(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
16 May 2019 (16.05.2019)

WIPO | PCT

(10) International Publication Number
WO 2019/090438 A1

(51) International Patent Classification:
G06F 21/60 (2013.01) *G06F 16/903* (2019.01)
G06F 16/901 (2019.01) *G06Q 10/06* (2012.01)

(21) International Application Number: PCT/CA2018/051433

(22) International Filing Date: 13 November 2018 (13.11.2018)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/585,295 13 November 2017 (13.11.2017) US

(71) Applicant: **YOPPWORKS INC.** [CA/CA]; 411 Legget Drive, Suite 701, Kanata, Ontario K2K 3C9 (CA).

(72) Inventors: **GULAS, Jack**; c/o 411 Legget Drive, Suite 701, Kanata, Ontario K2K 3C9 (CA). **MUMBY, Paul**; c/o 411 Legget Drive, Suite 701, Kanata, Ontario K2K 3C9 (CA).

(74) Agent: **SPRIGINGS, Mark** et al.; Gowling WLG (canada) LLP, Suite 2600, 160 Elgin Street, Ottawa, Ontario K1P 1C3 (CA).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,

(54) **Title:** VEHICLE ENTERPRISE FLEET MANAGEMENT SYSTEM AND METHOD

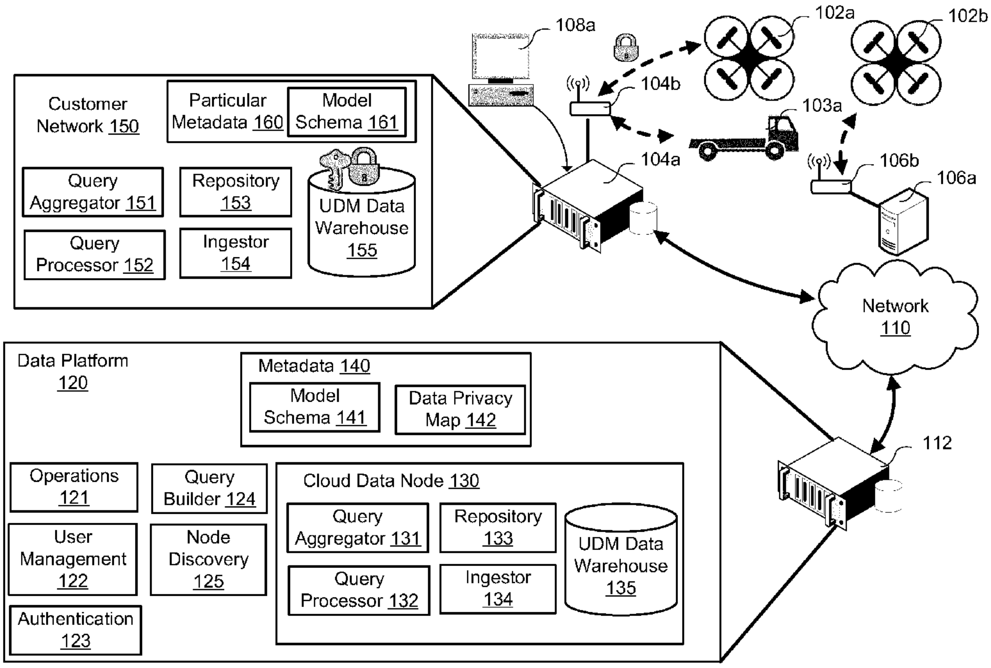

Fig. 1

(57) **Abstract:** The present disclosure relates to systems and methods for vehicle fleet management. The systems and methods disclosed herein allow for a back-end data platform to manage data collected via unmanned vehicles/robots or industrial vehicle, as well as one or more front-end applications to facilitate capturing, uploading, and reporting of data from the vehicles. Methods and systems for partitioning data received from one or more vehicles are disclosed that, among other things, allow customers to retain any sensitive data while also making any non-sensitive data publicly available. Methods and systems for responding to a query that request vehicles are disclosed that, among other things, allow reports and other data analytics tools to be generated comprising non-sensitive data from a plurality of customers/vehicles, as well as sensitive data for particular vehicles when the user has appropriate authorization to access such sensitive data.

[Continued on next page]

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))*

— *in black and white; the international application as filed contained color or greyscale and is available for download from PATENTSCOPE*

VEHICLE ENTERPRISE FLEET MANAGEMENT SYSTEM AND METHOD

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present disclosure is non-provisional application of United States Provisional Application No. 62/585,295 filed November 13, 2017, the entirety of which is hereby incorporated by reference for all purposes.

TECHNICAL FIELD

[0002] The present disclosure relates to systems and methods for managing a fleet of unmanned vehicles or industrial vehicles, and in particular to the management and analysis of data generated by one or more unmanned vehicles or industrial machinery.

BACKGROUND

[0003] Unmanned vehicles/robots are becoming increasingly popular, both for recreational, professional and military applications, such as for example unmanned aerial vehicles (UAVs) and autonomous automobiles. The data generated by such vehicles is invaluable for guiding decision-making, such as what models of unmanned vehicles may be best-suited for operational requirements of a desired activity, performance, maintenance planning, flight planning etc. In addition industrial vehicles or construction machinery, which may or may not be autonomous, are becoming more sophisticated and can collect and generate more data on their performance and operation.

[0004] From a manufacturer perspective for the data to be useful, a large population size of data is required. However, certain vehicle owners/users may have data that is deemed to be sensitive and which they do not wish to share outside of their organization. Similarly an operator may have vehicles from different manufacturers and wish to correlate data across vehicles which present data formatting challenges.

[0005] Accordingly, systems and methods that enable additional, alternative, and/or improved ways to manage and analyze data generated by one or more unmanned vehicles and industrial vehicles remains highly desirable.

SUMMARY

[0006] In accordance with an aspect of the present invention there is provided a method of partitioning data received from one or more vehicles, comprising: receiving vehicle raw data from an vehicle of the one or more vehicles, the vehicle raw data comprising both sensitive and non-sensitive data; converting the vehicle raw data into a plurality of universal data elements, each data element indicating the vehicle from which the raw data was received, a sensor reading of the vehicle, and a time that the sensor reading occurred; determining, from a data privacy map established by an owner of the vehicle, which of the universal data elements are sensitive data elements and which of the universal data elements are non-sensitive data elements; transmitting the non-sensitive data elements to an external device; and storing the sensitive data elements locally.

[0007] In accordance with an aspect of the present invention there is provided a method of responding to a query requesting vehicle data, comprising: receiving, from a user, the query requesting vehicle data, the query having a security level associated therewith; determining, based on the security level associated with the query and the vehicle data requested by the query, vehicle data elements that provide the vehicle data requested and that correspond to the security level associated with the query; determining one or more nodes storing the vehicle data elements; generating a sub-query request for each of the one or more nodes, the sub-query request indicating the vehicle data elements to be retrieved at the respective node, and an indication of an output node to which the vehicle data elements should be sent to; sending the sub-query requests to each of the one or more nodes; receiving the requested vehicle data elements from the one or more nodes; and aggregating the vehicle data elements to provide a report.

[0008] In accordance with an aspect of the present invention there is provided a system for partitioning data received from one or more vehicles and responding to a query requesting vehicle data, comprising: a first server hosting a customer network; and a second server hosting a cloud-based data platform operably coupled with the first server through a data network; wherein the first server is configured to: receive vehicle raw data from an vehicle of the one or more vehicles that is operably coupled with the customer network, the vehicle raw data comprising both sensitive and non-sensitive data; convert the vehicle raw data into a plurality of universal data elements, each data element indicating the vehicle from which the raw data was received, a sensor reading of the vehicle, and a time that the sensor reading occurred; determine, from a data privacy map established by an owner of the vehicle, which of the universal data elements are sensitive data elements and which of the universal data elements are non-sensitive data elements; transmit the non-sensitive data elements to the second server through the data network; and store the sensitive data elements locally.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] Further features and advantages of the present disclosure will become apparent from the following detailed description, taken in combination with the appended drawings, in which:

Fig. 1 shows a representation of a system for managing a fleet of unmanned vehicles and industrial vehicles;

Fig. 2 shows a representation of a system used for generating reports for a fleet of unmanned vehicles and industrial vehicles;

Fig. 3 shows a method for partitioning data received from one or more unmanned vehicles and industrial vehicles;

Fig. 4 shows a diagram representing the partitioning of data received from an unmanned vehicle and industrial vehicles;

Fig. 5 shows a data storage structure for storing unmanned vehicle data and industrial vehicles;

Fig. 6 shows a method for responding to a query that requests unmanned vehicle data and industrial vehicles;

Fig. 7 shows a diagram representing the generation of a report in response to a query; and

Fig. 8 shows a diagram representing the aggregation of unmanned vehicle data and industrial vehicles.

[0010] It will be noted that throughout the appended drawings, like features are identified by like reference numerals.

DETAILED DESCRIPTION

[0011] The present disclosure provides systems and methods for unmanned vehicle fleet management, such as drones, UAVs (Unmanned Aerial Vehicles), ROVs (Remotely Operated Vehicles), autonomous vehicles and industrial vehicles etc.

[0012] The systems and methods disclosed herein allow for a back-end data platform to manage data collected via unmanned vehicles/robots and can also be applicable to industrial vehicles, as well as one or more front-end applications to facilitate capturing of data from unmanned vehicles and industrial vehicles in the field, uploading of the data to the back-end platform, visualization and reporting of the data, and providing alerts and other useful information to guide decisions based on the data. The back-end data platform components may collectively offer a web or other service platform that can be accessed by the one or more front-end applications respectively providing user interfaces. One of the front-end applications may be a web application. Another front-end application may be associated with a type of mobile device. Application programming interfaces may be publicly available to allow third party integration as well. For industrial vehicles data may be collected for vehicles in a particular location or at multiple locations, for example a mine or construction site.

[0013] The systems and methods disclosed herein allow for the ability to handle data for any type of autonomous robot or drone including land, sea, and air, as well as the ability to handle data for individual robots, components thereof, and/or the aggregate of a large fleet of robots or industrial vehicles. Access to such a large and diverse set of data allows for the determination of emerging trends and complex predictive analytics, while also satisfying securing concerns of large enterprises, government, and military uses, for example.

[0014] Non-sensitive or unrestricted data may be uploaded to a data platform, and accessed/shared with other users. Any sensitive data is always maintained securely on the customer network and is not uploaded to the data platform or accessed/shared with other users external to the customer network. Secure hosting services may also be provided, which host a dedicated data node for a customer that does not want to operate their own customer network, but also does not want certain data mingling with the non-secure cloud data.

[0015] The system disclosed herein is built around two primary domains, with two primary focuses from a user interface perspective.

[0016] A Core Data Domain focuses on the data modelling, storage, distribution, and querying of robot generated data in a granular fashion, while supporting a hybrid cloud approach that enables customers and users to protect a sub-set of their data. A series of components and micro-services may be implemented specifically for dealing with the core data. The Core Data Domain comprises a Universal Data Model (UDM) that provides categorized and distributed columnar key/value storage of data, a granular breakdown of robot and/or sensor data, the characteristics of a key, and the characteristics of a value (including complex value types), and extensibility). The Core Data Domain further comprises Data Repositories, Data Classification (e.g. sensitive, non-sensitive data), and Data Repository Lookup (the location of certain data for a given customer or user). The Core Data Domain further comprises Raw File Ingestion, including decryption and decoding of data, and extraction of metadata. The Core Data Domain still further comprises Data Mapping and Storage, which converts raw data into UDM data and decides where to store it. Data Packets may be used for an efficient serialization of data, for shipping between repositories, and storage. The Core Data Domain also supports a mechanism to query the data stored in the UDM and distributed across multiple repositories. Security is implemented to prevent unauthorized access and/or modifications to various elements of the data. Specific types of objects within the Core Data Domain include vehicle vendors, robot vendors, robot models (models have a schema), components (batteries, motors, props, etc.), and missions (e.g. a collection of data objects, for a given combination of robot and components.

[0017] A Fleet Operations Domain focuses on the operations of robots, or fleets of robots, or vehicles as well as the operations of a business focusing on robots, or fleets of robots. A series of components and micro-services may be implemented specifically for providing fleet operations. The Fleet Operations Domain may provide user management, document management (maintenance records, official certifications, registration/licensing documents, insurance documents, etc.), pilots (stats, permissions, training, certifications, etc.), robots (details, serial numbers, purchase info., maintenance history, warranty information), parts (batteries, props, motors, etc.), assignment (asset management), business operation tools (customers, jobs, contracts, etc.), compliance reporting, and incidents (tracking, investigation, post-mortem analysis, etc.).

[0018] A user interface providing Fleet Operations and Data Acquisition may be implemented. The user interface includes end-user functionality and features that provide: Sign-In / User Management, Management of Pilots, Management of Drones, Management of Parts, Management of Clients/Jobs, Management of Documents, keeping notes on all of the above (as well as general notes), Management of Maintenance Logs, Asset Tracking, Assignment of Equipment/Pilot to a Job, Assignment of Equipment to a Pilot, return-by dates, check-out of equipment by pilots, or driver, (if enabled), manual upload of log files, and possible auto-detection of files based on model choice.

[0019] A user interface providing Data Reporting and Analytics may also be implemented. The user interface allows a user to perform analytics and visualization of data coming from the Core Data Domain, and displaying the data in a way that is relevant to the Fleet Operations Domain. From this perspective, there is end-user functionality and features that provide:

- View Charts/Graphs, Geospatial, and trends of data from:
    - Individual Drones/UVs

    - Individual Components
    - Individual Pilots
    - Individual Flights
    - Aggregated Drones/Components by:
        - Flight
        - Pilot
        - Job
        - Customer
        - Drone/Component
        - Model
        - etc.
    - Aggregated data from global statistics/trends to give greater value/insights
- Custom Derived data based on our own analysis
- Allow users to build “Reports”
- Reports can be viewed manually, or automatically run at intervals
- Notifications and Alerts based on triggers

[0020] Embodiments are described below, by way of example only, with reference to Figs. 1-8.

[0021] Fig. 1 shows a representation of a system for managing a fleet of unmanned vehicles. Data from one or more unmanned vehicles, such as drones 102a and 102b (collectively, unmanned vehicle 102), may be transmitted to server 104a of a Customer Network 150 and/or server 106a of a Global Network 110, such as the Internet. The Customer Network 150 may be restricted and may only be accessed, for example with computer 108a, where a user of the computer 108a has authorization to access the Customer Network 150. The Global Network 110 is a public network that may be accessed by anyone.

[0022] The data may be received at the servers 104a or 106a through a wireless interface such as wireless routers 104b and 106b, respectively, which are operably coupled with the servers. Data may also be sent from the servers 104a and 106b to the unmanned vehicles 102, and data may further be exchanged between the Customer Network 150 and the Global Network 110 through servers 104a and 106a.

The data to/from the unmanned vehicles 102 or industrial vehicle 103 may also be transmitted/received through a direct connection (not shown). The data may be provided by a defined protocol such as for example MAVLink (Micro Air Vehicle Link) or vendor specific link protocol. As depicted in Fig. 1, the transmitted data may be encrypted (such as the data between the drone 102a and the server 104a, or vehicle 102c). The transmitted data may be encrypted for security purposes to prevent unauthorized interception and access of sensitive data. Particularly, transmitted data may be encrypted when being sent to/from the Customer Network 150, for reasoning further described herein.

[0023] A central server 112 running a Data Platform 120 can exchange data through the Global Network 110. For example, the central server may 112 send and receive data to/from the drone 102b through the server 106a, and the central server 112 may send and receive data to/from the drone 102a, or vehicle 103a through the server 104a via the Global Network 110. The Data Platform 120 is configured to store any non-sensitive and/or public data pertaining to unmanned vehicles 102 or vehicles 103. The unmanned vehicles 102 may be any type of land, sea, or air robot and may respectively or collectively belong to individual customers.

[0024] The Data Platform 120 comprises various logical elements representative of functionality that the central server 112 is configured to perform, including but not limited to: Operations 121, User Management 122, Authentication 123, Query Builder 124, and Node Discovery 125. The Data Platform further comprises a Cloud Data Node 130, which itself comprises a Query Aggregator 131, Query Processor 132, Repository 133, Ingestor 134, and UDM Data Warehouse 135. The Data Platform may also comprise Metadata 140, which comprises a Model Schema 141 and Data Privacy Map 142. The Customer Network 150 is a data node that similarly comprises a respective Query Aggregator 151, Query Processor 152, Repository 153, Ingestor 154, and UDM Data Warehouse 155. These logical elements are further described below.

[0025] The UDM (Universal Data Model) Data Warehouse 135, 155, is how all data is stored. This Data Model supports breaking each data element down into individual records of Key->Value, where Value can be a complex data type. (eg: 3

Axis Double Precision Vector, or Radial array, GPS Data, etc). Each Key may belong to any number of Categories. While the individual record is being described as a Key->Value, in fact it's a Key-Value pair, blended with a link to a given "Mission" record, a particular "Robot", and possibly "Components", and will be attached to a timestamp (ie: a voltage reading, might be for a Battery, from flight #192 @ 12:45:45.0214 PM UDT). UDM Data Elements refers to a single data point record in the UDM model. (i.e: a single sensor reading's key->value, at a specific point in time, for a specific flight/pilot/robot/component, etc.). The timestamp precision must be high, as some drones sample upwards of 2000 times per second (or more), meaning that even millisecond precision may not be sufficient.

[0026] As depicted in Fig. 1, the UDM Data Warehouse 155 within the Customer Network 150 may encrypt a subset of the stored data. As noted above, and will be further described herein, unmanned vehicle data received at the Customer Network may contain sensitive data. Encrypting and storing the received data in the UDM Data Warehouse 155 adds a further layer of security to prevent unauthorized access.

[0027] Each Robot/Vehicle Model has a Schema, which is stored as Model Schema 141, 161. This maps UDM Keys to a model of robot, by providing information that identifies, e.g., which sensors does a particular model of drone have, what data does it generate, etc. There may be more than one of a given key, and the Model Schema assigns them unique names such as Sensor0, Sensor1, etc.

[0028] The Data Privacy Map 142 is a mapping of either Keys, or Categories of Keys to either "Sensitive Data", or "Not Sensitive". There may also be other states such as "Obfuscate", for example. Each customer has a Data Privacy Map applied to them.

[0029] The Data Node (e.g. Cloud Data Node 130) is a collective of the following sub-components: Ingestor, Repository, Query Processor, and Query Aggregator. The collection of these sub-components makes a stand-alone Node, which is capable of handling UDM data. A Data Node may run in the cloud (Cloud Data Node 130) or it may be a "Portable Data Node" which runs on-premises at a customer site (e.g. this collection of components is located in the Customer Network 150), in order to ingest, house, and query customer's sensitive as well as non-sensitive data.

[0030] The Ingestor 134, 154 is responsible for receiving raw files being uploaded for processing. These files can be in any format, though plugins may be required for various vendors. A given customer may have multiple Ingestors available (e.g. one in the Data Platform 120, and one on-premises in the Customer Network 150). In the event that a given customer has any sensitive data, all data uploads are redirected by Node Discovery to the On-Premise Node. The Ingestor consumes the raw files, and breaks them up into individual UDM Data Elements. Then using Node Discovery, decides how to separate the data into packets which are submitted either to the public Data Nodes, or to an on-premises Data Node.

[0031] The Repository 133, 153 is a datastore, and is made up of a read and write micro-service. The Repository handles storage and retrieval of UDM Data Elements.

[0032] The Node Discovery Component 125 is used to determine, for a given pilot/flight/customer/etc., whether or not they have sensitive data. The Node Discovery 125 is further configured to determine if the pilot/flight/customer/etc. is running a Data Node on-premises within their respective Customer Network 150, and if so, the location of the Data Node, how many data nodes exist on the Customer Network, etc. The determinations by the Node Discovery 125 help to determine how to partition data that is received by the corresponding pilot/flight/customer/etc. The Node Discovery 125 helps to route data from the client side to the appropriate Ingestor, and is also used by the Ingestors to route Data Packets into the right Repositories.

[0033] The Query Processor 132, 152, is responsible for querying UDM Data Elements from the local Data Node. The Query Processor is only responsible for fulfilling queries for which the local repository actually physically holds the data. If a query comes is received that is for data in which the local Data Node does not possess, the Query Processor 132, 152 returns an error.

[0034] The Query Builder 124 is configured to receive a higher level query for an overall set of UDM Data Elements. When a query is first received, it is not known where the data elements are located. Queries may originate in the client for querying analytics data. The Query Builder 124 uses Node Discovery 125 to determine where the subsets of the required data live. The Query Builder 124 selects an appropriate Query Aggregator to receive the aggregate data, and submits sub-queries to each Query Processor (within their respective Data Nodes). These sub-queries may include a header directing the Query Processor to return it's results to the specified Query Aggregator. The Query Builder 124 may also submit any aggregation stage query clauses to the Query Aggregator, along with the user ID associated with the new query. The Query Builder 124 may also return a response to the client, with a UID of the query, and the Query Aggregator where the client may retrieve the query results.

[0035] The Query Aggregator 131, 151 receives sub-query results from one or more Query Processors, identified with a UID for the client-side query that generated them. Once the Query Aggregator receives all the required parts of the sub-query results, any final aggregation clauses are applied. The sub-query results may, for example, be aggregated into a report in response to the query. The Query Aggregator may hold the final results for the client to request, and the client can then request the query results for a given UID directly from the Query Aggregator and the final results will be returned to the client for display/visualization. Alternatively, the query results and/or report may be pushed to the client, for example, through a user interface displayed on the computer 108a.

[0036] The Operations Component 121 comprises a cluster of micro-services, handling operational data such as: Pilots, Customers, Jobs, Missions, Documents, Notes, Assignment/Asset Tracking, etc. The functionality of the Operations Component 121 may leverage some data points from the UDM, however for the most part these services are stand-alone, and operate in a more conventional relational database management system (RDBMS) type back-end.

[0037] The Authentication Component 123 handles multi-tenant authentication of users from the client-side. This service may simply be implemented as a user login service. The Authentication Component 123 supports authorization of users of different roles. For example, some users may be pilots of UVs, while other users may be administrators. Users may also belong to more than one organization. For example a freelance pilot may belong to 3-4 different enterprises. In that case during authentication the user may be presented an option to choose which organization he is signing into.

[0038] The User Management Component 122 comprises one or more micro-services that handles activities such as: User Creation, User Deletion, Password Resets, inviting a user to an organization (e.g. enterprise admins could invite users by email to join their organization, a “Pilot” may belong to more than one organization), User Rights Management/Role Management, etc.

[0039] Some elements depicted as being part of the Data Platform 120 may instead, or in addition, be included in the Customer Network 150. For example, the Customer Network 150 shown in Fig. 1 is also depicted as comprising Particular Metadata 160, which itself comprises a Particular Model Schema 161. Having a Particular Model Schema 161 located within the Customer Network 150 may be implemented, for example, if the models of unmanned vehicles operated by the owner of Customer Network 150 are known and finite. Again, the system depicted in Fig. 1 may be implemented in various ways as would be appreciated by a person skilled in the art without departing from the scope of this disclosure.

[0040] The system depicted in Fig. 1 shows an exemplary system only, and the above logical elements are exemplary used to describe the functionality of the system. Modifications to the logical elements, including adding, combining, or removing logical elements in an implementation of the system that meets the functionality requirements of the system would be readily apparent to a person skilled in the art without departing from the scope of this disclosure. Some of the functionality requirements of the system are outlined below.

[0041] One functionality requirement is to provide a data platform configured to consume raw log files that have no “standard format”. The UDM may be created in a manner that can accommodate new data types from new vendors without re-write. Data ingested from raw logs are converted into elements that satisfy UDM. The UDM supports granular data storage (e.g. columnar), and each data value (column/field) in the data, must be able to be stored independently of each other.

[0042] The UDM may be implemented as a Key->Value Pair, with Keys established for each main data type, and values being able to support complex data types. For example:

- “Accelerometer” -> X,Y,Z Vector data value
- “Magnetometer” -> X,Y,Z Vector data value
- “Gyroscope” -> X,Y,Z Vector data value
- “Temperature” -> Single temperature value
- “GPS” -> Complex GPS Data type, eg:
  - ➢ GPSData(Lat,Long,Altitude,Sats,Quality,Confidence)
- “Radar Ranger” -> Single Distance Value
- “Radar Sweep” -> Array of Radial Distance Values
- “Visual Obstacle Avoidance” -> 2D Depth Map
- “LIDAR” – 2D Depth Map, or Spherical Array of Distance Values

[0043] The UDM may have a concept of “Derived Keys” which are calculated through some transformation of other keys. This may be represented generically, such as in an extensible API, which allows for new derived data to be easily developed. Derived Keys can similarly be identified as sensitive or non-sensitive. If a derived key includes any sensitive data for a customer, the derived key automatically becomes sensitive itself. The generation of derived key values may be a part of a scheduled job which runs on the data repository containing the derived key data. If for example a derived key contains one sensitive field, and does some calculation against a non-sensitive field, the data repository will need to query the other repo in the cloud which holds the non-sensitive field, and do the transformation as an async job.

[0044] The UDM may have a concept of a “Data Category”, which contains one or more keys. A key may belong to more than one category. For example:

- Positional Data Category
  - ➢ GPS Data
- IMU Data Category
  - ➢ Accelerometer
  - ➢ Gyro
  - ➢ Magnetometer

  - ➢ GPS
  - ➢ Visual
  - ➢ LIDAR
- Visual Sensors Category
  - ➢ Visual
  - ➢ Radar
  - ➢ LIDAR
  - ➢ etc.

[0045] The UDM may have a concept of a “robot model” which defines a robot master schema. The master schema defines which keys belong to a given vehicle. As noted above, a vehicle may have more than one of each key type. For each key mapping to the schema, the column is given a unique name, i.e.: Accelerometer0, Accelerometer1, Gyroscope0, Gyroscope1, etc.

[0046] Another requirement of the UDM is for it to be extensible. Plugins to support new value types may be implemented.

[0047] The system is configured for data sensitivity to be decided on a per-customer basis. Data can be flagged as either sensitive or not-sensitive. Data can be flagged on either a category, or individual key.

[0048] If a customer wishes to store data on their own network, for example to store sensitive data, the customer is able to have their own data storage repository that runs on-premises, including the raw-log ingestor component.

[0049] A system requirement is also to provide a mechanism for the discovery of data repositories, and if a customer has any sensitive data, then all logs are sent to an on-premises ingestor, rather than to the cloud. The data repository of which the data should be sent to may be determined at the time of data collection on a per-key value when ingested if it’s stored in on-premises repository or in a cloud. The Keys/Values are bundled into packets, and shuttled to the appropriate storage repositories during ingestion.

[0050] The data repository is a component which stores a given subset of key/value pairs and is a long-term datastore. The data repository comprises a back-end storage platform (such as a database), a microservice for submission of new key/value packets, and a microservice for retrieval/query of the repository for data. The ingestor is a component which receives the raw data logs and ingests them into the UDM format. The ingestor is able to determine the vehicle model from the files being submitted, with the detection process/requirements varying based on vendor. The ingestor accepts robot model plugins, which will allow it to interpret/decode a given robot model's raw log format into UDM key/value pairs.

[0051] An App Layer Datastore provides an additional relational datastore with one or more of the following concepts: users/pilots; robots, batteries, components/parts; payloads; customers, jobs, etc.; flights/missions/etc. Any sensor data lives in a different datastore and is tied to a flight/operation that uses a particular unmanned vehicle and/or components that in turn define the sensor data that may be recorded.

[0052] Fig. 2 shows a representation of a system used for generating reports for a fleet of unmanned vehicles. The system of Fig. 2 further highlights the security of sensitive data as described with reference to Fig. 1.

[0053] Data 202a from a customer's drone 102a is received at the server 104a on the Customer Network 150 through the wireless router 106a. As depicted, some of the data 202a received from the drone 102a may be sensitive and/or encrypted. As described with reference to Fig. 1 and will be further described herein, the sensitive elements of the data received from the drone 102a is stored within the Customer Network 150 at the server 104a, and non-sensitive data may be sent to the central server 112 through the Global Network 110 for storage in the Data Platform 120, and/or the non-sensitive data may also be stored within the Customer Network 150 at the server 104a.

[0054] Data 202b from a customer's drone 102b is received at the server 104b on the Customer Network 150 through the wireless router 106b. As depicted, all of the data 202b received from the done 102b is non-sensitive. As described with reference to Fig. 1 and will be further described herein, the non-sensitive may be sent to the central server 112 through the Global Network 110 for storage in the Data

Platform 120, and/or the non-sensitive data may also be stored within the Customer Network 150 at the server 104a.

[0055] If a pilot/customer wishes to access data and/or perform data analytics, the pilot/customer is restricted based on the security levels and authorization requirements associated with the data. For example, if a user of computer 108a on the Customer Network 150 wishes to access data and/or perform data analytics, the user, in accordance with any additional security requirements, will have access to data and reports 204a containing both non-sensitive data from several pilots/customers (both internal and external to the Customer Network 150) and sensitive data particular to the pilot/customer that sends data to the Customer Network 150. If a user of computer 108b that connects to the central server 112 through the Global Network 110 wishes to access data and/or perform data analytics, that user only has access to data and reports 204b containing non-sensitive data from the various pilots/customers. A collection of reports 204c containing non-sensitive data may also be stored and accessible at the central server 112. A further description of the method for generating reports is provided in Figs. 6 and 7.

[0056] It is noted that the systems depicted in Figs. 1 and 2 are simplified for explanatory purposes, and in implementation the system may support several Customer Networks 150 each configured to receive data from one or more unmanned vehicles, and the data received at the server 104b on the Global Network 110 is similarly configured to receive data from one or more vehicles.

[0057] Fig. 3 shows a method for partitioning data received from one or more unmanned vehicles. The method 300 may be performed, for example, by the server located within the Customer Network 150.

[0058] Unmanned vehicle (UV) or industrial vehicle data is received at the server (302). The received UV data is raw data received from an UV owned/operated by an owner/user of a data node within the customer network. The raw UV data may comprise one or both of sensitive data and non-sensitive data. The raw data is converted to a plurality of universal data elements in accordance with the UDM (304). Each universal data element may, for example, indicate the UV from which the raw data was received, a sensor reading of the UV, and a time that the sensor reading occurred. The conversion of raw data to the universal data elements may be performed, for example, by the Ingestor within the Data Node and based on Model Schema for that UV or vehicle type.

[0059] The universal data elements may be divided (306), for example into sensitive data and non-sensitive data. Other states may also be available, such as obfuscate, etc. Dividing the data into states may be performed by the Ingestor 154, for example by using the Node Discovery 125 and Data Privacy Map 142. The division of sensitive and non-sensitive data may be implemented in several ways without departing from the scope of this disclosure. For example, as will be described with reference to Fig. 5 below, both the non-sensitive data and sensitive data may be stored in the UDM 155 within the customer network, where the non-sensitive data and sensitive data may be divided based on their storage location within the UDM. In other embodiments, an index and/or label indicating the data sensitivity may be generated and associated with the data. Non-secure data is transmitted (308), for example to a Repository 133 in the Data Platform 120, to a Repository in a different Data Node, or to a Query Aggregator. The non-sensitive data may be transmitted while also storing the data in the UDM 155 of the Customer Network 150. The non-sensitive data may be transmitted and not stored in the UDM 155 of the Customer Network 150. The non-sensitive data may only be transmitted when requested, such as during a query building process as further described with reference to Fig. 6 below. The sensitive data is encrypted and securely stored (310) in the UDM 155 of the Customer Network 150.

[0060] Fig. 4 shows a diagram representing the partitioning of data received from an unmanned vehicle or industrial vehicle. Raw data is received from the vehicle 102, for example a drone, at the Ingestor 154 of the Customer Network (401). The Ingestor accesses the Model Schema 141, 161 (403), and uses the mapping of the Model Schema to create UDM elements (405). The Model Schema may be particular to the type or model of UV or vehicle from which the raw data was received from, and the Ingestor may first need to identify or determine the type or model of the UV. The Ingestor determines if the data is sensitive or non-sensitive by communicating with the Node Discovery 125 (407). The Node Discovery determines how to divide the data by accessing a Data Privacy Map 142 (409).

[0061] Sensitive data is sent from the Ingestor 154 to a Repository 153 within the Customer Network (411), and the sensitive data is written from the Repository 153 to the UDM 155 within the Customer Network (413). The process of writing the sensitive data to the UDM within the Customer Network may comprise encrypting the sensitive data. The encryption of the sensitive data may be in accordance a security level required to access the data, which may be set forth in the Data Privacy Map 142. Non-sensitive data is sent from the Ingestor 154 to a non-secure Repository 133 (415), such as within the Data Platform 120, and the non-sensitive data is written by the non-secure Repository to the UDM 135 of the Data Platform 120 (417).

[0062] The diagram shown in Fig. 4 represents one possible embodiment for how the data is partitioned, and variations may exist without departing from the scope of this disclosure. For example, the locations of the metadata comprising the Model Schema and Data Privacy Map may be one or both of the Customer Network and the Data Platform. In another embodiment, the Customer Network may comprise a plurality of Data Nodes and a Node Discovery element may be located within the Customer Network. In yet another embodiment, as previously described, the non-sensitive data may alternatively or additionally be stored within the UDM of the Customer Network.

[0063] Fig. 5 shows a data storage structure for storing vehicle data. The data storage structure shown in Fig. 5 represents UDM 154 within the Customer Network 150. Shown within the UDM 155 is an exemplary data table 502. The data table 502 is divided into columns containing non-sensitive data and columns containing sensitive data. Each row in the data table 502 may correspond to a universal data element.

[0064] As previously described, each universal data element may indicate an identifier for the UV associated with the data element, a sensor reading from the UV, and a time that the sensor reading occurred. As depicted, the vehicle identifier may be considered non-sensitive information, or may be represented by a hash. The time at which a sensor reading occurs may also be considered non-sensitive information, and stored for example in column X. Depending on the type of sensor reading, the sensor reading may be considered non-sensitive information and stored for example in column Y, or may be considered sensitive information and stored for example in column Z.

[0065] As described above, the determination of whether or not data is sensitive can be based on a Data Privacy Map established by the owner/user of the vehicle and/or Data Node. For example, the Data Privacy Map may establish that a sensor reading corresponding to the speed of a drone, or a sensor reading corresponding to a voltage reading of a drone battery, may be non-sensitive data. On the other hand, the Data Privacy Map may establish that a sensor reading corresponding to the GPS of a drone, or a payload of the drone, may be sensitive data. In an example of data that is received from military drones, it may be undesirable for data indicating where the drone is flying, and data possibly indicative of weaponry the drone is equipped with, to be made available to the public and hence the GPS data and payload may be classified as sensitive data.

[0066] The data table 502 contains, as a whole, all universal data elements for a fleet of vehicles associated with the Customer Network. When transmitting data to the Data Platform or Query Aggregator of another Data Node, the data table 502 facilitates transmitting only the non-sensitive data according to the columns that are identified as storing non-sensitive data. For further security, the sensitive data may be encrypted when stored in the UDM 155 so that it cannot be accessed by unauthorized users. In comparison, the UDM 135 of the Data Platform 120 would contain only non-sensitive data.

[0067] The data table 502 is shown for exemplary purposes, and may be implemented in numerous ways without departing from the scope of this disclosure. As previously described, additional or alternative states for classifying data may be envisioned. For example, it may be desirable for the ID of the vehicle to be obfuscated. In this regard, an obscure ID may be transmitted, thus making it unidentifiable as to where the data came from (e.g., it cannot be determined that the data element is for a military drone), but the non-secure data can still be effectively analyzed (e.g., by examining several voltage readings taken for the battery of Drone1234, it may be determined that the battery-life of such drone type is worse than other models of drones, but the general identifier provides no indication that the drone is a military drone).

[0068] Fig. 6 shows a method for responding to a query that requests unmanned vehicle data. In some embodiments, the method 600 may be performed by the server 104a located within the Customer Network 150. In other embodiments, the method 600 may be performed by the server 106a providing the Data Platform 120. Querying the Data Platform may, for example, be performed by a recreational user who does not have their own Customer Network with a Data Node.

[0069] A query is received at the server (602). The query may originate, for example, from a user of a web or mobile application supported by the system, and is received at the Query Builder. The query may, for example, request analytics data, and the query may have a security level associated therewith. The security level may be determined from login credentials of the user, the location from which the query originated, etc. Based on the security level associated with the query and the analytics data requested by the query, the locations of nodes that have data relevant to the query and where the security level of the query complies with the security level and sensitivity of the data stored thereon is determined (604). This determination may be performed using Node Discovery and the Data Privacy Map for the respective nodes. If no data elements are available for responding to the query (for example, the security level associated with the query is below the security level of the data that is being queried), an error is returned.

[0070] A sub-query is generated for the determined nodes (606) that requests the data elements relevant for the query. The sub-queries also indicate a Query Aggregator to which the data elements should be sent to. The specified Query Aggregator is generally located within the Data Node from which the query was received/originated, but does not have to be limited to such. These sub-queries are sent to the respective nodes (608) determined to contain data relevant to the query, with the sub-queries being sent from the Query Builder to respective Query Processors.

[0071] Depending on the query and the locations of the nodes having data elements relevant to the query, the number of sub-queries and the locations of nodes that the sub-queries are sent to may greatly vary. For example, military personnel with a high level of security clearance may query for sensitive data from their fleet of UVs. The query would be received at the server on the Customer Network and, determining that the relevant data is only stored in the UDM on the Customer Network and that the security level associated with the query is compliant with the sensitivity of the data being requested, the sub-query may only be sent to the Query Processor at the Data Node in the Customer Network.

[0072] In another example, the same military personnel may wish to assess non-sensitive data, such as in order to determine what model of drone has the longest flying range before needing to be recharged. The query would be received at the server on the Customer Network and the Query Builder determines all of the nodes that have data relevant to the query. If all non-sensitive data is pushed to the Data Platform, the sub-query might only be sent to the Query Processor in the Data Platform. If some nodes have relevant non-sensitive data stored in the local UDM, the sub-query may also be generated and sent to Query Processors on different networks.

[0073] The Query Processor receives a sub-query, and the local UDM is queried for the relevant data elements (610). If a sub-query is received for data elements that the local UDM does not possess, the Query Processor returns an error. The data elements are output (612) from the respective Query Processors to the Query Aggregator that was identified in the sub-query. The Query Aggregator receives (614) and aggregates (616) the data elements. The data elements may be aggregated into a report that facilitates analysis of the retrieved data. The user from which the query first originated may be provided with a location of the Query Aggregator from which to view the aggregated data/report.

[0074] Fig. 7 shows a diagram representing the generation of a report in response to a query. The diagram in Fig. 7 is representative of one scenario for a query requesting both non-sensitive and sensitive data, and where the user is an authorized user (i.e., a user that has a security level appropriate to access the sensitive data). However, as described with reference to Fig. 6, various types of queries may exist and various implementations for storing the data may exist, and therefore the diagram shown in Fig. 7 represents only one embodiment and variations of such may be implemented without departing from the scope of this disclosure.

[0075] A query is received at the Query Builder 155 (701). Note that the Query Builder 155 is not shown in Fig. 1, however as previously noted several logical elements that are in the Data Platform 120 may similarly be located within the Customer Network 150, allowing for various implementations of the system to be possible without affecting the functionality of the system. The Query Builder 155 determines data elements that are relevant to the query and the locations of nodes where the relevant data elements are stored based on the Node Discovery 125 (703). The Node Discovery 125 may access the Data Privacy Map 142 (705) to determine the sensitivity of data elements for different customers, thereby identifying the non-sensitive data elements of those customers that may be relevant to the query and retrieved. Based on the determination of node locations where the relevant data elements are stored, the Query Builder 155 generates a sub-query or sub-queries (707). In the example of this diagram, all of the sensitive data to be retrieved is stored in a single UDM 155 within the Customer Network 150, and all of the non-sensitive data to be retrieved is stored in the UDM 135 of the Data Platform 120.

[0076] The sub-query for sensitive data is sent from the Query Builder 155 to the Query Processor 152 on the secure Customer Network 150 (709). The Query Processor 152 queries the relevant sensitive data elements from the local UDM 155 and retrieves the data elements (711). The sensitive data elements are sent from the Query Processor 152 to the Query Aggregator 151 (713).

[0077] The sub-query for non-sensitive data is sent from the Query Builder 155 to the Query Processor 132 on the non-secure Data Platform 120 (715). The Query Processor 132 on the Data Platform 120 queries the relevant non-sensitive data elements from the UDM 135 of the Data Platform 120 and retrieves the data elements (717). The non-sensitive data elements are sent from the Query Processor 132 to the Query Aggregator 151 (719).

[0078] Having received the data elements, the Query Aggregator 151 aggregates the data and generates a report in response to the query (721). The report is accessed by / sent to the user from which the query originated (723).

[0079] Fig. 8 shows a diagram representing the aggregation of unmanned vehicle data. Continuing with the exemplary query received in Fig. 7 (i.e. a query requesting both non-sensitive and sensitive data, and where the user is an authorized user), the UDM 802 corresponds to a UDM on the Customer Network from which the query was received, and the UDM 804 corresponds to the UDM on the Data Platform. For exemplary purposes, further consider that there is data relevant to the query and which is stored in a UDM 806 of a Data Node on a separate Customer Network. In accordance with the security restrictions established by the customers/owners through their Data Privacy Map, a report 810 is generated in response to the query.

[0080] The relevant data element within the UDM 802 contains both sensitive and non-sensitive information, all of which is provided in the report 810 because the user from which the query originated has access / security clearance to view the sensitive data.

[0081] The relevant data element within the UDM 804 only contains non-sensitive information. The unmanned vehicle data is identified as belonging to “Tim”, and Tim does not mind sharing all of his UV data, nor having his data being identifiable as belonging to him, so Tim uploads all of his data to the Data Platform and places no restrictions on what may be shared in analytics reports. All of the data within the data element from the UDM 804 is provided in the report 810.

[0082] The relevant data element within the UDM 806 contains non-sensitive information, sensitive information, and information which is designated to be obfuscated. In this case, the identifier indicating the vehicle associated with the data element is designated to be obfuscated. Perhaps the vehicle belongs to a corporation which is concerned that the identifier of their unmanned vehicles may be identifiable as belonging to them, and the corporation does not wish to make public how many drones are in their fleet, for example. The identifier is not obfuscated in the UDM 806, because an employee at the corporation may wish to access the UDM and uniquely identify the different vehicles. However, when the data is provided in the report 810, the identifier has been obfuscated and now is output as “123”. The non-sensitive data from the data element in UDM 806 is also provided in the report, however no sensitive data is provided from the UDM 806.

[0083] Although certain components and steps have been described, it is contemplated that individually described components, as well as steps, may be combined together into fewer components or steps or the steps may be performed sequentially, non-sequentially or concurrently. Further, although described above as occurring in a particular order, one of ordinary skill in the art having regard to the current teachings will appreciate that the particular order of certain steps relative to other steps may be changed. Similarly, individual components or steps may be provided by a plurality of components or steps. One of ordinary skill in the art having regard to the current teachings will appreciate that the system and method described herein may be provided by various combinations of software, firmware and/or hardware, other than the specific implementations described herein as illustrative examples.

[0084] It is understood that the specific order or hierarchy of steps in the processes disclosed is an example of exemplary approaches. Based upon design preferences, it is understood that the specific order or hierarchy of steps in the processes may be rearranged while remaining within the scope of the present disclosure. The accompanying method claims present elements of the various steps in a sample order, and are not meant to be limited to the specific order or hierarchy presented.

[0085] Each element in the embodiments of the present disclosure may be implemented as hardware, software/program, or any combination thereof. Software codes, either in its entirety or a part thereof, may be stored in a computer readable medium or memory (e.g., as a ROM, for example a non-volatile memory such as flash memory, CD ROM, DVD ROM, Blu-ray™, a semiconductor ROM, USB, or a magnetic recording medium, for example a hard disk). The program may be in the form of source code, object code, a code intermediate source and object code such as partially compiled form, or in any other form.

[0086] It would be appreciated by one of ordinary skill in the art that the system and components shown in Figures 1-8 may include components not shown in the drawings. For simplicity and clarity of the illustration, elements in the figures are not necessarily to scale, are only schematic and are non-limiting of the elements structures. It will be apparent to persons skilled in the art that a number of variations and modifications can be made without departing from the scope of the invention as defined in the claims.

CLAIMS:

1. A method of partitioning data received from one or more vehicles, comprising:

receiving vehicle raw data from an vehicle of the one or more vehicles, the vehicle raw data comprising both sensitive and non-sensitive data;

converting the vehicle raw data into a plurality of universal data elements, each data element indicating the vehicle from which the raw data was received, a sensor reading of the vehicle, and a time that the sensor reading occurred;

determining, from a data privacy map established by an owner of the vehicle, which of the universal data elements are sensitive data elements and which of the universal data elements are non-sensitive data elements;

transmitting the non-sensitive data elements to an external device; and

storing the sensitive data elements locally.

2. The method of claim 1, wherein the data privacy map further indicates which of the universal data elements are to be obfuscated, the method further comprising:

storing the universal data elements that are indicated to be obfuscated locally;

obfuscating the universal data elements that are indicated to be obfuscated; and

transmitting the obfuscated data elements to the external device.

3. The method of claims 1 or 2, wherein converting the vehicle raw data into a plurality of universal data elements comprises:

identifying a model of the vehicle from which the vehicle raw data was received;

accessing a pre-determined mapping particular to the model of the vehicle, the mapping indicating how to convert the vehicle raw data for the model of the vehicle to the plurality of universal data elements.

4. The method of claim 3, wherein the universal data elements have the same format for all models of the vehicle.

5. The method of any one of claims 1 to 4, further comprising storing the non-sensitive data elements locally.

6. The method of claim 5, wherein the non-sensitive data elements are transmitted to the external device in response to a query from the external device.

7. The method of any one of claims 1 to 6, further comprising transmitting the sensitive data elements to a device within a local network.

8. The method of any one of claims 1 to 7, wherein the data privacy map further indicates a security level required to access the sensitive data.

9. The method of any one of claims 1 to 8, wherein the vehicle raw data received from the vehicle is encrypted, and the method further comprises decrypting the vehicle raw data.

10. The method of any one of claims 1 to 9, wherein storing the sensitive data elements locally comprises encrypting the sensitive data elements.

11. The method of any one of claims 1 to 10, wherein the vehicle is an unmanned aerial vehicle.

12. The method of any one of claims 1 to 10, wherein the vehicle is an industrial vehicle.

13. A method of responding to a query requesting vehicle data, comprising:

receiving, from a user, the query requesting vehicle data, the query having a security level associated therewith;

determining, based on the security level associated with the query and the vehicle data requested by the query, vehicle data elements that provide the vehicle data requested and that correspond to the security level associated with the query;

determining one or more nodes storing the vehicle data elements;

generating a sub-query request for each of the one or more nodes, the sub-query request indicating the vehicle data elements to be retrieved at the respective node, and an indication of an output node to which the vehicle data elements should be sent to;

sending the sub-query requests to each of the one or more nodes;

receiving the requested vehicle data elements from the one or more nodes; and

aggregating the vehicle data elements to provide a report.

14. The method of claim 13, wherein the vehicle data requested by the query is one or both of sensitive and non-sensitive.

15. The method of claims 13 to 14, wherein a sub-query for the sensitive vehicle data is sent to a node on a local network.

16. The method of any one of claims 13 to 15, wherein a sub-query for the non-sensitive vehicle data is sent to one or both of a node on a local network and a node on an external network.

17. The method of any one of claims 13 to 16, further comprising:

querying a local data storage for the requested vehicle data elements comprising one or both of the sensitive data and the non-sensitive data; and

outputting the vehicle data to the output node.

18. The method of any one of claims 13 to 17, wherein the security level is based on log-in credentials of the user.

19. The method of any one of claims 13 to 18, further comprising:

receiving an error from a node of the one or more nodes in response to the sub-query when the data requested from the node is not available.

20. The method of any one of claims 13 to 19, further comprising:

returning an error when there are no vehicle data elements on any node that provide the vehicle data requested and that correspond to the security level associated with the query.

21. The method of any one of claims 13 to 20, wherein the query is received through a user interface of a supported application.

22. The method of claim 21, wherein the report is presented in the user interface.

23. The method of any one of claims 13 to 22, further comprising:

storing the report locally at the output node; and

notifying the user of the output node for which an aggregation of the vehicle data elements is stored.

24. The method of any one of claims 13 to 23, wherein the vehicle is an unmanned aerial vehicle.

25. The method of any one of claims 13 to 24, wherein the vehicle is an industrial vehicle.

26. A system for partitioning data received from one or more vehicles and responding to a query requesting vehicle data, comprising:

a first server hosting a customer network; and

a second server hosting a cloud-based data platform operably coupled with the first server through a data network;

wherein the first server is configured to:

receive vehicle raw data from an vehicle of the one or more vehicles that is operably coupled with the customer network, the vehicle raw data comprising both sensitive and non-sensitive data;

convert the vehicle raw data into a plurality of universal data elements, each data element indicating the vehicle from which the raw data was received, a sensor reading of the vehicle, and a time that the sensor reading occurred;

determine, from a data privacy map established by an owner of the vehicle, which of the universal data elements are sensitive data elements and which of the universal data elements are non-sensitive data elements;

transmit the non-sensitive data elements to the second server through the data network; and

store the sensitive data elements locally.

27. The system of claim 26, wherein the second server is configured to:

receive the non-sensitive data elements from the first server; and

store the non-sensitive data elements locally.

28. The system of claim 27, wherein the second server is configured to:

receive vehicle raw data from an vehicle of the one or more vehicles that is not connected to the customer network, the vehicle raw data comprising only non-sensitive data;

convert the vehicle raw data into a plurality of universal data elements, each data element indicating the vehicle from which the raw data was received, a sensor reading of the vehicle, and a time that the sensor reading occurred; and

store the non-sensitive data elements locally.

29. The system of claim 27, wherein the first server is further configured to:

receive a query from a user of a device on the customer network requesting vehicle data, the query having a security level associated therewith;

determine, based on the security level associated with the query and the vehicle data requested by the query, vehicle data elements that provide the vehicle data requested and that correspond to the security level associated with the query;

generate a sub-query request for each of the one or more nodes, the sub-query request indicating the vehicle data elements to be retrieved at the respective node, and an indication of an output node to which the vehicle data elements should be sent to;

send the sub-query requests to each of the one or more nodes;

receive the requested vehicle data elements from the one or more nodes; and

aggregate the vehicle data elements to provide a report.

30. The system of claim 26, wherein the vehicle data requested by the query is one or both of sensitive and non-sensitive, and wherein the first server is further configured to:

send a sub-query request for sensitive vehicle data to a node within the customer network; and

send a sub-query request for non-sensitive vehicle data to a node external to the customer network,

wherein the node external to the customer network is either located within the data platform hosted by the second server or is located within a separate customer network.

31. The system of claim 26, wherein the second server is further configured to:

receive a query from a user of a device on the data network requesting vehicle data;

determine, based on the vehicle data requested by the query, vehicle data elements that provide the vehicle data requested and that correspond to the security level associated with the query;

generate a sub-query request for each of the one or more nodes, the sub-query request indicating the vehicle data elements to be retrieved at the respective node, and an indication of an output node to which the vehicle data elements should be sent to;

send the sub-query requests to each of the one or more nodes;

receive the requested vehicle data elements from the one or more nodes; and

aggregate the vehicle data elements to provide a report.

32. The system of claim 31, wherein the vehicle data requested by the query is non-sensitive and wherein there is no security level associated with the query, and wherein the second server is further configured to:

send a sub-query request for non-sensitive vehicle data to one or both of a node within the customer network and a node within the data platform.

33. The system of any one of claims 26 to 32, wherein the vehicle is an unmanned aerial vehicle.

34. The system of any one of claims 26 to 32, wherein the vehicle is an industrial vehicle.

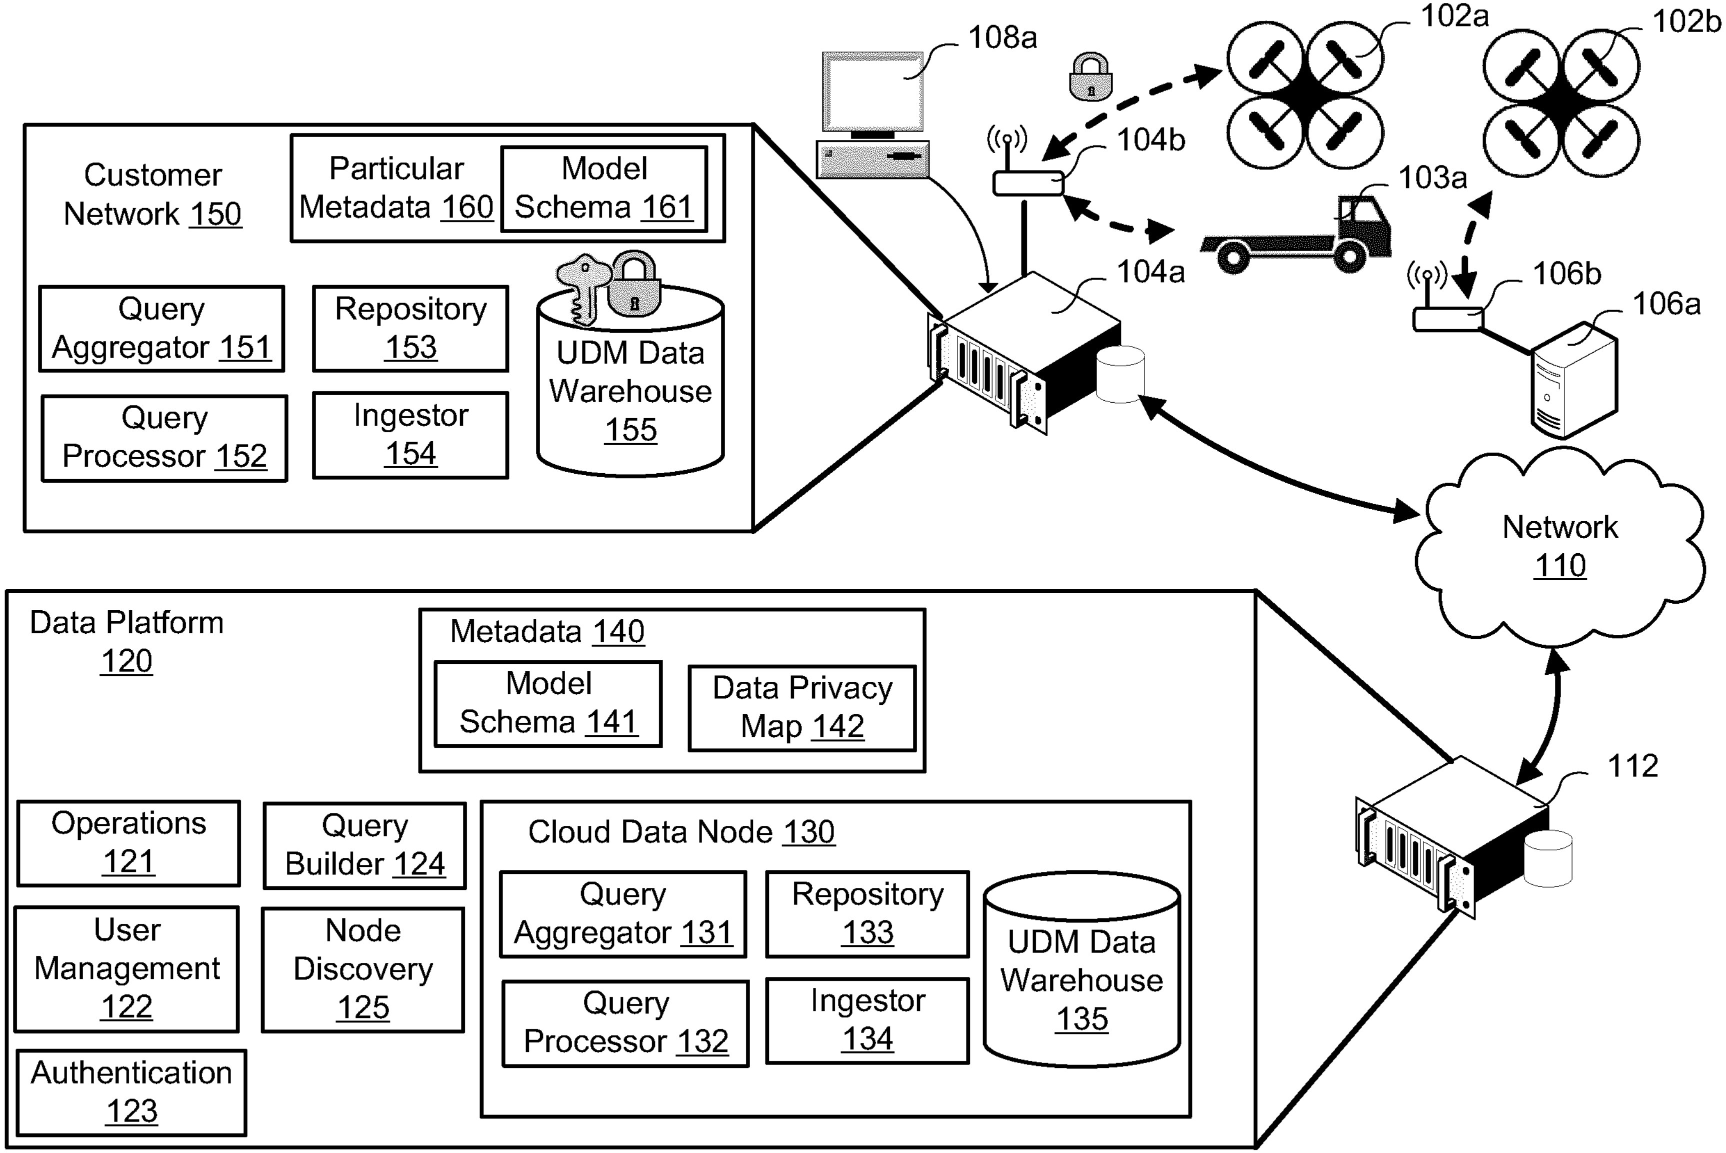
Customer Network 150
Particular Metadata 160
Model Schema 161
Query Aggregator 151
Repository 153
UDM Data Warehouse 155
Query Processor 152
Ingestor 154
108a
104b
102a
102b
103a
104a
106b
106a
Network 110
112
Data Platform 120
Metadata 140
Model Schema 141
Data Privacy Map 142
Operations 121
Query Builder 124
Cloud Data Node 130
Query Aggregator 131
Repository 133
UDM Data Warehouse 135
User Management 122
Node Discovery 125
Query Processor 132
Ingestor 134
Authentication 123

Fig. 1

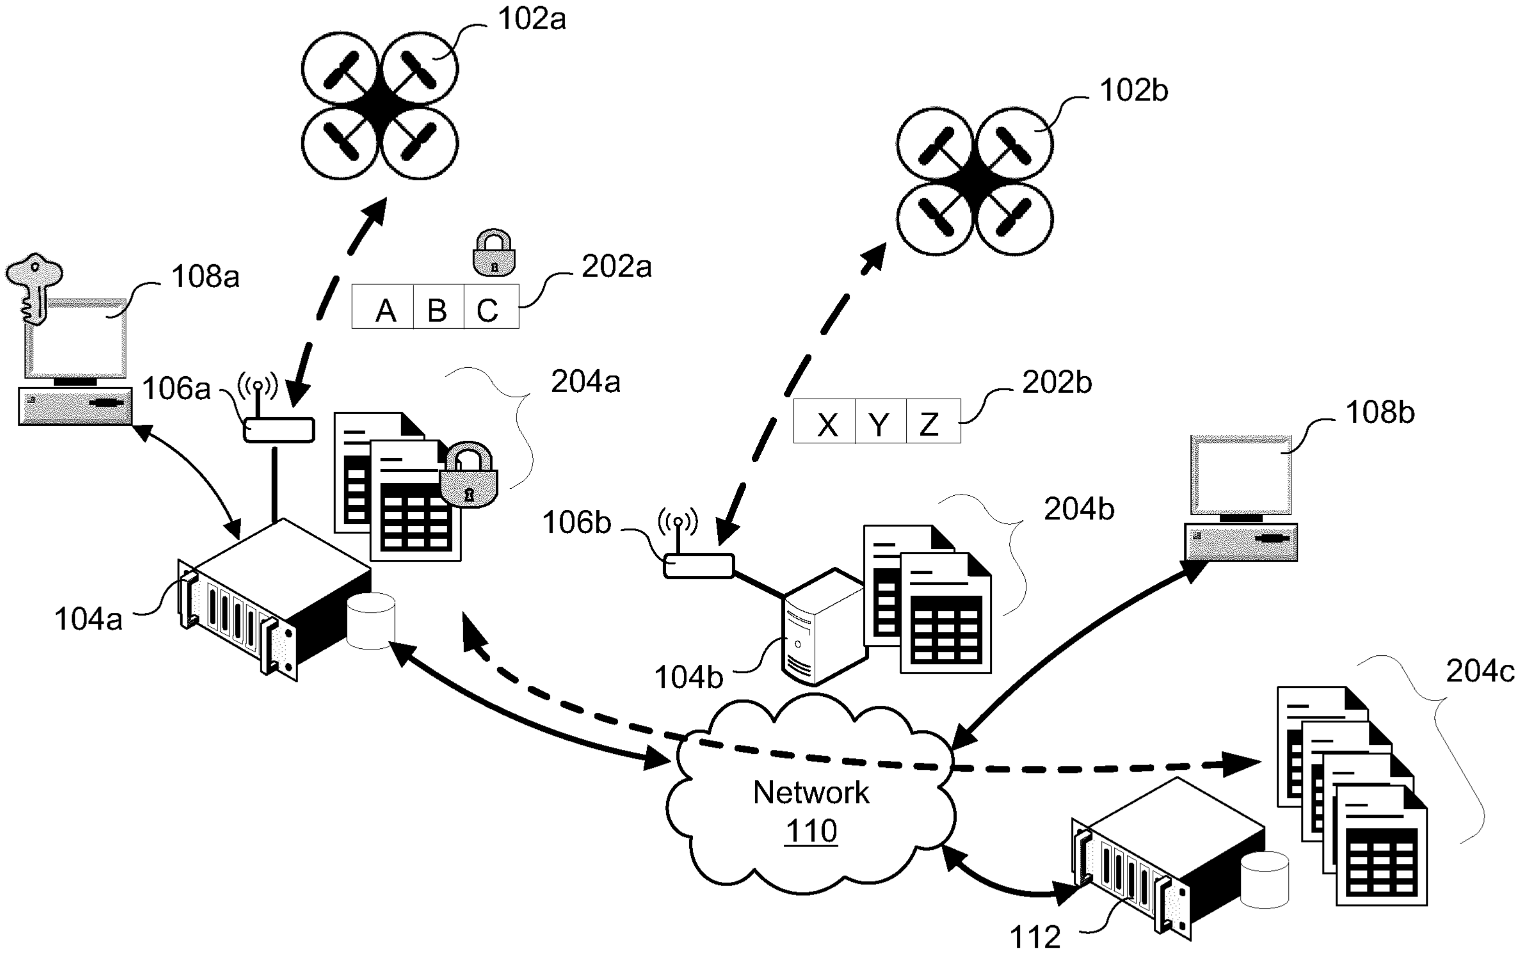
102a
102b
108a
A B C
202a
106a
204a
X Y Z
202b
108b
106b
204b
104a
104b
204c
Network
110
112

FIG. 2

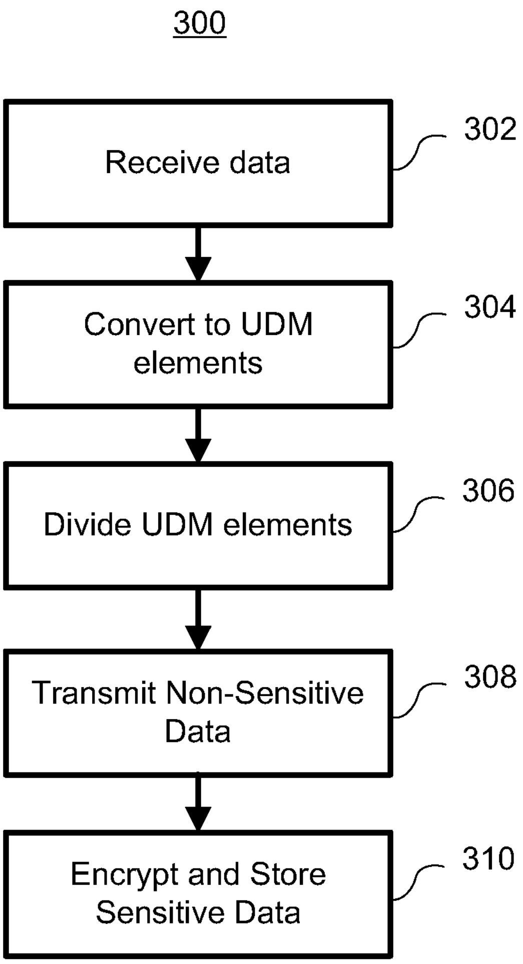
300
Receive data
302
Convert to UDM elements
304
Divide UDM elements
306
Transmit Non-Sensitive Data
308
Encrypt and Store Sensitive Data
310

Fig. 3

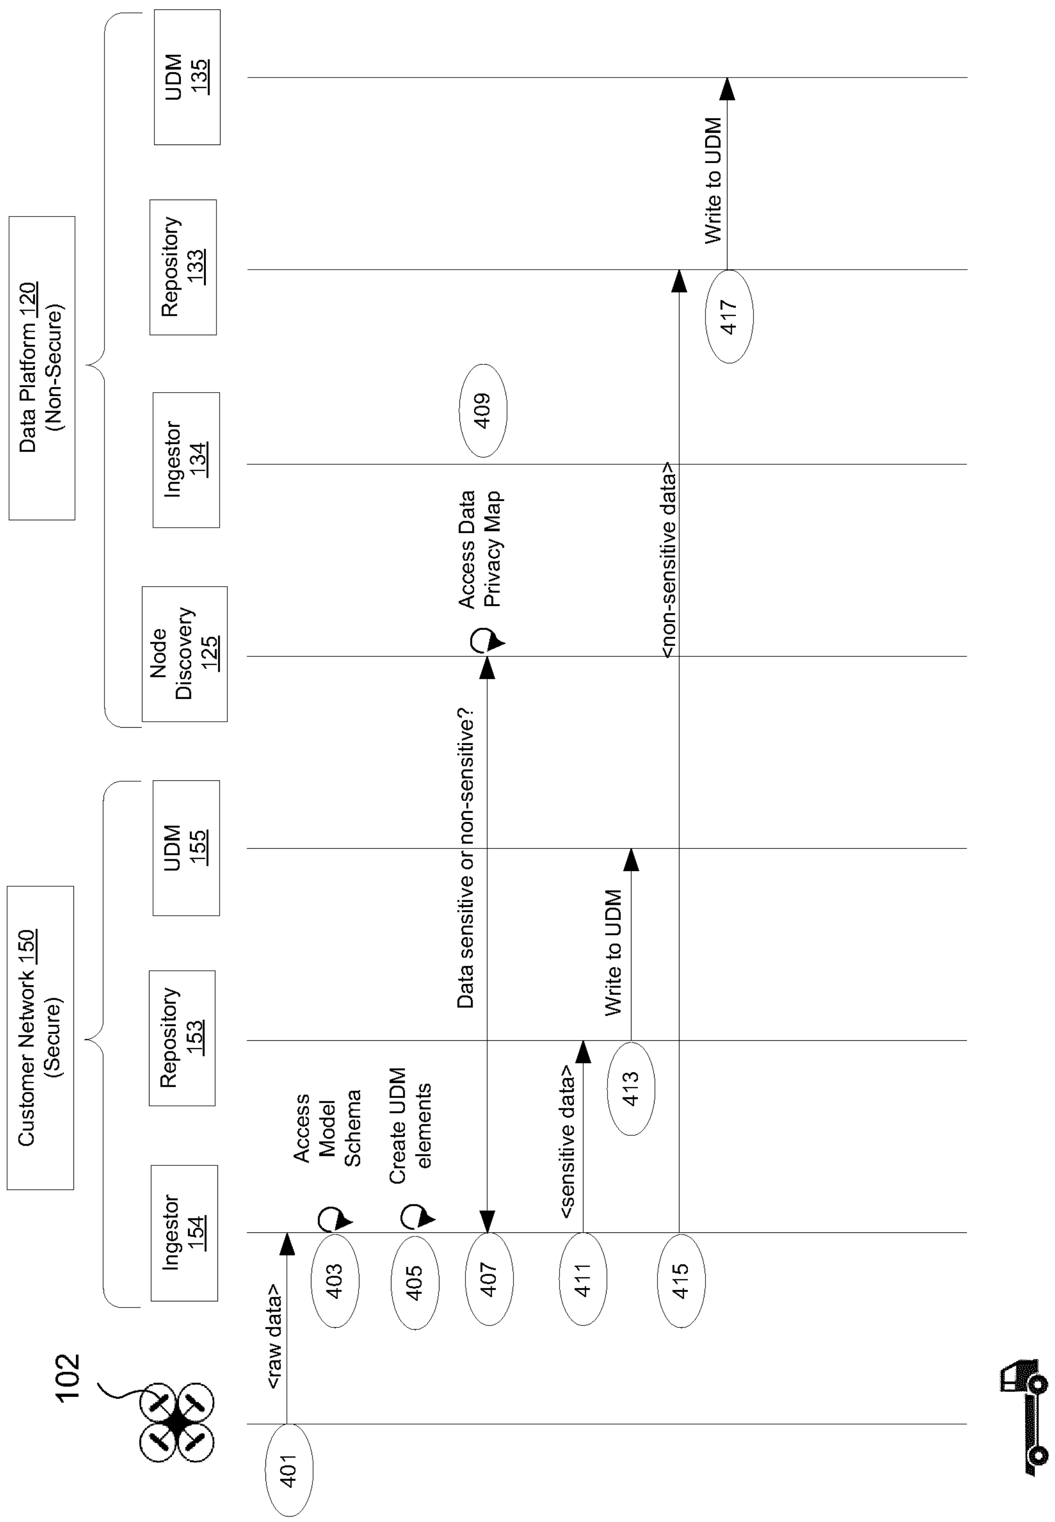
Customer Network 150
(Secure)
Data Platform 120
(Non-Secure)
102
Ingestor 154
Repository 153
UDM 155
Node Discovery 125
Ingestor 134
Repository 133
UDM 135
401
<raw data>
403
Access Model Schema
405
Create UDM elements
407
Data sensitive or non-sensitive?
Access Data Privacy Map
409
411
<sensitive data>
413
Write to UDM
415
<non-sensitive data>
417
Write to UDM

Fig. 4

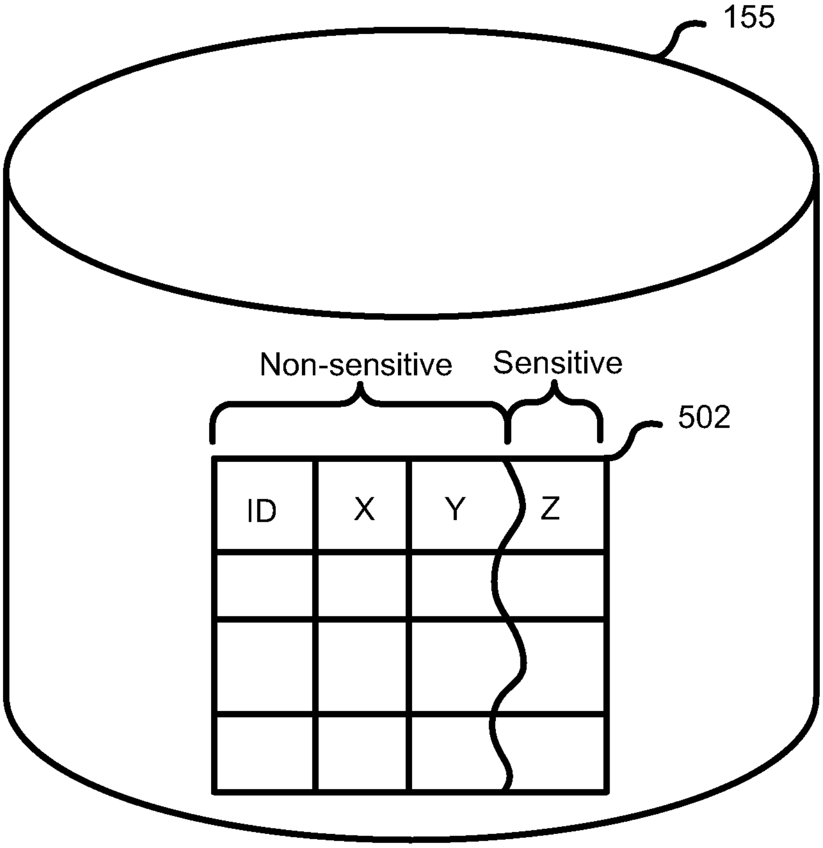

Fig. 5

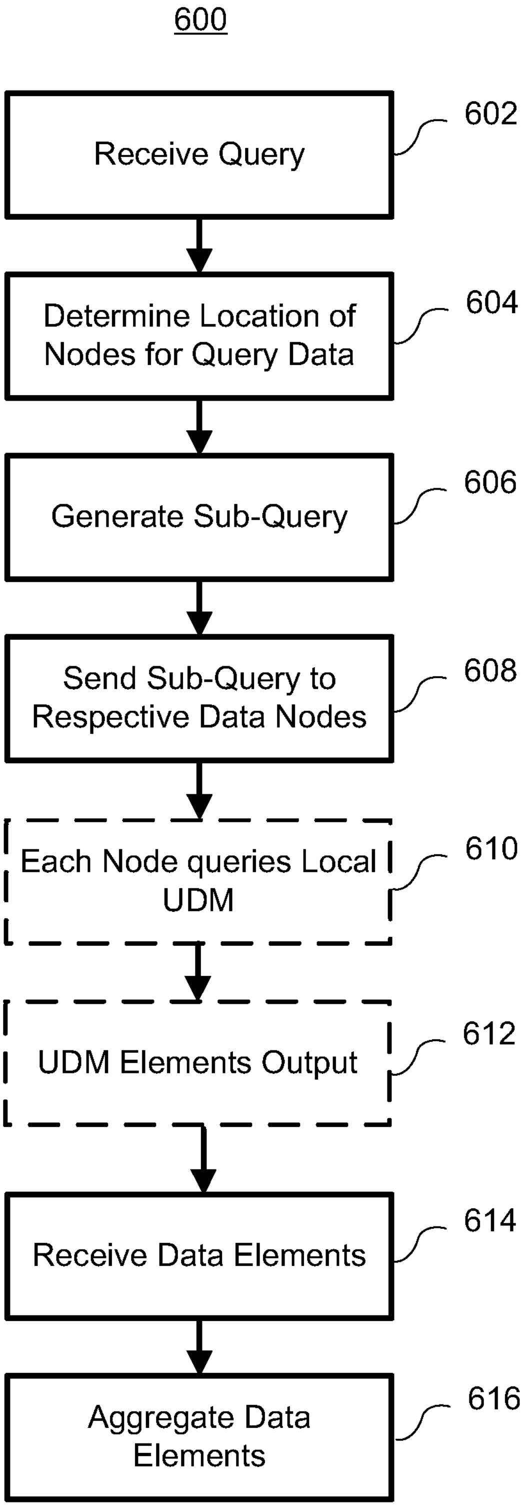
600
Receive Query
602
Determine Location of Nodes for Query Data
604
Generate Sub-Query
606
Send Sub-Query to Respective Data Nodes
608
Each Node queries Local UDM
610
UDM Elements Output
612
Receive Data Elements
614
Aggregate Data Elements
616

Fig. 6

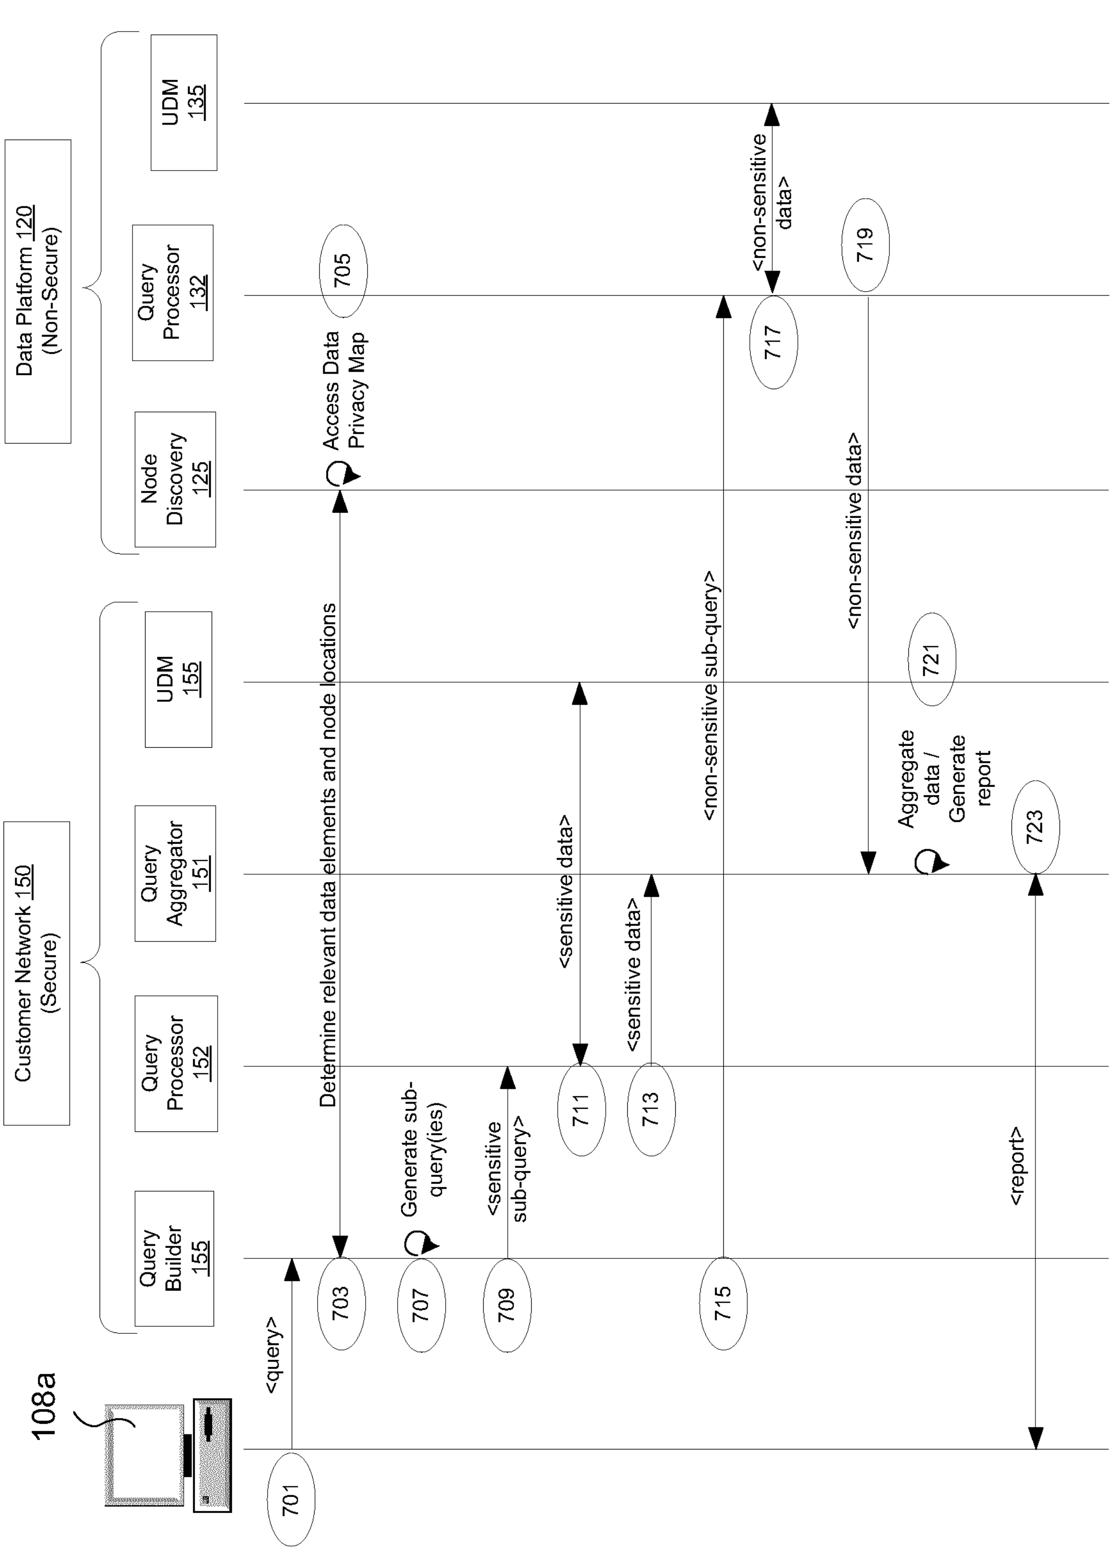
108a
Customer Network 150 (Secure)
Query Builder 155
Query Processor 152
Query Aggregator 151
UDM 155
Data Platform 120 (Non-Secure)
Node Discovery 125
Query Processor 132
UDM 135
701
<query>
703
Determine relevant data elements and node locations
705
Access Data Privacy Map
707
Generate sub-query(ies)
709
<sensitive sub-query>
711
<sensitive data>
713
<sensitive data>
715
<non-sensitive sub-query>
717
<non-sensitive data>
719
<non-sensitive data>
721
Aggregate data / Generate report
723
<report>

Fig. 7

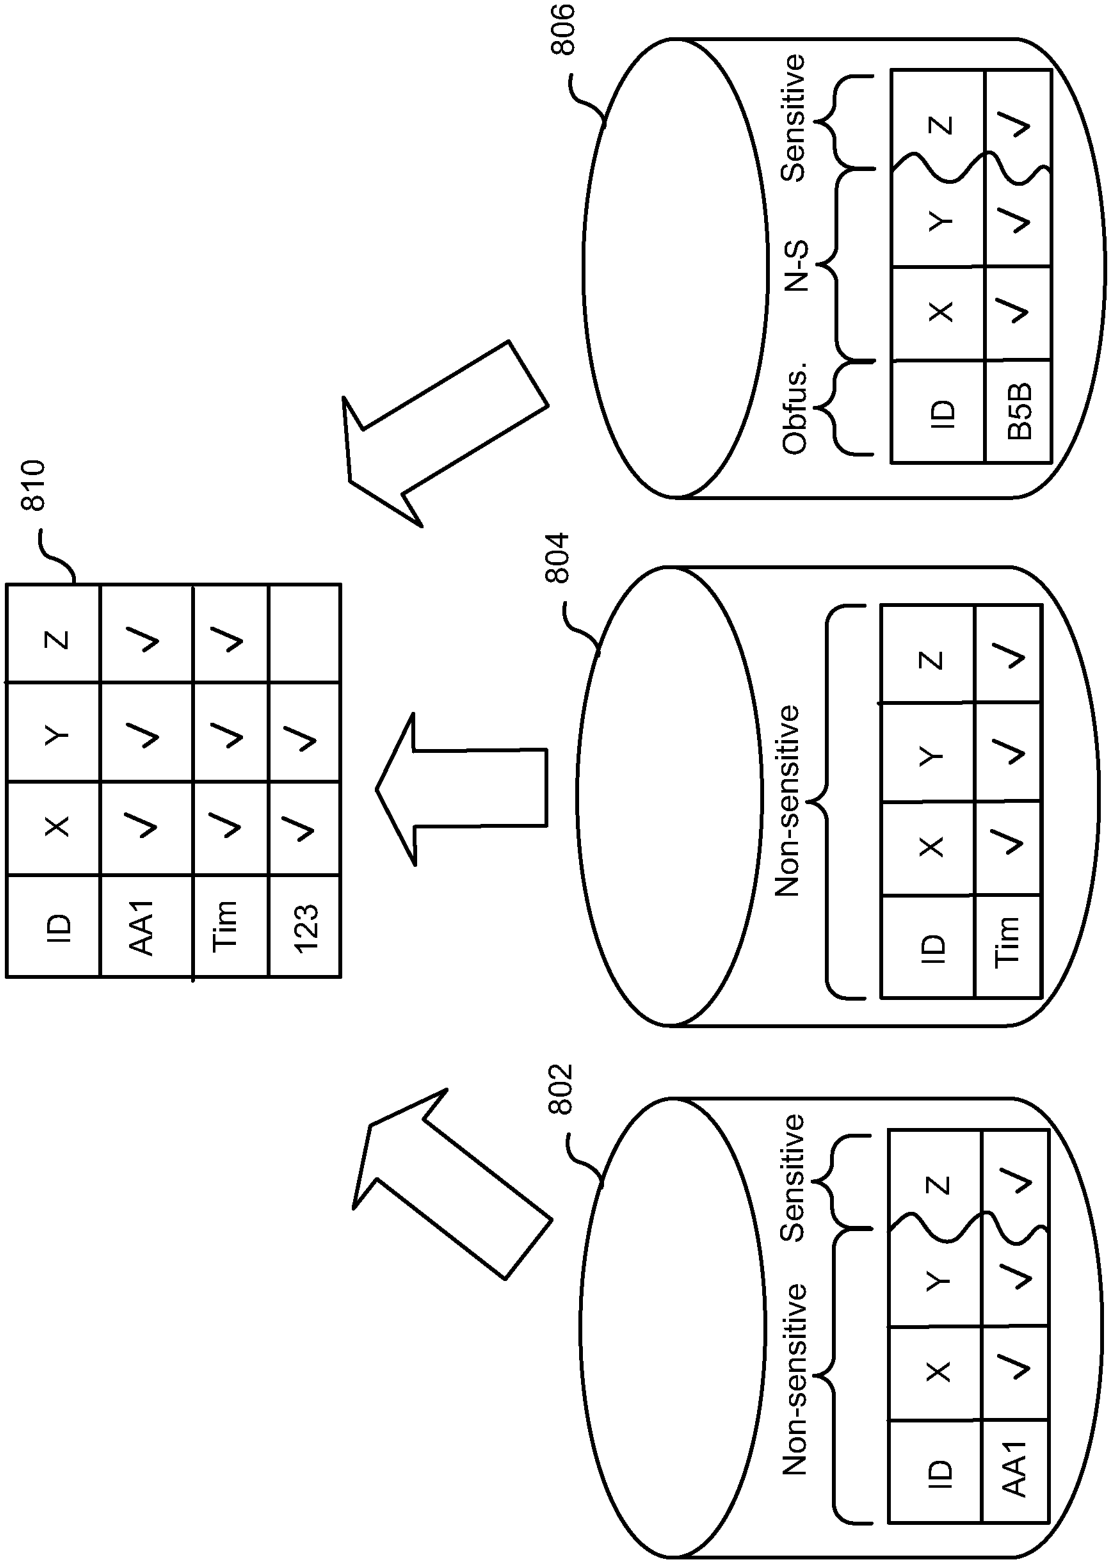

ID	X	Y	Z
AA1	✓	✓	✓

ID	X	Y	Z
Tim	✓	✓	✓

ID	X	Y	Z
B5B	✓	✓	✓

ID	X	Y	Z
AA1	✓	✓	✓
Tim	✓	✓	✓
123	✓	✓	

Fig. 8

INTERNATIONAL SEARCH REPORT

International application No.
PCT/CA2018/051433

A. CLASSIFICATION OF SUBJECT MATTER
IPC: ***G06F 21/60*** (2013.01) , ***G06F 16/901*** (2019.01) , ***G06F 16/903*** (2019.01) , ***G06Q 10/06*** (2012.01)

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 21/60 (2013.01), G06F 16/901 (2019.01), G06F 16/903 (2019.01), G06Q 10/06 (2012.01)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used)
Databases: Questel-Orbit (Fampat), Canadian Patents Database, IEEEXplore, Google
Keywords: partition/segment/divide/separate/split data, convert/format/normalize, unmanned/ autonomous/industrial vehicle, fleet management, common/fixed/standardized format, sensitive/private/confidential, non-sensitive/public, obfuscate, encrypt, sensor, security level, report, query, distributed storage, node, cloud

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2015/0269790 A1 (BATCHELLER, D.C. et al.) 24 September 2015 (24-09-2015) *paragraphs [0023]-[0057], [0075]-[0126]; claims 1, 2, 11, 12, 19; figs. 1, 9, 10A, 10B*	1-34
Y	US 2016/0300078 A1 (WOOLDRIDGE, J.B.) 13 October 2016 (13-10-2016) *paragraphs [0063]-[0093]; figs. 1, 2, 6-8; claims 1-12*	1-12 and 26-34
Y	CLARK, G.J. et al., "Multi-platform Airplane Health Management", 2007 IEEE Aerospace Conference, 3-10 March 2007, Big Sky, MT, USA, 13 pages, 3 March 2007 (03-03-2007) *section 2: Concept of Operation; section 5: Interoperable Database Schema, Access Methods, and Translators*	1-12 and 26-34
Y	US 2010/0299313 A1 (ORSINI, R.L. et al.) 25 November 2010 (25-11-2010) *paragraphs [0011]-[0033], [0093]-[0094], [0416]-[0559]; figs. 27-32, 43-47*	13-25 and 29-32

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

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

Date of the actual completion of the international search
31 January 2019 (31-01-2019)

Date of mailing of the international search report
22 February 2019 (22-02-2019)

Name and mailing address of the ISA/CA
Canadian Intellectual Property Office
Place du Portage I, C114 - 1st Floor, Box PCT
50 Victoria Street
Gatineau, Quebec K1A 0C9
Facsimile No.: 819-953-2476

Authorized officer

Daniela Savin (819) 635-6286

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2010/0174576 A1 (NAYLOR, D.G.) 8 July 2010 (08-07-2010) *whole document*	1-34
A	US 2014/0226010 A1 (MOLIN, H.M. et al.) 14 August 2014 (14-08-2014) *whole document*	1-34
A	US 2016/0232721 A1 (SINGH, P.S. et al.) 11 August 2016 (11-08-2016) *whole document*	1-34
A	US 2005/0187940 A1 (LORA, B. et al.) 25 August 2005 (25-08-2005) *whole document*	1-34
A	US 2017/0104746 A1 (NAIR, A.K. et al.) 13 April 2017 (13-04-2017) * whole document *	1-34
A	US 2013/0145483 A1 (DIMURO, J.D. et al.) 6 June 2013 (06-06-2013) *whole document*	1-34
A	US 2007/0124189 A1 (STOUGHTON, C. et al.) 31 May 2007 (31-05-2007) *whole document*	1-34
A	US 2014/0359552 A1 (MISRA, P. et al.) 4 December 2014 (04-12-2014) *whole document*	1-34
A	US 9 628 488 B1 (DESAI, P. et al.) 18 April 2017 (18-04-2017) *whole document*	1-34

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

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

1. □ Claim Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. □ Claim Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. □ Claim Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

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

This International Searching Authority found multiple inventions in this international application, as follows:
The claims are directed to a plurality of alleged inventions as follows:
Group A: Claims 1-12 and 26-34 are directed to partitioning data received from one or more vehicles; and
Group B: Claims 13-25 are directed to responding to a query requesting vehicle data.
The claims must be limited to one inventive concept as set out in Rule 13 of the PCT.

1. □ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. ☑ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.

3. □ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claim Nos.:

4. □ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claim Nos.:

Remark on Protest

□ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.

□ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.

□ No protest accompanied the payment of additional search fees.

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
US2015269790A1	24 September 2015 (24-09-2015)	US9202318B2	01 December 2015 (01-12-2015)
		US2008074854A1	27 March 2008 (27-03-2008)
		US7616449B2	10 November 2009 (10-11-2009)
		US2010020512A1	28 January 2010 (28-01-2010)
		US7957152B2	07 June 2011 (07-06-2011)
		US2008234936A1	25 September 2008 (25-09-2008)
		US8116922B2	14 February 2012 (14-02-2012)
		US2008077290A1	27 March 2008 (27-03-2008)
		US8565943B2	22 October 2013 (22-10-2013)
		US2014081483A1	20 March 2014 (20-03-2014)
		US9047717B2	02 June 2015 (02-06-2015)
		US2014024395A1	23 January 2014 (23-01-2014)
		US9172481B2	27 October 2015 (27-10-2015)
US2016300078A1	13 October 2016 (13-10-2016)	CA2974836A1	13 October 2016 (13-10-2016)
		CN107209820A	26 September 2017 (26-09-2017)
		WO2016164210A1	13 October 2016 (13-10-2016)
US2010299313A1	25 November 2010 (25-11-2010)	US8654971B2	18 February 2014 (18-02-2014)
		AU2010249631A1	12 January 2012 (12-01-2012)
		AU2010249631B2	18 February 2016 (18-02-2016)
		BRPI1013062A2	05 April 2016 (05-04-2016)
		CA2760251A1	25 November 2010 (25-11-2010)
		CN102428686A	25 April 2012 (25-04-2012)
		CN104079573A	01 October 2014 (01-10-2014)
		EP2433409A2	28 March 2012 (28-03-2012)
		JP2012527838A	08 November 2012 (08-11-2012)
		JP5757536B2	29 July 2015 (29-07-2015)
		JP2015146587A	13 August 2015 (13-08-2015)
		JP6120895B2	26 April 2017 (26-04-2017)
		JP2014238599A	18 December 2014 (18-12-2014)
		US2014150120A1	29 May 2014 (29-05-2014)
		US9064127B2	23 June 2015 (23-06-2015)
		US2015249687A1	03 September 2015 (03-09-2015)
		WO2010135412A2	25 November 2010 (25-11-2010)
		WO2010135412A3	03 February 2011 (03-02-2011)
US2010174576A1	08 July 2010 (08-07-2010)	US8548669B2	01 October 2013 (01-10-2013)
		CA2689744A1	08 July 2010 (08-07-2010)
		CA2689744C	05 May 2015 (05-05-2015)
		US2014095047A1	03 April 2014 (03-04-2014)
		US9053585B2	09 June 2015 (09-06-2015)
US2014226010A1	14 August 2014 (14-08-2014)	US9922567B2	20 March 2018 (20-03-2018)
		DE112012003061T5	15 May 2014 (15-05-2014)
		US2018225974A1	09 August 2018 (09-08-2018)
		WO2013012926A1	24 January 2013 (24-01-2013)
US2016232721A1	11 August 2016 (11-08-2016)	US10032317B2	24 July 2018 (24-07-2018)
US2005187940A1	25 August 2005 (25-08-2005)	US7917536B2	29 March 2011 (29-03-2011)
		WO2005081848A2	09 September 2005 (09-09-2005)
		WO2005081848A3	01 February 2007 (01-02-2007)
US2017104746A1	13 April 2017 (13-04-2017)	US9948637B2	17 April 2018 (17-04-2018)
		US2018191705A1	05 July 2018 (05-07-2018)
		US10057253B2	21 August 2018 (21-08-2018)

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/CA2018/051433

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
US2013145483A1	06 June 2013 (06-06-2013)	None	
US2007124189A1	31 May 2007 (31-05-2007)	AU2006320836A1	07 June 2007 (07-06-2007)
		CA2631159A1	07 June 2007 (07-06-2007)
		EP1963938A2	03 September 2008 (03-09-2008)
		EP1963938A4	27 April 2011 (27-04-2011)
		IL191700D0	29 December 2008 (29-12-2008)
		JP2009517777A	30 April 2009 (30-04-2009)
		KR20080073312A	08 August 2008 (08-08-2008)
		WO2007064509A2	07 June 2007 (07-06-2007)
		WO2007064509A3	15 November 2007 (15-11-2007)
US2014359552A1	04 December 2014 (04-12-2014)	AU2012338372A1	27 March 2014 (27-03-2014)
		AU2012338373A1	27 March 2014 (27-03-2014)
		BR112014006445A2	04 April 2017 (04-04-2017)
		BR112014006446A2	04 April 2017 (04-04-2017)
		CA2848988A1	23 May 2013 (23-05-2013)
		CA2848988C	22 May 2018 (22-05-2018)
		CA2848995A1	23 May 2013 (23-05-2013)
		CN103890730A	25 June 2014 (25-06-2014)
		CN103890730B	11 July 2017 (11-07-2017)
		CN103891201A	25 June 2014 (25-06-2014)
		CN103891201B	30 March 2018 (30-03-2018)
		EP2758879A2	30 July 2014 (30-07-2014)
		EP2758879A4	22 July 2015 (22-07-2015)
		EP2759093A2	30 July 2014 (30-07-2014)
		EP2759093A4	15 July 2015 (15-07-2015)
		JP2014534487A	18 December 2014 (18-12-2014)
		JP5840786B2	06 January 2016 (06-01-2016)
		JP2015501459A	15 January 2015 (15-01-2015)
		MX2014003168A	22 August 2014 (22-08-2014)
		MX337513B	09 March 2016 (09-03-2016)
		MX2014003171A	09 July 2014 (09-07-2014)
		MX350877B	19 September 2017 (19-09-2017)
		TW201328921A	16 July 2013 (16-07-2013)
		TWI627084B	21 June 2018 (21-06-2018)
		TW201333731A	16 August 2013 (16-08-2013)
		TWI630493B	21 July 2018 (21-07-2018)
		US2014380264A1	25 December 2014 (25-12-2014)
		US9990182B2	05 June 2018 (05-06-2018)
		WO2013072925A2	23 May 2013 (23-05-2013)
		WO2013072925A3	11 July 2013 (11-07-2013)
		WO2013072926A2	23 May 2013 (23-05-2013)
		WO2013072926A3	27 February 2014 (27-02-2014)
		ZA201401803B	25 February 2015 (25-02-2015)
		ZA201401804B	25 February 2015 (25-02-2015)
US9628488B1	18 April 2017 (18-04-2017)	US9819655B1	14 November 2017 (14-11-2017)