 may be provided by one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. In an illustrative example, the processing device 1002 may comprise a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or a processor implementing other instruction sets or processors implementing a combination of instruction sets. The processing

device 1002 may also comprise one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), a network processor, or the like. The processing device 1002 may be configured to execute the methods 300A-300C for managing datasets produced by alert-triggering search queries, in accordance with one or more aspects of the present disclosure.

**[000123]** The computing device 1000 may further include a network interface device 1008, which may communicate with a network 1020. The computing device 1000 also may include a video display unit 1010 (*e.g.*, a liquid crystal display (LCD) or a cathode ray tube (CRT)), an alphanumeric input device 1012 (*e.g.*, a keyboard), a cursor control device 1014 (*e.g.*, a mouse) and an acoustic signal generation device 1016 (*e.g.*, a speaker). In one embodiment, video display unit 1010, alphanumeric input device 1012, and cursor control device 1014 may be combined into a single component or device (*e.g.*, an LCD touch screen).

**[000124]** The data storage device 1018 may include a computer-readable storage medium 1028 on which may be stored one or more sets of instructions (*e.g.*, instructions of the methods 300A-300C for managing datasets produced by alert-triggering search queries, in accordance with one or more aspects of the present disclosure) implementing any one or more of the methods or functions described herein. Instructions implementing methods 300A-300C may also reside, completely or at least partially, within main memory 1004 and/or within processing device 1002 during execution thereof by computing device 1000, main memory 1004 and processing device 1002 also constituting computer-readable media. The instructions may further be transmitted or received over a network 1020 via network interface device 1008.

**[000125]** While computer-readable storage medium 1028 is shown in an illustrative example to be a single medium, the term “computer-readable storage medium” should be taken to include a single medium or multiple media (*e.g.*, a centralized or distributed database and/or associated caches and servers) that store one or more sets of instructions. The term “computer-readable storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by the machine and that cause the machine to perform the methods described herein. The term “computer-readable storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, optical media and magnetic media.

**[000126]** Unless specifically stated otherwise, terms such as "updating," "identifying," "determining," "sending," "assigning," or the like refer to actions and processes performed or implemented by computing devices that manipulate and transform data represented as physical (electronic) quantities within the computing device's registers and memories into other data similarly represented as physical quantities within the computing device memories or registers or other such information storage, transmission or display devices. Also, the terms "first," "second," "third," "fourth," etc. as used herein are meant as labels to distinguish among different elements and may not necessarily have an ordinal meaning according to their numerical designation.

**[000127]** Examples described herein also relate to an apparatus for performing the methods described herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computing device selectively programmed by a computer program stored in the computing device. Such a computer program may be stored in a computer-readable non-transitory storage medium.

**[000128]** The methods and illustrative examples described herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used in accordance with the teachings described herein, or it may prove convenient to construct more specialized apparatus to perform the required method operations. The required structure for a variety of these systems will appear as set forth in the description above.

**[000129]** The above description is intended to be illustrative, and not restrictive. Although the present disclosure has been described with references to specific illustrative examples, it will be recognized that the present disclosure is not limited to the examples described. The scope of the disclosure should be determined with reference to the following claims, along with the full scope of equivalents to which the claims are entitled.

## CLAIMS

What is claimed is:

1. A method, comprising:

executing, by one or more processing devices, a search query on a portion of searchable data associated with a time window to produce a dataset comprising one or more results;

responsive to determining that at least a portion of the dataset satisfies a triggering condition defining an alert associated with the search query, generating an instance of the alert;

associating, by a memory data structure, the instance of the alert with an identifier of the search query and a time parameter specifying the time window;

receiving, from a client computing device, a request for the portion of the dataset; and

responsive to determining that the portion of the dataset is not stored in the memory in a manner associating it with the instance of the alert, reproducing the portion of the dataset by re-executing the search query in view of the time parameter.

2. The method of claim 1, further comprising:

storing, in a memory associated with the computer system, the portion of the dataset and an association of the stored portion of the dataset with the instance of the alert.

3. The method of claim 1, further comprising:

implementing a file retention policy with respect to datasets stored in the memory, wherein the file retention policy requires deleting certain datasets responsive to evaluating corresponding file retention conditions.

4. The method of claim 1, further comprising:  
transmitting the copy of the portion of the dataset to the client computing device.
5. The method of claim 1, further comprising associating the instance of the alert with an identifier of the triggering condition.
6. The method of claim 1, wherein the searchable data includes time-stamped events having portions of raw machine data.
7. The method of claim 1, further comprising:  
transmitting, to the client computing device, a notification of the instance of the alert.
8. The method of claim 1, wherein the client computing device includes at least one of: a desktop computing device or a mobile computing device.
9. The method of claim 1, wherein executing the search query on the portion of searchable data includes applying a late binding schema to the data, the late binding schema associated with one or more extraction rules defining one or more fields.
10. The method of claim 1, wherein the portion of searchable data includes machine data generated by at least one of a server, a database, an application, or a network.
11. The method of claim 1, wherein the search query is execute in near real-time.

12. The method of claim 1, wherein the search query is executed on a schedule that is associated with the alert.

13. The method of claim 1, wherein the search query and triggering condition together evaluate portions of the searchable data falling within a rolling time window.

14. The method of claim 1, wherein the triggering condition requires that the portion of the dataset includes at least a predetermined number of results.

15. The method of claim 1, wherein the triggering condition comprises a secondary conditional search on the dataset produced by the search query.

16. The method of claim 1, further comprising:  
performing at least one action associated with the alert, wherein the action includes:  
sending an electronic mail message, creating a Really Simple Syndication (RSS) feed, executing a script, or causing visual display of the alert instance.

17. A computer system comprising:  
a memory; and  
one or more processing devices, coupled to the memory, to:  
execute a search query on a portion of searchable data associated with a time window to produce a dataset comprising one or more results;

responsive to determining that at least a portion of the dataset satisfies a triggering condition defining an alert associated with the search query, generate an instance of the alert;

associate, by a memory data structure, the instance of the alert with an identifier of the search query and a time parameter specifying the time window;

receive, from a client computing device, a request for the portion of the dataset;  
and

responsive to determining that the portion of the dataset is not stored in the memory in a manner associating it with the instance of the alert, reproduce the portion of the dataset by re-executing the search query in view of the time parameter.

18. The computer system of claim 17, wherein the processing devices are further to:  
store, in a memory associated with the computer system, the portion of the dataset and an association of the stored portion of the dataset with the instance of the alert.

19. The computer system of claim 17, wherein the processing devices are further to:  
implement a file retention policy with respect to datasets stored in the memory, wherein the file retention policy requires deleting certain datasets responsive to evaluating corresponding file retention conditions.

20. The computer system of claim 17, wherein the processing devices are further to:  
transmit the copy of the portion of the dataset to the client computing device.

21. The computer system of claim 17, wherein the processing devices are further to:  
associate the instance of the alert with an identifier of the triggering condition.

22. The computer system of claim 17, wherein the searchable data includes time-stamped  
events having portions of raw machine data.

23. The computer system of claim 17, wherein the processing devices are further to:  
transmit, to the client computing device, a notification of the instance of the alert.

24. The computer system of claim 17, wherein executing the search query on the portion  
of searchable data includes applying a late binding schema to the data, the late binding schema  
associated with one or more extraction rules defining one or more fields.

25. A computer-readable non-transitory storage medium comprising executable  
instructions that, when executed by a computer system, cause the computer system to perform  
operations comprising:

executing a search query on a portion of searchable data associated with a time window  
to produce a dataset comprising one or more results;

responsive to determining that at least a portion of the dataset satisfies a triggering  
condition defining an alert associated with the search query, generating an instance of the alert;

associating, by a memory data structure, the instance of the alert with an identifier of the  
search query and a time parameter specifying the time window;

receiving, from a client computing device, a request for the portion of the dataset; and

responsive to determining that the portion of the dataset is not stored in the memory in a manner associating it with the instance of the alert, reproducing the portion of the dataset by re-executing the search query in view of the time parameter.

26. The computer-readable non-transitory storage medium of claim 25, further comprising executable instructions causing the computer system to:

store, in a memory associated with the computer system, the portion of the dataset and an association of the stored portion of the dataset with the instance of the alert.

27. The computer-readable non-transitory storage medium of claim 25, further comprising executable instructions causing the computer system to:

implement a file retention policy with respect to datasets stored in the memory, wherein the file retention policy requires deleting certain datasets responsive to evaluating corresponding file retention conditions.

28. The computer-readable non-transitory storage medium of claim 25, further comprising executable instructions causing the computer system to:

transmit the copy of the portion of the dataset to the client computing device.

29. The computer-readable non-transitory storage medium of claim 25, further comprising executable instructions causing the computer system to:

associate the instance of the alert with an identifier of the triggering condition.

30. The computer-readable non-transitory storage medium of claim 25, further comprising executable instructions causing the computer system to:  
transmit, to the client computing device, a notification of the instance of the alert.

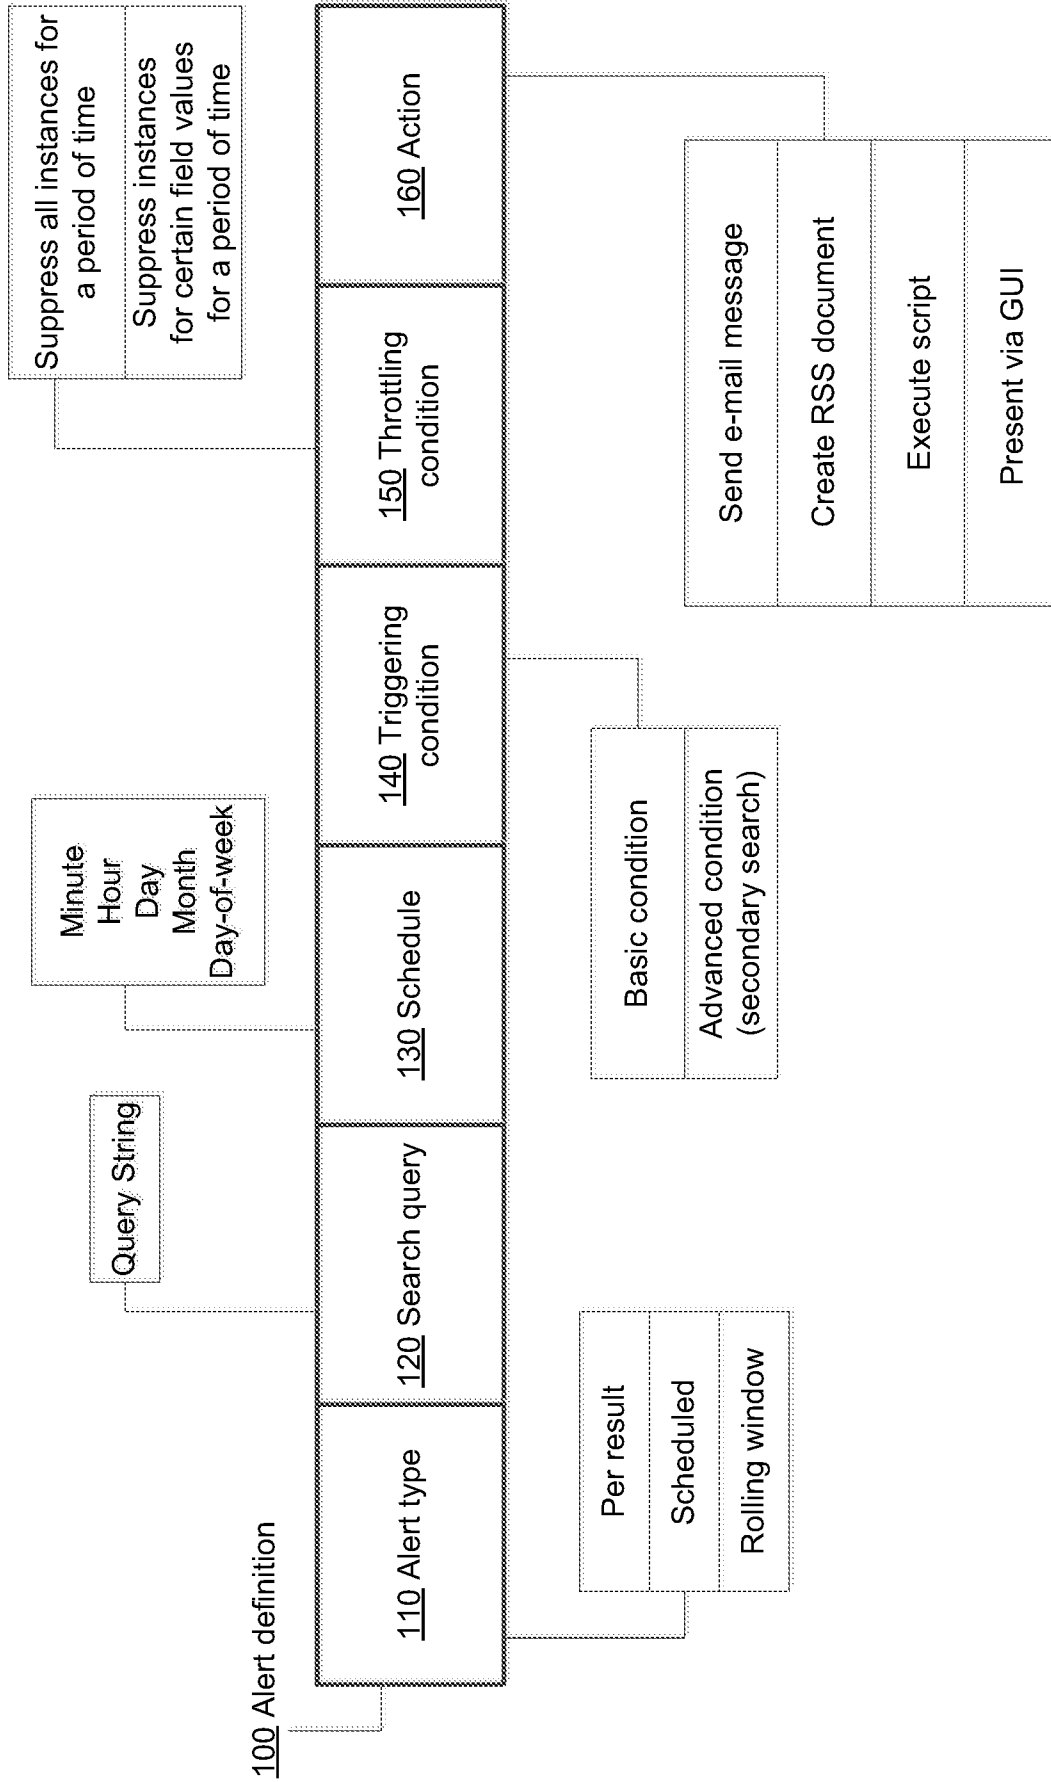


FIG. 1

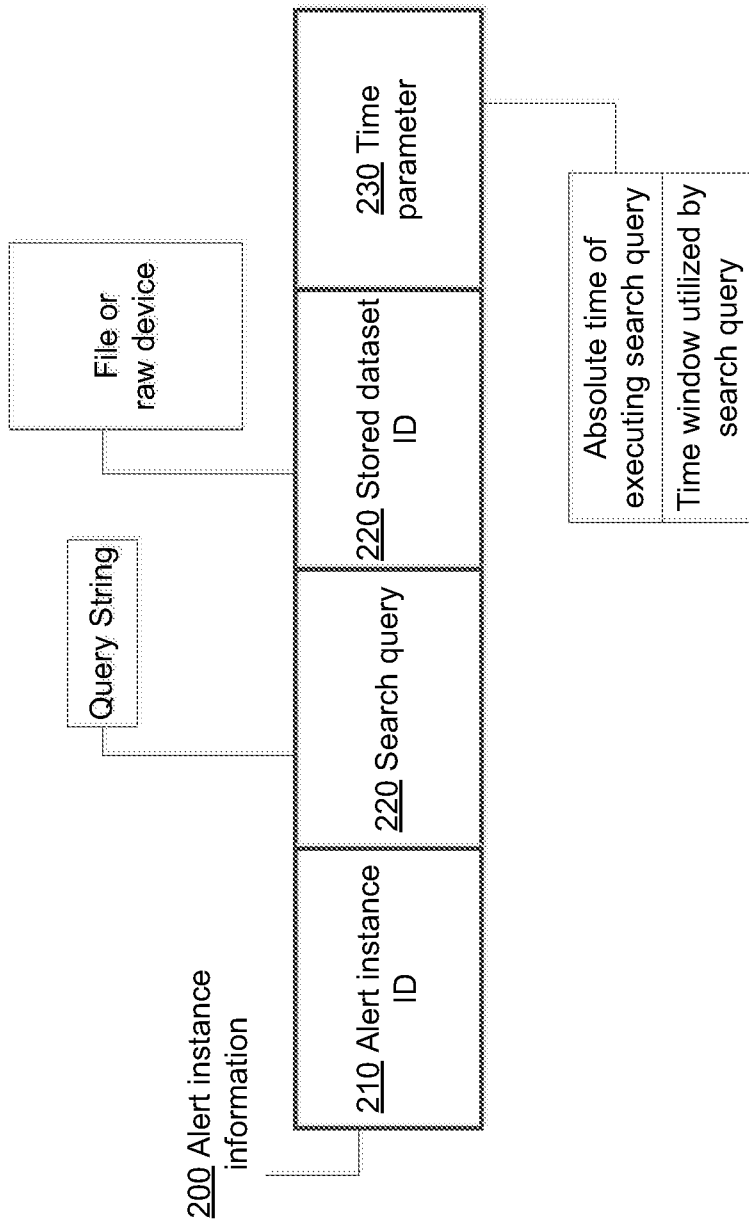


FIG. 2

300A

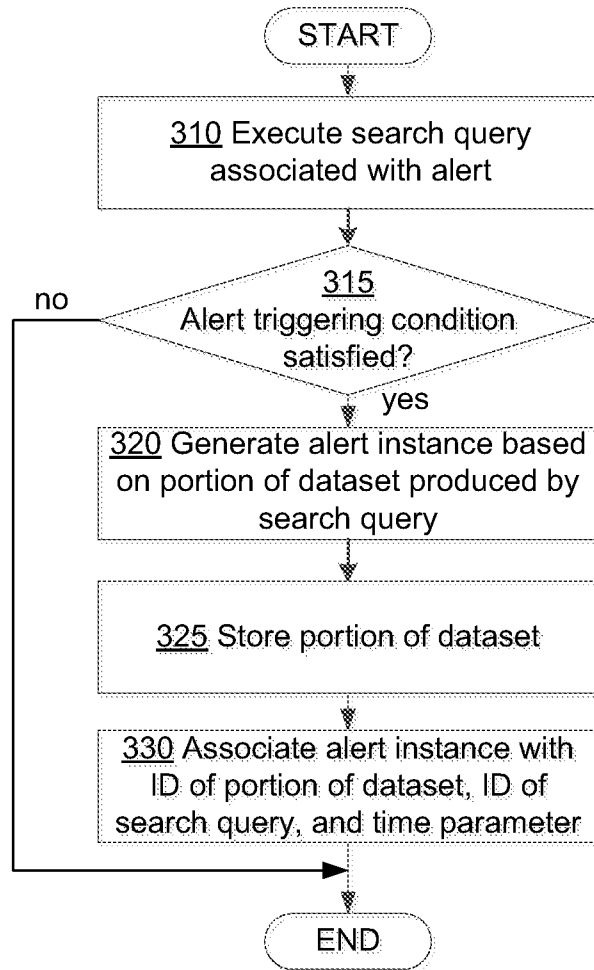


FIG. 3A

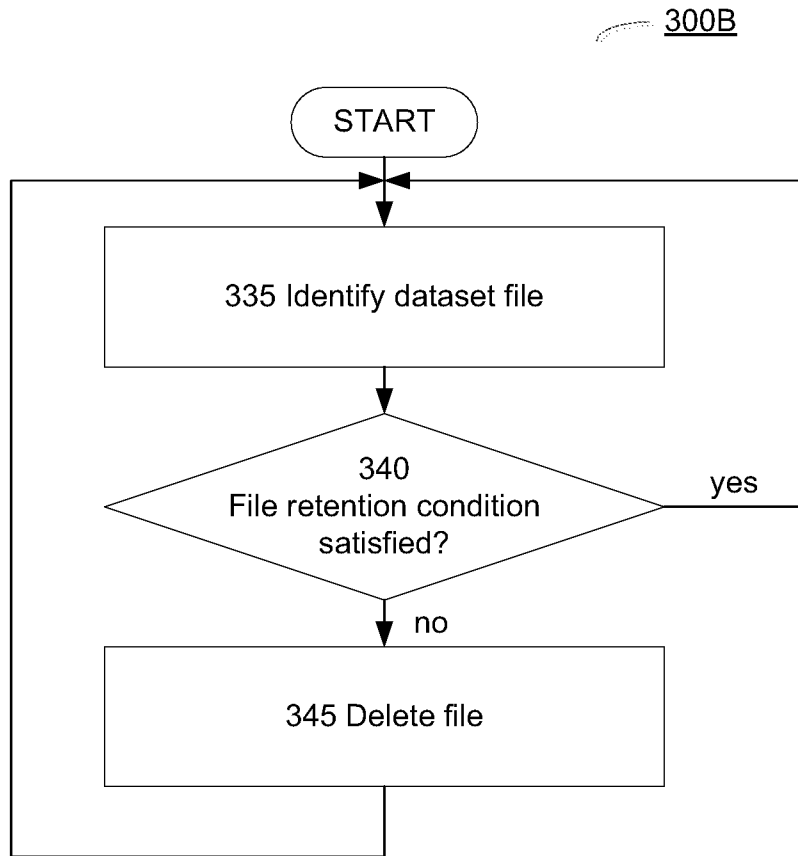


FIG. 3B

300C

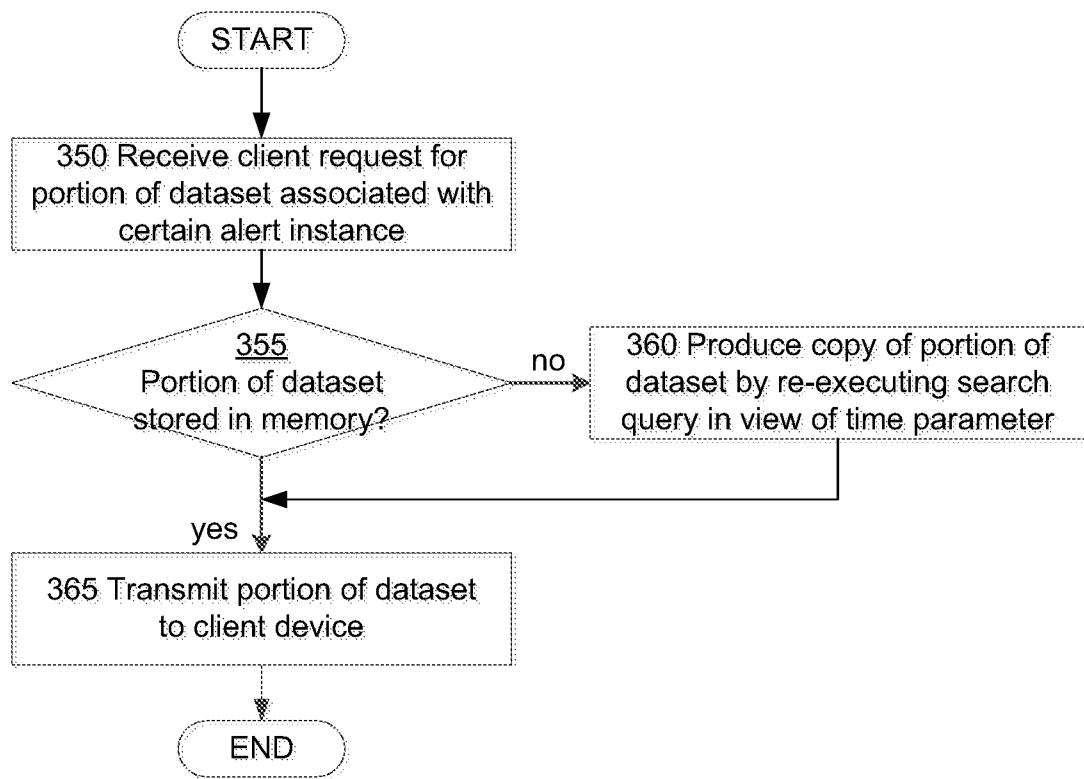


FIG. 3C

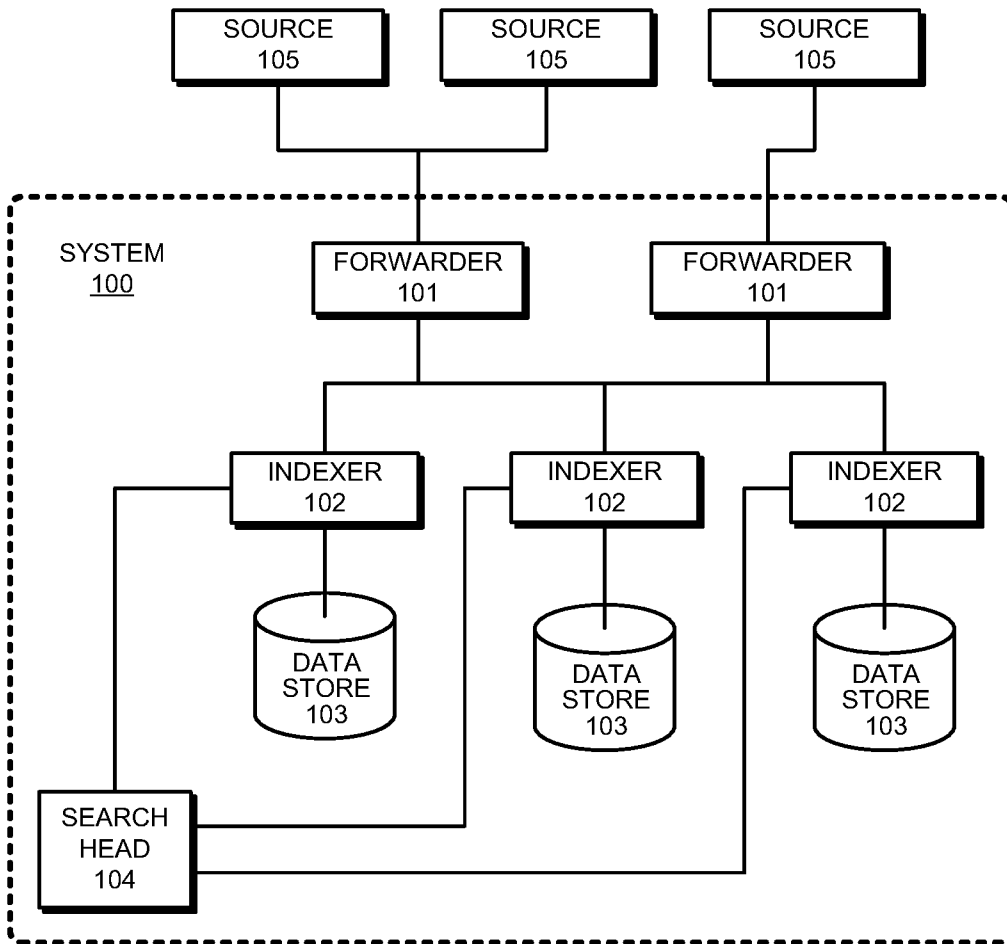
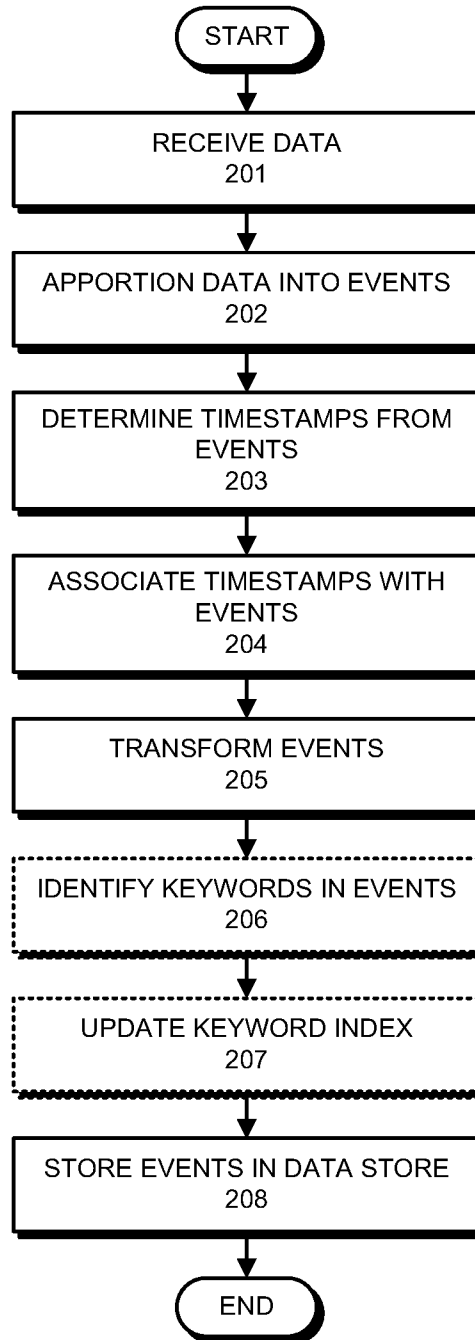
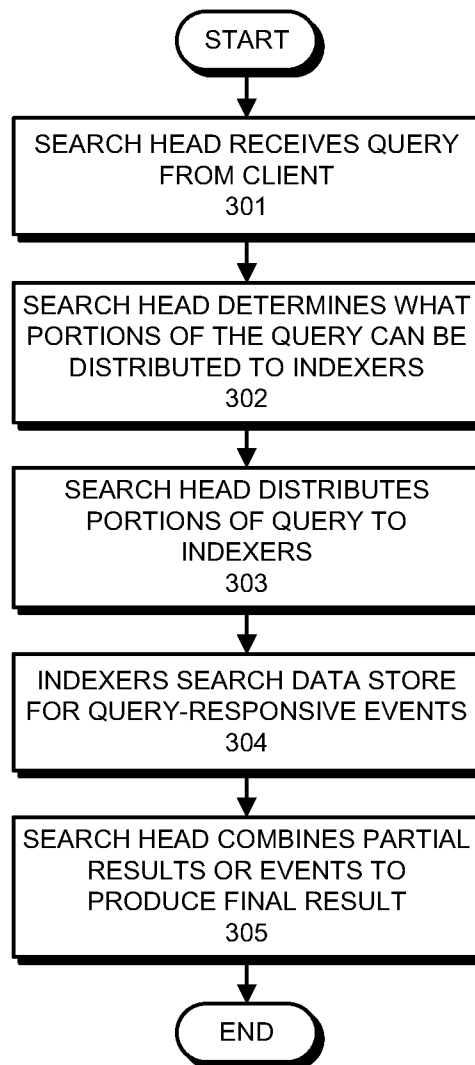


FIG. 4

**FIG. 5**

**FIG. 6**

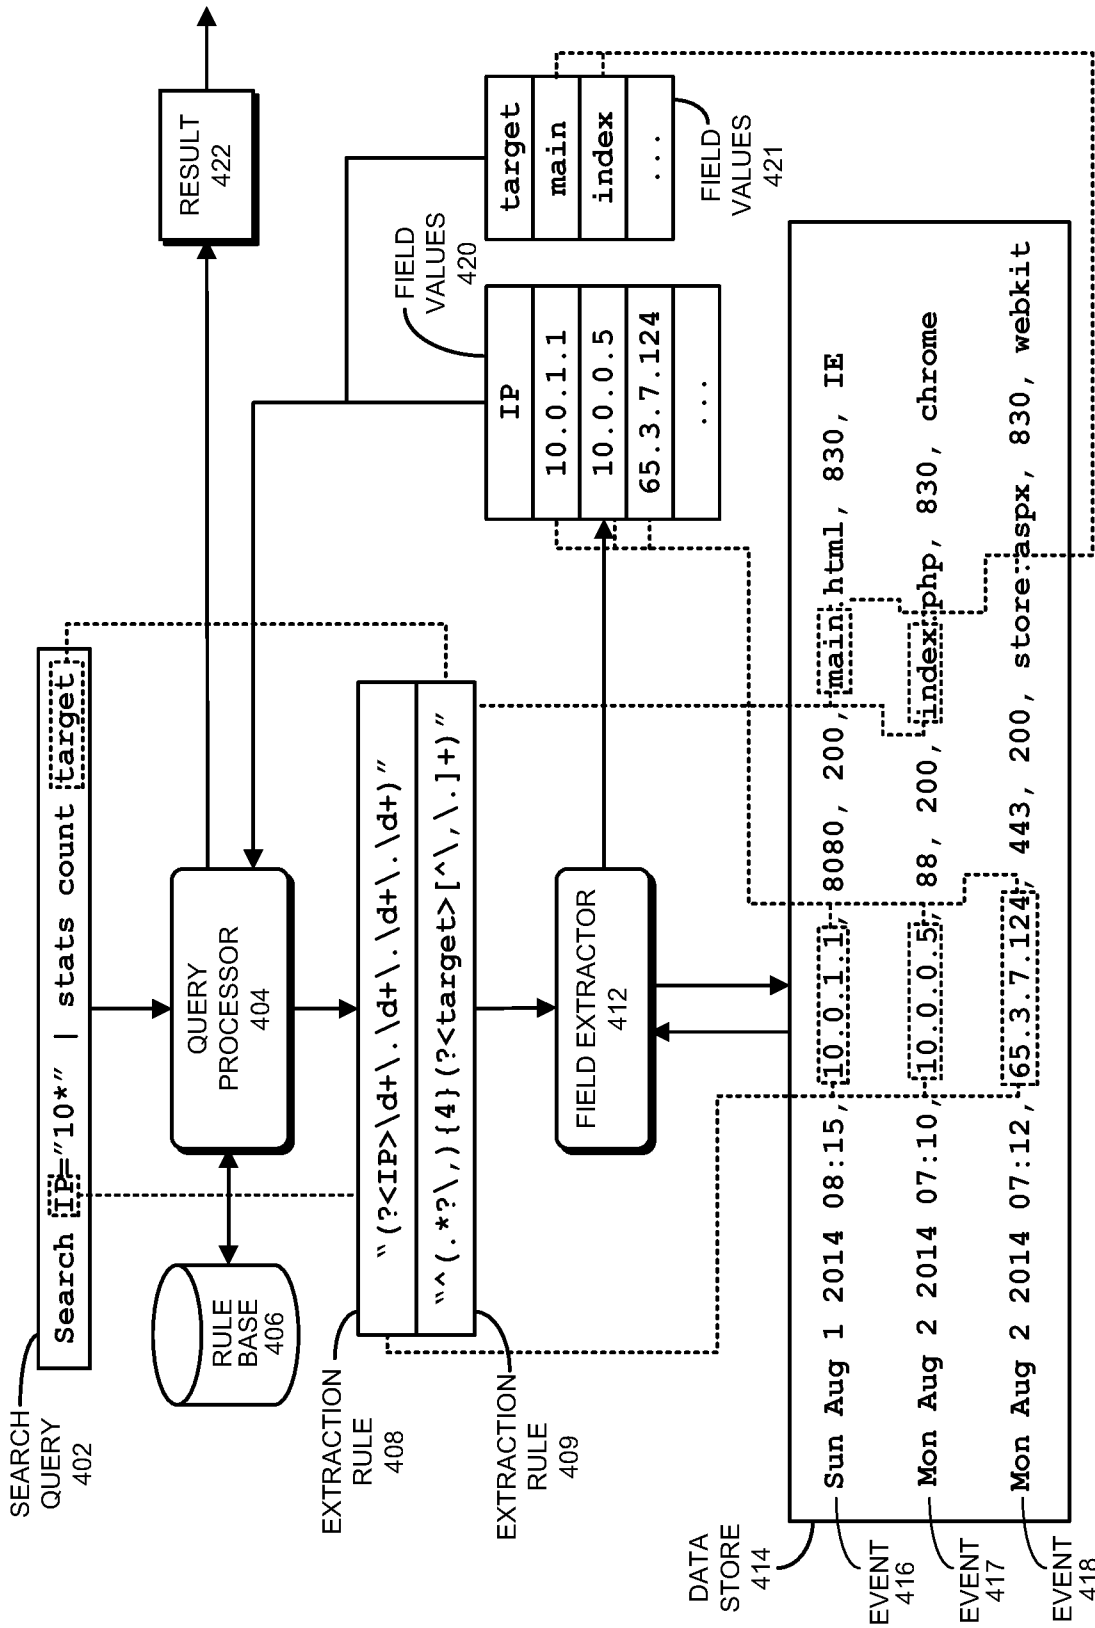


FIG. 7

Original Search: 501  
search error | stats count BY host

Sent to peers: 502  
search error | prestats count BY host

Executed by search head: 503  
Merge prestats results received from peers (reduce)

**FIG. 8**

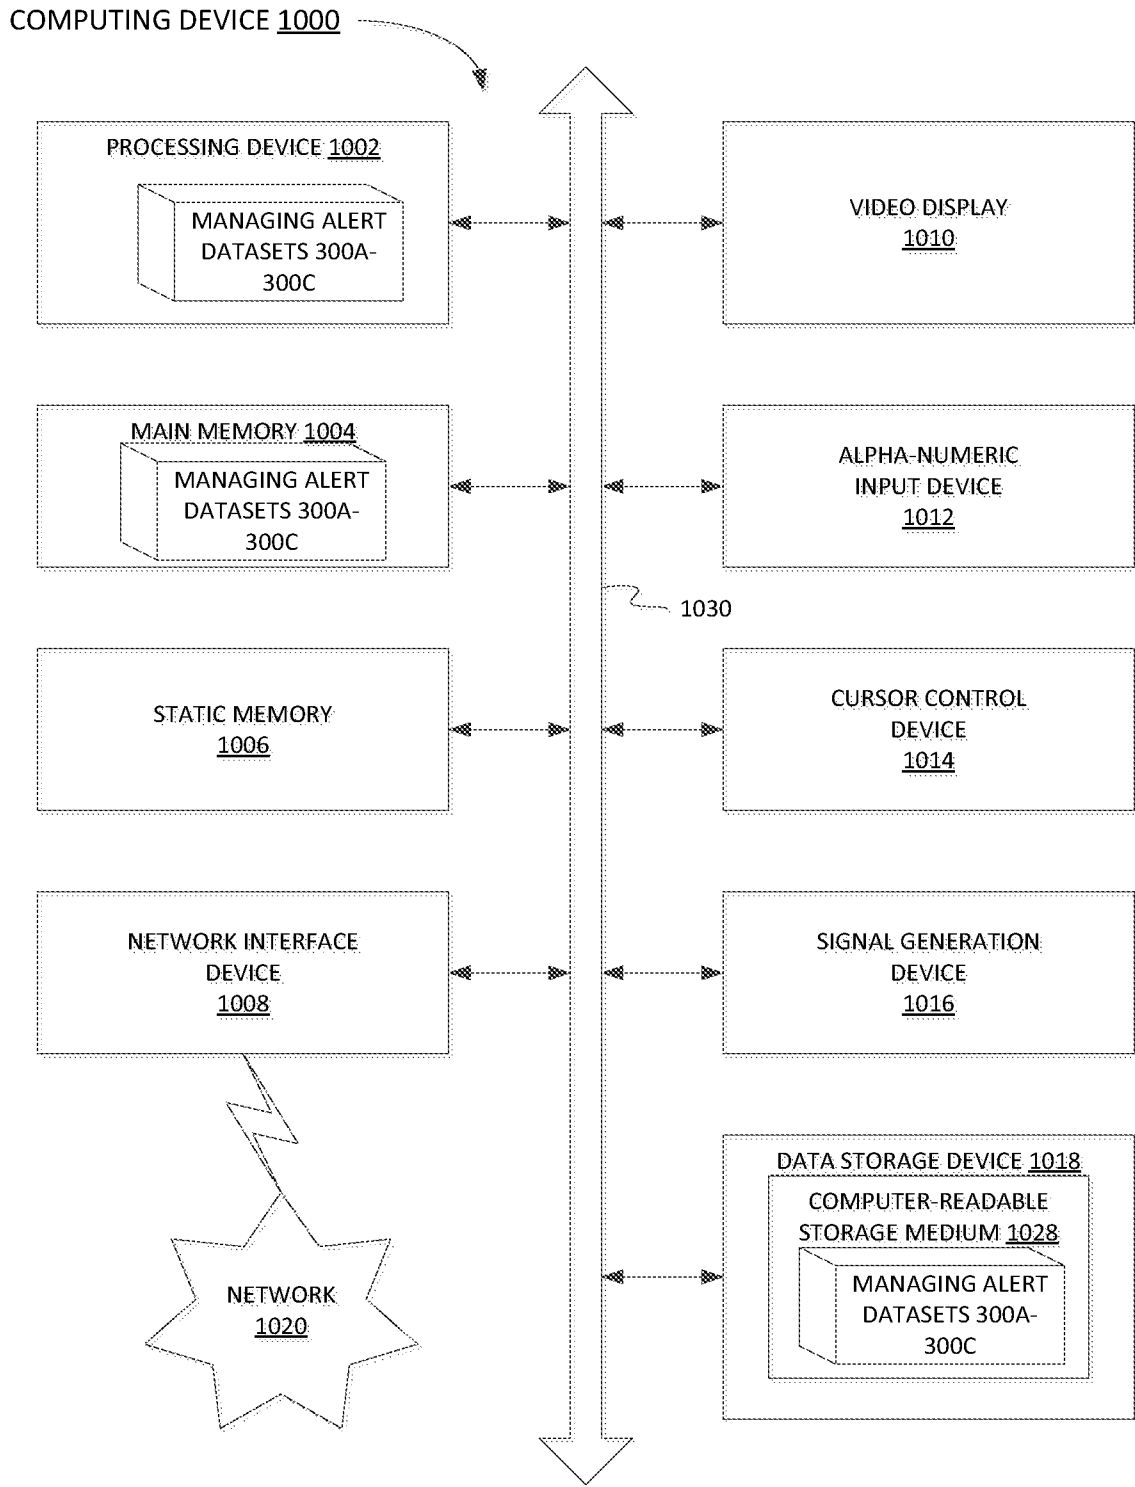


FIG. 9