



US 20170053295A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2017/0053295 A1**

Tiell et al.

(43) **Pub. Date: Feb. 23, 2017**

(54) **PLATFORM DATA MARKETPLACE**

(52) **U.S. Cl.**

(71) Applicant: **Accenture Global Solutions Limited**,
Dublin (IE)

CPC **G06Q 30/0201** (2013.01); **G06Q 30/0601**
(2013.01)

(72) Inventors: **Steven C. Tiell**, San Francisco, CA
(US); **Shimon Modi**, Washington, DC
(US)

(57) **ABSTRACT**

(21) Appl. No.: **15/095,440**

Techniques are described for receiving data generated by multiple platforms of different types, and providing the data, aggregate data, and/or data analysis results to data consumer (s) through a marketplace associated with a super-platform. The data from various individual platforms may be received, ingested, stored, analyzed, aggregated, and/or otherwise processed by a super-platform. A data consumer may interact with a marketplace that is a component of, or in communication with, the super-platform. The data consumer may request one or more data set(s), and the marketplace may assemble and provide the data set(s) in response to the request. In some examples, the data consumer may request an insight or answer provided by an analysis job that analyzes the data collected from the individual platforms. The analysis job may be executed on the super-platform a device of the data consumer, and analysis results may be provided to the data consumer.

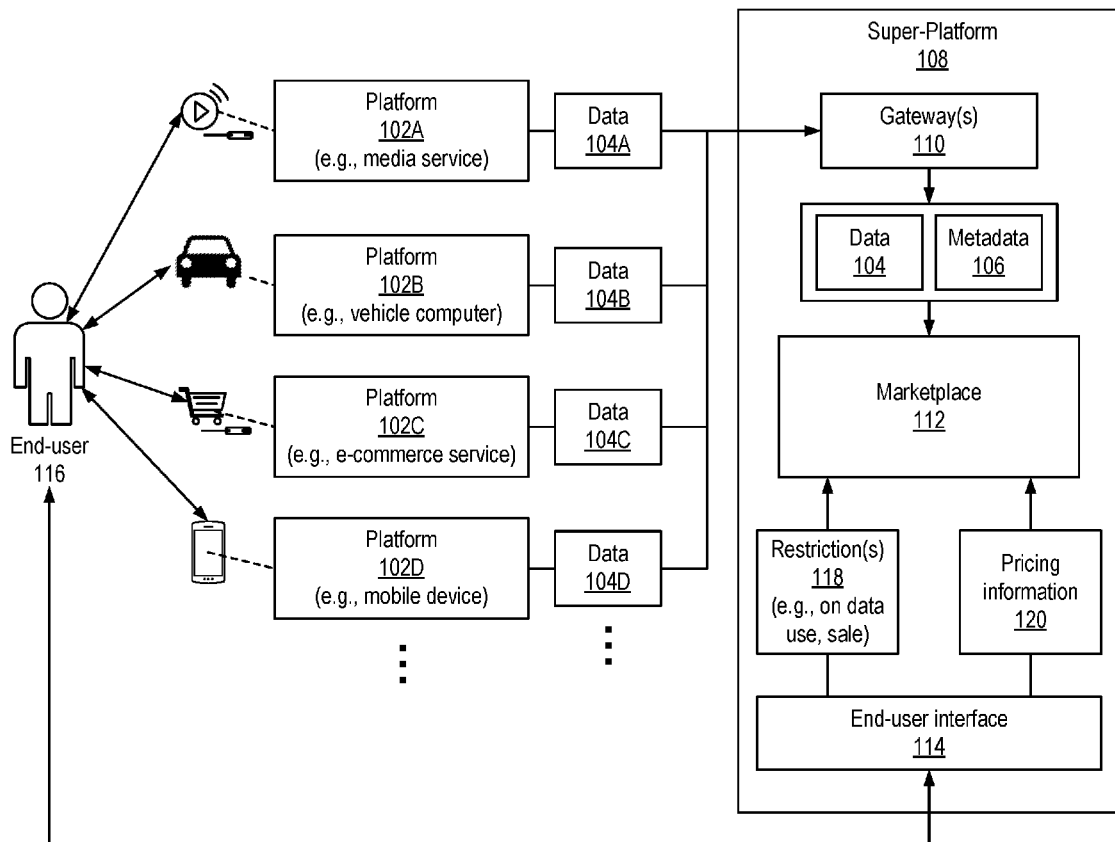
(22) Filed: **Apr. 11, 2016**

Related U.S. Application Data

(60) Provisional application No. 62/206,078, filed on Aug. 17, 2015.

Publication Classification

(51) **Int. Cl.**
G06Q 30/02 (2006.01)
G06Q 30/06 (2006.01)



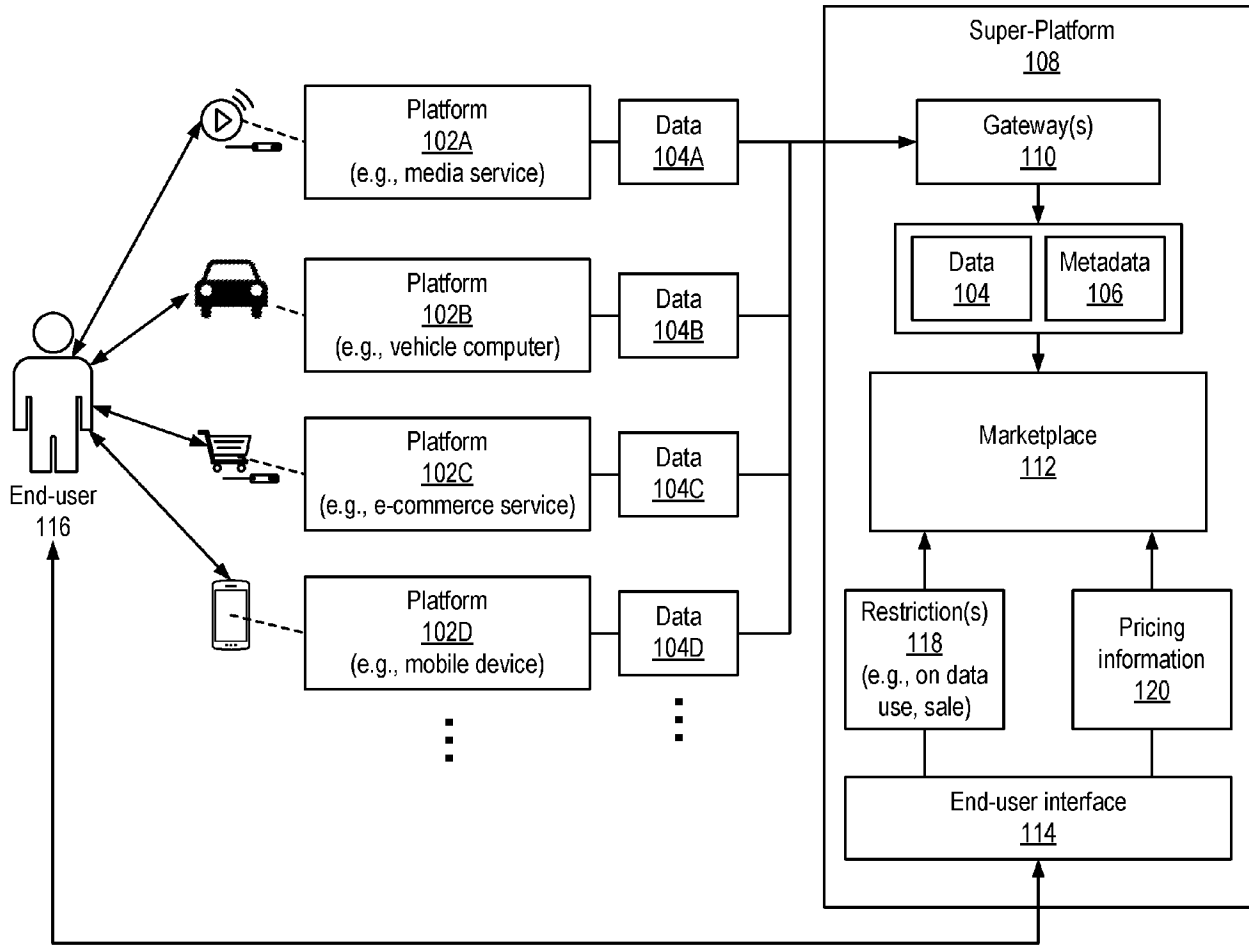


FIG. 1

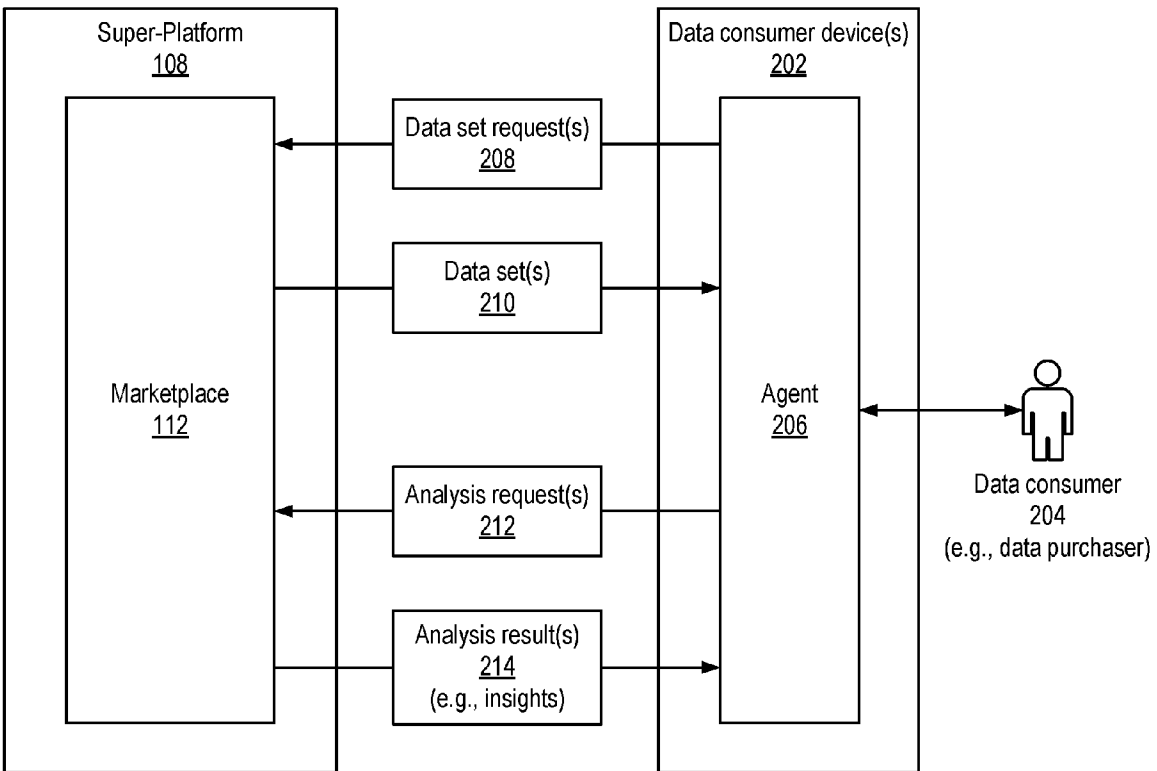


FIG. 2

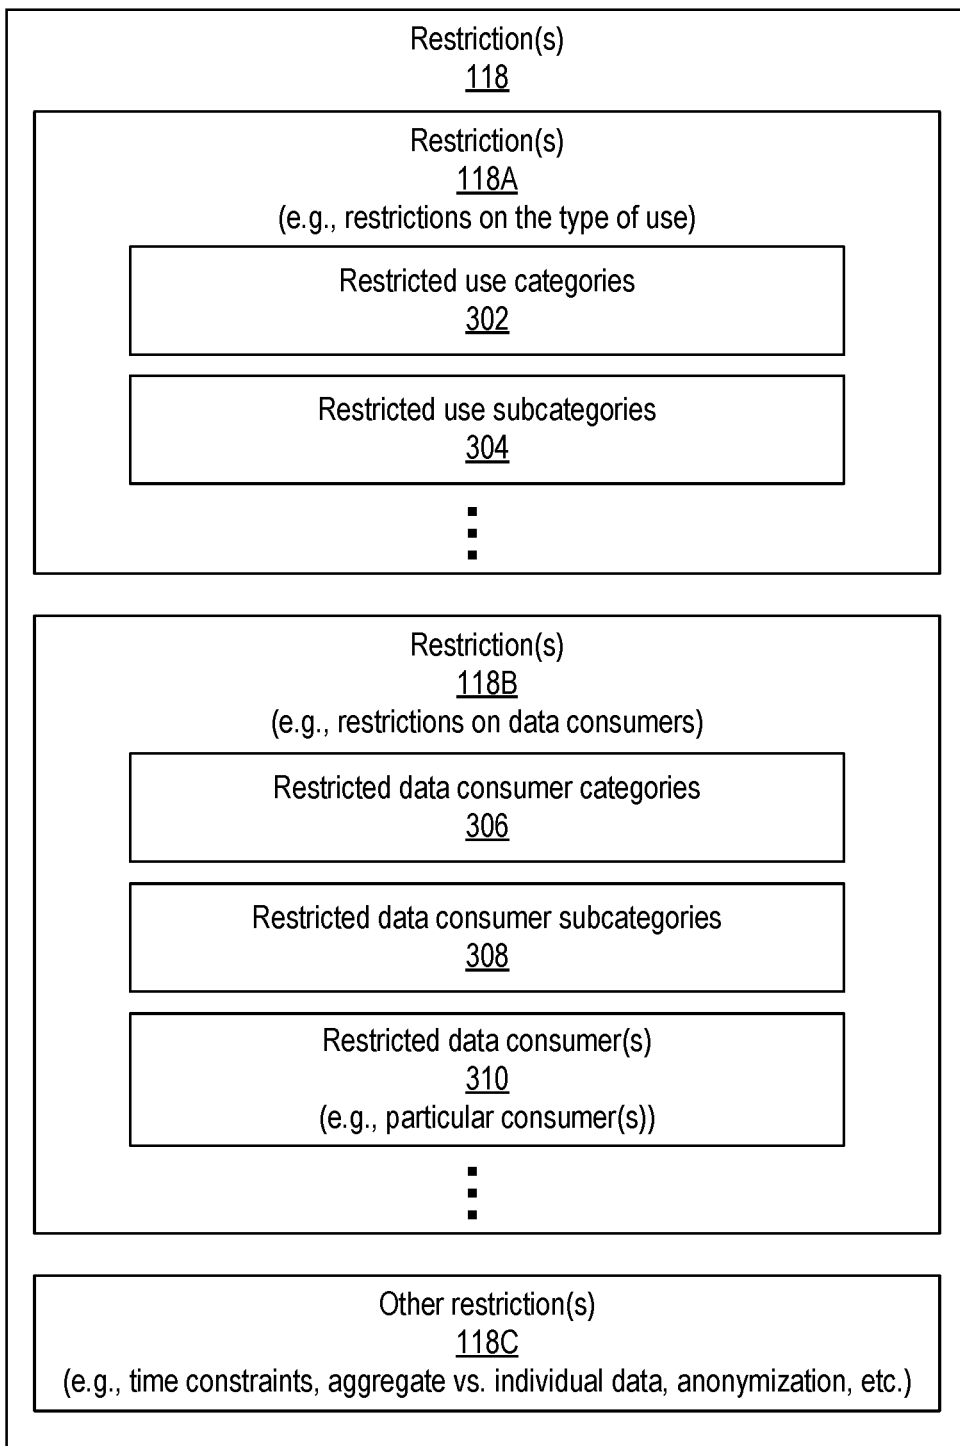


FIG. 3

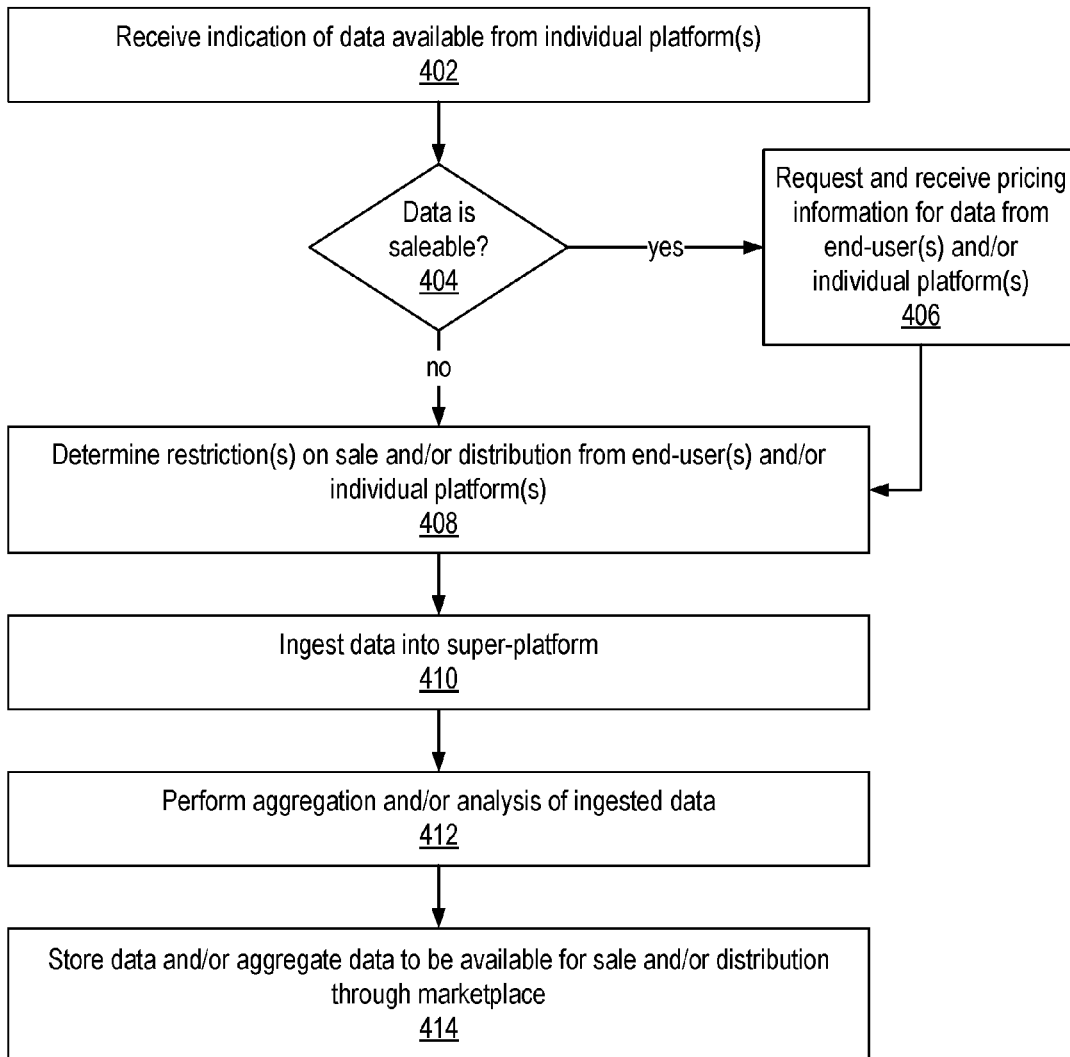


FIG. 4

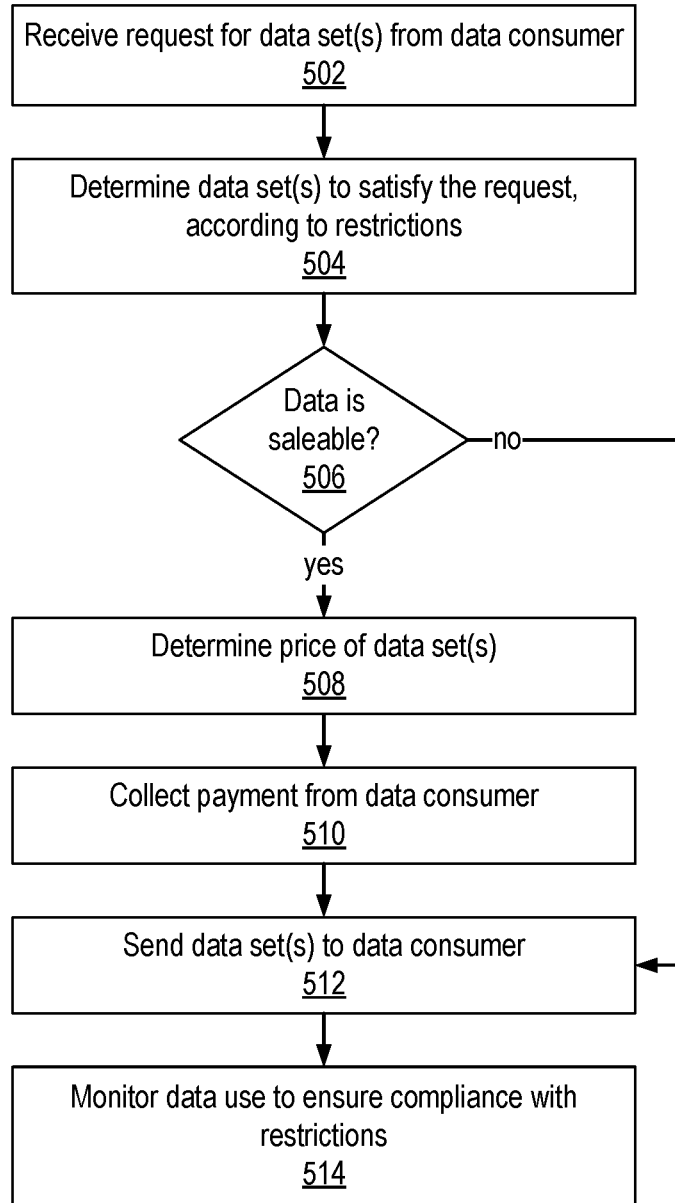


FIG. 5

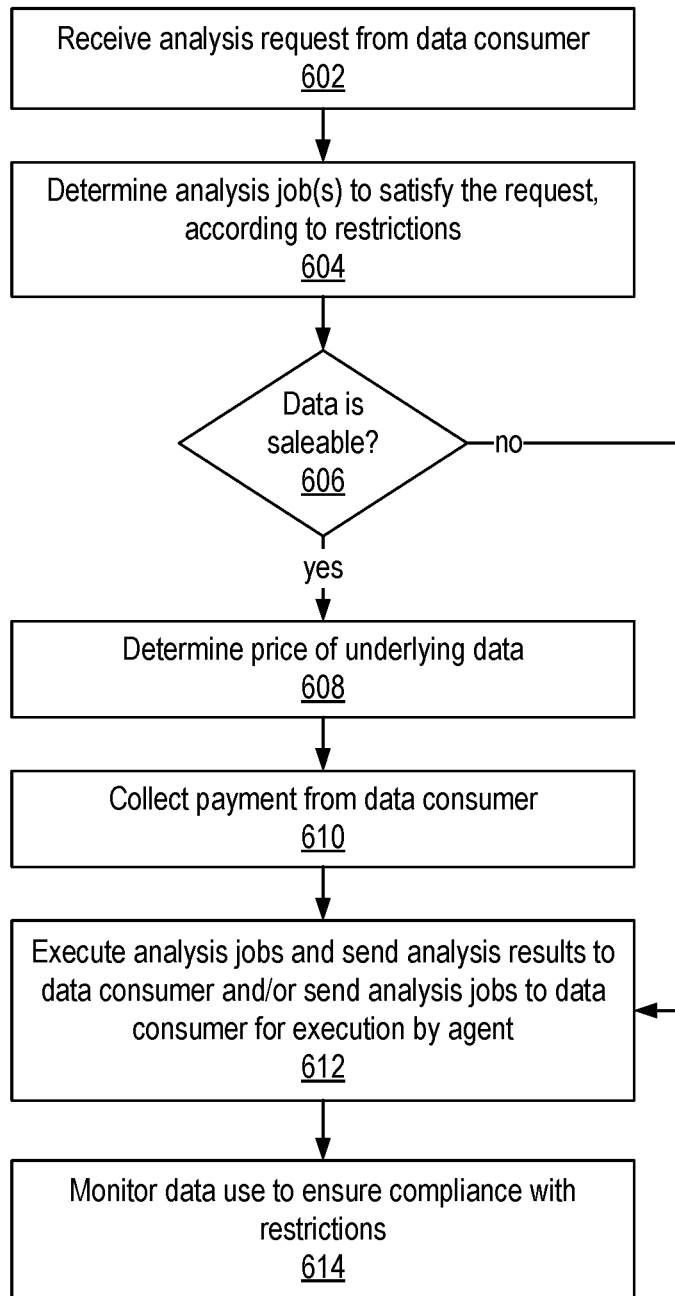


FIG. 6




FIG. 7

PLATFORM DATA MARKETPLACE

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims priority to U.S. Provisional Patent Application Ser. No. 62/206,078 filed on Aug. 17, 2015, titled "Super-Platform," which is hereby incorporated by reference in the entirety.

BACKGROUND

[0002] With the rapid growth of ubiquitous computing and the Internet of Things (IoT), users are interacting with many different computing systems in the course of their lives. In addition to interacting with devices whose primary function is computing, such as home computers, mobile computing devices, server computers, and so forth, users may also interact with devices that include computers as components, such as smart appliances, automotive computers, or any number of IoT devices. The various computing devices may generate data regarding users, and such data may be generated and delivered using different formats, models, or communications channels.

SUMMARY

[0003] Implementations of the present disclosure are generally directed to the providing of data, aggregate data, analysis results, analysis jobs, and/or other information to data consumers. More specifically, implementations are directed to receiving and aggregating, at a super-platform, data generated by multiple individual platforms of various types, and selling or otherwise providing the aggregate data, analysis results, analysis jobs, and/or other information to data consumers through a marketplace associated with the super-platform.

[0004] In general, innovative aspects of the subject matter described in this specification can be embodied in methods that includes actions of: receiving data generated at a plurality of individual platforms, the data associated with one or more end-users of the plurality of individual platforms; receiving a request for at least a portion of the data, the request provided by a data consumer through a marketplace; providing at least the portion of the data to the data consumer at a price that is based at least partly on pricing information provided by the one or more end-users associated with the data.

[0005] Implementations can optionally include one or more of the following features: the actions further include aggregating the data from multiple individual platforms to generate aggregate data; the actions further include providing the aggregate data to the data consumer at the price; the actions further include receiving an analysis request provided by the data consumer through the marketplace; the actions further include executing at least one analysis job that is determined to satisfy the analysis request; the actions further include providing at least one analysis result of the at least one analysis job to the data consumer; the actions further include receiving an analysis request provided by the data consumer through the marketplace; the actions further include determining at least one analysis job to satisfy the analysis request; the actions further include providing the at least one analysis job for execution on a data consumer device of the data consumer; the plurality of individual platforms include at least two of: an e-commerce platform

configured to provide an e-commerce service, a retail platform configured to support retail operations of a physical store, a media platform configured to provide an online media service, a mobile computing device, an Internet of Things (IoT) computing device, or an in-vehicle computer; the actions further include receiving one or more restrictions on the use of the data, the one or more restrictions provided by the one or more end-users; the actions further include receiving, from the data consumer, an indicated use of the data; the actions further include determining at least the portion of the data to be provided to the data consumer, such that the indicated use complies with the one or more restrictions on the use of at least the portion of the data; the actions further including monitoring the use of at least the portion of the data provided to the data consumer to ensure compliance with the one or more restrictions; the actions further include receiving one or more restrictions on data consumers to be provided the data, the one or more restrictions provided by the one or more end-users; the actions further include determining at least one characteristic of the data consumer; and/or the actions further include determining at least the portion of the data to be provided to the data consumer, such that the at least one characteristic complies with the one or more restrictions on data consumers to be provided the data.

[0006] Other implementations of any of the above aspects include corresponding systems, apparatus, and computer programs that are configured to perform the actions of the methods, encoded on computer storage devices. The present disclosure also provides a computer-readable storage medium coupled to one or more processors and having instructions stored thereon which, when executed by the one or more processors, cause the one or more processors to perform operations in accordance with implementations of the methods provided herein. The present disclosure further provides a system for implementing the methods provided herein. The system includes one or more processors, and a computer-readable storage medium coupled to the one or more processors having instructions stored thereon which, when executed by the one or more processors, cause the one or more processors to perform operations in accordance with implementations of the methods provided herein.

[0007] Implementations provide the following advantages. By ingesting, aggregating, and distributing data from multiple platforms of various types, implementations provide data consumers with a more comprehensive view of the characteristics and/or activities of end-users compared to traditional systems that may provide a more limited view based on a particular type of data or based on data collected at a single platform. Moreover, because implementations provide data analysis jobs that analyze data from multiple platforms of various types, implementations provide data consumers with more accurate insights into end-user characteristics and behaviors compared to traditional systems which may only analyze data collected by a single platform. By providing more accurate and/or more comprehensive analysis of data based on the collected data of multiple individual platforms, implementations provide technical advantages over traditional data analysis platforms. For example, a traditional data analysis platform may require a larger amount of data from a single source, and thereby consume more processing, storage, memory, and/or communications resources than the super-platform provided by the implementations described herein. In some implemen-

tations, the super-platform may provide for centralized coordination (e.g., orchestration) of distributed analytics jobs, such that data that is required to remain in specific geographic regions can still be included in super-platform coordinated analysis jobs. Implementations also provide for relatively clean data (e.g., relative to traditional systems) and annotated data sets. In some implementations, the super-platform provides lifelong and/or immutable data provenance that is achieved through the automated integration and addition of metadata to new and existing data assets. Implementations may also provide for extensible in-flight analytics, which can also function as a platform, allowing for third party developers to publish models and algorithms. Implementations provide aggregation features that allow for on-the-fly assembly of data assets based upon requirements. At least some implementations also support the ability to autonomously scrub data prior to sharing to minimize re-identification risks. Implementations provide the super-platform based and/or cloud based analytics. Moreover, scalability may be achieved through a platform-of-platforms architectural approach as described herein.

[0008] It is appreciated that aspects and features in accordance with the present disclosure can include any combination of the aspects and features described herein. That is, aspects and features in accordance with the present disclosure are not limited to the combinations of aspects and features specifically described herein, but also include any combination of the aspects and features provided.

[0009] The details of one or more implementations of the present disclosure are set forth in the accompanying drawings and the description below. Other features and advantages of the present disclosure will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

[0010] FIG. 1 depicts an example system for collecting data from multiple data generating platforms, in accordance with implementations of the present disclosure.

[0011] FIG. 2 depicts an example system for providing the collected data through a marketplace to one or more data consumers, in accordance with implementations of the present disclosure.

[0012] FIG. 3 depicts an example of restriction(s) on data use and/or distribution, in accordance with implementations of the present disclosure.

[0013] FIG. 4 depicts an example process for ingesting data from multiple data generating platforms, in accordance with implementations of the present disclosure.

[0014] FIG. 5 depicts an example process for providing data set(s) to data consumer(s), in accordance with implementations of the present disclosure.

[0015] FIG. 6 depicts an example process for providing result(s) of analysis job(s) to data consumer(s), in accordance with implementations of the present disclosure.

[0016] FIG. 7 depicts an example computing system architecture, in accordance with implementations of the present disclosure.

[0017] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0018] Implementations of the present disclosure include systems, devices, methods, and computer-readable media

for receiving data generated by multiple platforms of different types, and providing the data, aggregate data, and/or data analysis results to data consumer(s) through a marketplace associated with a super-platform. An end-user may interact with multiple individual (e.g., siloed) platforms of different types or that support different business purposes or industries. The individual platforms may generate data describing, and/or resulting from, these interactions with end-user(s). The data from the various individual platforms may be received, ingested, stored, analyzed, aggregated, and/or otherwise processed by a super-platform. A data consumer may interact with a marketplace that is a component of, or in communication with, the super-platform. In some examples, the data consumer may request one or more data set(s), and the marketplace may assemble and provide the data set(s) in response to the request. In some examples, the data consumer may request an insight or an answer provided by an analysis job that analyzes the data collected from the individual platforms. The analysis job may be executed on the super-platform, or on a computing device of the data consumer, and the analysis results may be provided to the data consumer. In some implementations, an end-user and/or the individual platform(s) may indicate a price for data that is associated with the end-user and/or generated by the individual platform(s), and the marketplace may sell the data, aggregate data, and/or analysis results to data consumers for a price that is based on the price indicated by the end-user and/or individual platforms. In some implementations, the end-user and/or the individual platform(s) may provide one or more restrictions on how the collected data may be used, sold, distributed, or otherwise handled. In such examples, the sale and/or distribution of data through the marketplace, and/or the analysis of the data, may be performed in accordance with the restriction(s).

[0019] As used herein, super-platform refers to a computing system that includes any number of computing devices executing any number of software modules to receive, analyze, aggregate, store, distribute, and/or otherwise process data generated by any number of individual platforms. In some instances, the super-platform may be separate from the individual platform(s) and configured to communicate with the platform(s) using one or more networks. In some instances, the super-platform may share computing hardware and/or software resources with at least one of the individual platforms that is generating data. The super-platform may be configured to process data generated by multiple individual platforms that provide different types of services, as described further below.

[0020] In some implementations, various types of entities (e.g., individuals, computing devices, processes, etc.) may interact with the super-platform. Such entities may include data disclosers, data manipulators, and data consumers. A data discloser may be any entity that generates and/or provides the data to the super-platform. A data discloser may also be any entity that causes the data to be generated and/or provided to the super-platform. For example, an end-user of individual platform(s) may be a data discloser. An individual platform may also, itself, be a data discloser. A data manipulator may be any entity that analyzes, aggregates, and/or otherwise processes the data. For example, the super-platform and/or the various processes executing on the super-platform may be described as data manipulators. In some instances, the individual platforms and/or end-users may be data manipulators as well as possibly data disclosers, e.g., in

scenarios where the data discloser manipulates the data in some way prior to disclosing the data to the super-platform. A data consumer may be any entity that receives the data and/or any information that is generated through manipulation of the data (e.g., aggregate data, analysis results, etc.). A data consumer may purchase or otherwise acquire information from the super-platform. In some instances, a data consumer may also be a data manipulator and/or data discloser. For example, a data consumer may purchase a set of data from the super-platform, analyze (e.g., manipulate) the data to produce analysis result(s) and/or other data set(s), and sell (e.g., disclose) the analysis result(s) and/or other data set(s) to the super-platform, individual platform(s), and/or other entities. Accordingly, a particular entity may operate in accordance with one or more of the roles of data discloser, data manipulator, and/or data consumer depending on the particular circumstances.

[0021] FIG. 1 depicts an example system for collecting data from multiple data generating platforms, in accordance with implementations of the present disclosure. An end-user 116 may interact with multiple, different individual platforms 102. Each of the individual (e.g., siloed) platforms 102 may generate data 104 describing the end-user's interactions with the platform 102. In some implementations, the individual platforms 102 may provide different types of services, and/or may be different types of platforms that support businesses or other organizations in different types of industries. The platforms 102 may include, but are not limited to, one or more of the following:

[0022] A computing device, which may be a mobile or portable computing device such as a smartphone, tablet computer, notebook computer, personal data assistant (PDA), wearable computer (e.g., watch device or glasses device), and so forth;

[0023] An in-vehicle (e.g., automotive) computer that is a component of a commercial or personal-use vehicle or otherwise in communication with a vehicle, the in-vehicle computer configured to collect telematics data, data regarding the operations (e.g., speed, braking, gas mileage, etc.) of the vehicle, characteristics of the vehicle (e.g., temperature, fuel status, fluid status, etc.), location information, environment information (e.g., temperature, air pressure, etc.), and so forth;

[0024] Online services for streaming or otherwise delivering content over a network (e.g., the internet), such as video content, audio content (e.g., music), audio-visual content (e.g., films, television, etc.), games, graphics, images, or other media;

[0025] Online shopping or other e-commerce services that offer products or services through a web site, a native application, or other type of interface;

[0026] Computing systems, such as point-of-sale terminals, servers, data storage devices, and so forth that are associated with physical (e.g., brick-and-mortar) stores or other commercial enterprises, the computing systems collecting or otherwise generating data based on in-store customer activities;

[0027] Internet-of-things (IoT) computing devices, such as smart appliances, home entertainment devices, or other devices present in a home, business setting, or elsewhere, and configured to include processing, memory, and communications capabilities in a device used for other purposes,

such as smart washing machines, heating and air conditioning systems, audio or video presentation systems, and so forth; or

[0028] An analytics cloud or other distributed computing system, such as a social media analytics cloud (SMAC).

[0029] A platform 102 may include a technical (e.g., hardware and/or software) architecture, governance rules, and one or more technology services that support the creation and execution of applications. Such applications may be specific to a particular industry associated with the platform 102. For example, platforms 102 that are vehicle computing systems, brick-and-mortar retail support systems, health care provider systems, or health insurance portals may provide and support applications associated with vehicle operations, retail operations, health care operations, or health insurance operations respectively. A platform 102 may provide a blueprint for how a business or other organization may build, connect, and deliver applications specific to problems in their associated industry. A platform 102 may serve as a pool of reusable functionality and capabilities to make building and evolving applications fast and easy, and may help companies or other organizations achieve better business outcomes. A platform 102 may also serve as a business model strategy that creates competitive differentiation. In a platform-based business, others outside a company may create value for the enterprise and, in some examples, enable new digital models for the company.

[0030] A platform 102 may include any number of applications (e.g., apps), and service layers or plug-in modules that support one or more features. Example features include financial transactions, authentication, logon, navigation, search, advertisement, marketplaces, or any other appropriate feature. An individual platform 102 may include hardware, software, or both hardware and software components. For example, a platform 102 that is an in-vehicle (e.g., telematics) or other type of computing system may include hardware components such as processor(s), memory, network interfaces, and so forth, as well as software components such as an operating system, firmware, or any number of applications or processes executing on the hardware component(s). As another example, a platform 102 that is an e-commerce or media streaming service may include any number of web servers, application servers, applications, or other software component(s) that execute on one or more computing devices to enable an end-user 116 to view content or interact with the platform 102 through a user interface (UI).

[0031] The individual platforms 102 may each generate data 104 that is associated with end-user(s) 116, such as data describing user identity, user characteristics, user activities, user purchases, content browsed by the user, and so forth. In some examples, the data 104 may be provided to the platform(s) 102 by the end-user 116, or may be otherwise generated based on the end-user's interactions with the platform 102. The data 104 may include personally identifiable information (PII) of the end-user 116, such as user name(s), address(es), telephone number(s), login information (e.g., credentials) to access the individual platform 102, or other information that may enable identification of the particular end-user 116 or that is particularly associated with the end-user 116. The data 104 may also include other types of information that are associated with the particular end-user 116, but not necessarily explicitly provided by the end-user 116, such as information describing inferred or

measured user characteristics (e.g., demographic characteristics), user activities on the platform 102, user purchase history, browsing data, clickstream data, and so forth. In some examples, the collected data 104 may include data that is generated by an individual platform 102 but that is not specific to a particular end-user 116. For example, an individual platform 102 that is a vehicle or in-vehicle telematics device may provide data 104 that is collected by sensors of the vehicle, such as data describing temperature, air pressure, location, velocity, vehicle operational state, or other information. As another example, a platform 102 that is a computing device may generate and provide data 104 describing the location of the device, used and available memory capacity, used and available processing capacity, used or available storage capacity, installed or executing applications, hardware configuration, software configuration, and/or other information that is not specifically associated with an end-user.

[0032] The data 104 generated by the platform(s) 102 may include any number of data objects. As used herein, a data object may include any size or amount of data that is arranged according to any format, or unformatted. A data object may include data of any type, and may include data that is compressed or uncompressed, encrypted or unencrypted. In some examples, a data object may include a particular amount of data and may be arranged according to a format such that the data object has a particular size when communicated. Data object(s) may also include a stream of data that continues for any duration to communicate any amount of data until the end of the stream. Data object(s) may be communicated in a sequence corresponding to an order of the data included in the data object(s). The data object(s) may also be communicated out of sequence, and may include sequence indicators or order indicators that enable a receiving process to reorder the received data object(s) into the order in which they were generated.

[0033] The data 104 generated by the individual platforms 102 may be sent to a (e.g., centralized) super-platform 108. For each individual platform 102, the super-platform 108 may include a gateway 110 that operates as an interface to receive the data 104 from the individual platform 102. The super-platform 108 may store the data 104, and may perform operations on the data 104 to process, store, analyze, aggregate, or otherwise handle the data 104. The super-platform 108 may include any number of software modules executing on any number of computing devices. Such software module(s) may include a marketplace 112, the gateway(s) 110, and/or an end-user interface 114. In some examples, the super-platform 108 may be implemented as a cluster or other group of server computers configured to execute the software module(s) that provide and support a super-platform service to ingest data 104, analyze, aggregate, or otherwise process the data 104, and provide the data 104 through the marketplace 112 to data consumer(s).

[0034] In some implementations, the super-platform 108 may execute one or more modules that aggregate the data 104 to generate aggregate data. Data aggregation may aggregate data 104 associated with a single user of one or more platforms 102. Data aggregation may also aggregate data 104 from multiple users. For example, data aggregation may generate data sets for particular segments of users, such as users having particular demographic characteristics, location, demonstrated interest in products or services, and so forth. Aggregate data may include, or be generated based on,

data that describes or identifies users, such as PII for one or more users. Aggregate data may also include, or be generated based on, data that is not associated with users, such as sensor data. The super-platform 108 may aggregate data 104 received from different platforms 102 that support different types of operations or that are associated with different industries. For example, the super-platform 108 may aggregate data 104 from an end-user's smartphone, from a brick-and-mortar retail business where the end-user shops, from the end-user's in-vehicle telematics device, and from a streaming video service used by the end-user.

[0035] In some implementations, data 104 may be tagged or otherwise associated with metadata 106 that describes various aspects of the data 104. For example, a temperature value in the data 104 may be associated with metadata 106 identifying that particular portion of data 104 as a temperature value, or the metadata 106 may indicate the unit of measure of that portion of data 104. In some examples, metadata 106 may describe the source of the data 104, such as the individual platform 102 that generated the data 104, the particular end-user whose interactions with the individual platform 102 caused the generation of the data 104, and so forth. The metadata 106 describing the source platform 102 may enable the super-platform 108 to perform various (e.g., unique or targeted) functions with regard to the data 104. For example, metadata 106 identifying the source platform 102 may enable the data 104 to be retracted from use, or recalled from distribution, if it is later determined that the data 104 from the platform 102 is inaccurate in some way. Metadata 106 may also identify a particular hardware or software component of the platform 102 that generated the data 104. For example, after ingestion of data 104 the individual platform 102 that supplied the data 104 may determine that a particular sensor that generated the data 104 is faulty, and the platform 102 may request that the super-platform 108 prevent further use of data 104 from that particular sensor. As another example, the super-platform 108 may perform operations to remove the faulty data 104 from aggregate data, statistically correct for the inclusion of faulty data 104, or other operations, such as to offset a calibration abnormality of the sensor.

[0036] For a particular portion of data 104, the super-platform 108 may determine as much metadata 106 as possible to describe the data 104. In this way, the super-platform 108 may be described as a metadata-rich environment that enables the data 104 to be interrogated or otherwise analysis in various ways, even in ways not anticipated when the data 104 was generated or received at the super-platform 108. Metadata 106 may indicate a source of the data 104, such as one or both of the end-user or individual platform 102 that generated the data 104. Metadata 106 may also indicate an owner of the data 104 at various stages in the data lifecycle of generation. The data lifecycle may include: ingestion by the super-platform 108; processing (e.g., analysis, aggregation, etc.); storage; distribution to data consumer(s); and deletion or retirement of the data 104. In some implementations, the metadata 106 is formatted according to a Resource Description Framework (RDF) or any other metadata format.

[0037] In some implementations, an individual platform 102 may provide metadata 106 with the data 104 that is ingested into the super-platform 108. In examples where an individual platform 102 does not provide metadata 106, or provides incomplete metadata 106, the super-platform 108

may generate metadata **106** for the ingested data **104**. This generation of metadata **106** may be based on a natural language (NL) or semantic analysis of the ingested data **104**. In some examples, the super-platform **108** may generate metadata **106** such as location tags or time tags for the ingested data **104**, e.g., based on other information indicating a location of the platform **102** or a time when the metadata **106** was generated or received. In some examples, the super-platform **108** may infer or otherwise determine metadata **106** for a portion of data **104** based on a semantic or format-based analysis of the data **104** to match the data **104** with one or more known patterns or formats of data. For example, if the data matches a known pattern or format for (e.g., “looks like”) a value for air pressure, the data **104** may be tagged as air pressure data. The super-platform **108** may also add metadata **106** based on context information that is known about the individual platform **102** or the end-user.

[0038] In some examples, the super-platform **108** may also flag data **104** that the super-platform **108** is unable to reliably tag automatically with an adequate degree of confidence in the tagging. The flagging of data **104** that cannot be tagged may lead to a manual review and manual determination of metadata **106** for the data **104**. In some examples, metadata tagging may be performed at a higher level or with less detail describing the data **104**. For example, it may be sufficient to know that a particular portion of data **104** is describing a movie, without knowing which particular aspect (e.g., title, genre, plot synopsis, cast, etc.) is being described by the data **104**.

[0039] In some implementations, data ingestion includes the normalization of data **104** received from different individual platforms **102**. Such normalization may impose a common format for the metadata **106**, or for the data **104** itself, in situations where different individual platforms **102** provide data **104** or metadata **106** in different formats, styles, data types, and so forth. Normalization may also include determination that different sets of data **104** from different individual platforms **102** are associated with the same end-user, in examples where the different individual platforms **102** employ a different user identifier for the end-user. In such examples, normalization may generate or alter metadata **106** to use a same user identifier to identify the same end-user. Normalization may normalize data **104** received from different countries or other geographical regions, to normalize the data **104** or metadata **106** to employ a same system of measurement, language, or other format.

[0040] In some implementations, the super-platform **108** may determine and employ a semantic data model that describes the ingested data **104** for a single platform **102** or for multiple platforms **102**. The semantic data model may sit on top of the metadata **106** and provide a framework or set of relationships that describes multiple portions of data **104** having different metadata **106**. For example, a semantic data model may describe a movie and may describe relationships between differently metadata-tagged portions of data, such as a relationship between movie title, description, cast, studio, format, length, genre, and so forth. The semantic data model may be described as a structure or map of the data **104**, and may be determined based on the metadata **106**. In some examples, the semantic data model may be hierarchical. For example, data **104** describing a refinery may include multiple gauges, valves, or other components, and those components may have sub-components, which may themselves have sub-components, and so forth.

[0041] The super-platform **108** may be configured to dynamically adapt to the addition of new individual platforms **102** by determining how the data **104** provided by a new individual platform **102** fits into an existing semantic data model that describes previously ingested data **104** from other platform(s) **102**. Capabilities of the newly added individual platform **102** may be determined by analogizing or otherwise comparing to the new platform **102** to existing, similar individual platforms. For example, a new platform **102** may be associated with the semantic data model of a currently supported platform **102** that is similar to the new platform **102** with regard to type of platform, hardware or software configuration, type of service provided by the platform, format of the data, mechanism(s) for accessing and/or receiving the data from the platform, or other characteristics of the platforms. In some implementations, machine learning may be employed to determine the semantic data model and/or other capabilities of a newly available platform **102**. Machine learning algorithm(s) may examine data sets that are sufficiently described by metadata **106** and, using the metadata **106**, may apply a semantic data model to a new data set from a new platform **102** that might extend, augment, improve, or stand alongside an existing data set from current platform(s) **102**.

[0042] The super-platform **108** may also execute one or more software modules to provide a marketplace **112** that enables potential data consumers to request various data set(s) of the collected data **104**, set(s) of aggregate data, and/or result(s) of data analysis job(s). The end-user **116** may access an end-user interface **114** provided by the super-platform **108**. Through the end-user interface **114**, the end-user **116** may specify one or more restrictions **118** on the use, sale, distribution, or other handling of the data **104** associated with the end-user **116**. The end-user **116** may also employ the end-user interface **114** to specify pricing information **120** indicating a price for one or more portions of data **104** associated with the end-user **116**.

[0043] FIG. 2 depicts an example system for selling, distributing, or otherwise providing data **104**, aggregate data, or analysis results through the marketplace **112** to one or more data consumers **204**, in accordance with implementations of the present disclosure. As shown in the example of FIG. 2, the system may include one or more data consumer devices **202**. The data consumer device(s) **202** may include any number and type of computing device. Each of the data consumer device(s) **202** may be owned by, operated by, and/or otherwise associated with data consumer **204**. A data consumer **204** may include any individual, process, system, and/or other entity that seeks to access information provided by the super-platform **108** through the marketplace **112**, such as the data **104**, aggregate data that is generated based on the originally collected data **104**, and/or analysis result(s) **214** of executed analysis job(s). In some examples, the data consumer **204** may purchase information from the marketplace **112**. Implementations also provide for the distribution of information through the marketplace **112** for free, outside a purchase.

[0044] The collected data **104**, aggregate data, and/or analysis result(s) **214** may be offered for sale or other distribution to data consumer(s) **204**, such as data purchaser(s), through the marketplace **112**. In some examples, an end-user **116** may opt in to allow their user-specific data **104** to be sold or otherwise distributed through the marketplace **112**. The end-user **116** may provide a price for their data **104**,

indicated in the pricing information 120. The end-user 116 may also provide restriction(s) 118 on the uses of their data 104 by potential data consumer(s) 204. For example, an end-user 116 may specify that their data may not be provided to data consumer(s) 204 that intend to use the data 104 for a particular purpose, such as research associated with a particular political position, campaign, or cause. As another example, an end-user 116 may specify that their data 104 may not be provided to particular types of organizations. For example, an end-user 116 may state that their data 104 may be sold to non-profit research organizations but not to for-profit corporations. An end-user 116 may also specify the manner in which their data 104 may be used. For example, an end-user 116 may specify that their data 104 may be used after it has been anonymized to remove and/or obfuscate any PII.

[0045] In some examples, the end-user 116 may specify restrictions 118 when they opt in or otherwise give permission for their data 104 to be sold or otherwise distributed, e.g., prior to the distribution or other use of the data 104. In some examples, the end-user 116 may access information presented through the end-user interface 114, the information indicating that their data 104 has been sold or otherwise provided for a particular purpose. Based on this information, the end-user 116 may indicate that their data 104 may not be subsequently sold for a similar purpose. Similarity may be determined based on a semantic (e.g., natural language) analysis of the stated purpose of a data consumer 204. In some implementations, the end-user 116 may have access to information (e.g., in a block chain) describing the various uses to which their data 104 has been put, who it has been sold to, how it has been analyzed, and so forth. Based on such information, the end-user 116 may learn that their data 104 has been provided to a data consumer 204 and the end-user 116 may request some remuneration such as a percentage of the sale price.

[0046] In some examples, the end-user 116 may be able to restrict the sale and use of their data 104 based on a static set of categories. The set of categories may also be dynamic in that the end-user 116 may be able to add new categories for possible restriction of data sale or use. Categories may be arranged hierarchically. For example, an end-user 116 may be able to restrict data usage within a broad category (e.g., political research) or within a more specific category (e.g., political research related to a particular campaign, party, or candidate). The super-platform 108 may generate push notifications or other types of messages that are sent to an end-user 116 to indicate that the end-user's data 104 has been purchased or used for a particular purpose by one or more data consumers 204. In some implementations, the end-user 116 may access the end-user interface 114 in response to receiving the push notifications or other messages, and the end-user 116 may provide further restriction(s) 118 to prevent similar use and/or distribution in the future, or to terminate the current use and/or distribution described in the message(s).

[0047] FIG. 3 depicts an example of restriction(s) 118 on data use and/or distribution, in accordance with implementations of the present disclosure. As shown in the example of FIG. 3, the restriction(s) 118 may include one or more restrictions 118A on a type of use of the data 104. The restriction(s) 118A may prevent the data 104 from being used with respect to one or more restricted use categories 302 and/or one or more restricted use subcategories. For

example, the end-user 116 may specify a restricted use category 302 such that their data 104 is not to be used for political research generally. The end-user 116 may also specify a restricted use subcategory 304 such that their data 104 is not to be used for a more particular type of political research, such as research related to a particular party, candidate, cause, and/or issue.

[0048] The restriction(s) 118 may also include one or more restrictions 118B on the data consumer(s) 204, and/or the types of data consumer(s) 204, that may be provided with the data 104. The restriction(s) 118B may include restricted data consumer categories 306 and/or restricted data consumer subcategories 308. For example, the end-user 116 may specify a restricted data consumer category 306 such that their data 104 is not to be provided for to market research firms generally. The end-user 116 may also specify a restricted data consumer subcategory 308 such that their data 104 is not to be provided to a more particular type of data consumer 204, such as a data consumer 204 that engages in market research associated with pharmaceuticals, cars, or other particular types of products. The end-user 116 may also specify one or more restricted data consumers 310, e.g., particular entities who are not to be provided with the end-user's data 104. The restriction(s) 118B may specify particular characteristic(s) of data consumers 204 that are not to be provided with the data 104 of the end-user 116. For example, the end-user 116 may indicate that data consumer(s) 204 may not receive the data 104 if the data consumer(s) 204 are in a particular location, are engaged in a particular business or other activity, are above (or below) a particular size (e.g., number of employees of a business or other organization), are a for-profit (or non-profit) entity, or based on other characteristic(s).

[0049] The restriction(s) 118 may also include other restriction(s) 118C. Such other restriction(s) 118C may include, but are not limited to: time constraints, indicating a particular time period during which the data 104 may be accessed and used; constraints indicating whether the data 104 may be used in an aggregate form with other data, or whether the data 104 may be used in its individual, originally collected, form; constraints indicating whether the data 104 is to be anonymized prior to sale, distribution, and/or use; and so forth.

[0050] Although examples herein may describe the restriction(s) 118 as a black list, e.g., a set of restricted uses and/or data consumer(s) 204, implementations are not so limited. In some implementations, the restriction(s) 118 may be a white list, in which the end-user 116 specifies a list of allowed uses and/or allowed data consumer(s) 204. In such examples, a requested use may be blocked or a requesting data consumer 204 may be denied access to the data 104 unless the use and/or the data consumer 204 is included in the white list of allowed uses and/or data consumer(s) 204. To account for a variety of possible restrictions on data use, the data and/or its permissions for use may be stored in an immutable block chain or other suitable type of immutable data storage. This allows the data to retain its complete provenance from its initial creation through every step that the data takes along its journey. Such provenance information may be stored on block chain(s), and may provide the various reasons for the sharing of the data and its use by various systems, processes, and/or individuals. By storing data and its restrictions based on disclosure of (e.g., human or machine) preferences in this way, implementations facili-

tate auditing as a feature of this system. Using a block chain also allows data owners to dynamically change permissions that govern the access to and use of existing data in the system. While the use of data in this block chain can be openly audited, the data assets themselves may not be retrievable outside of the super-platform without a client-side daemon on the data consumer's analytics system. Such a daemon may specify the conditions under which the data will be used, download the data assets from the super-platform, decrypt these data assets, and/or append the meta-data for each data asset on the block chain with the specific use. To audit for use compliance of specific data assets, a simple query of the block chain may return this information.

[0051] Returning to FIG. 2, in some examples the data seller may be an end-user 116. In some examples, the data seller may be an individual platform 102 through which the end-user's data 104 is collected. In some examples, an individual platform 102 may be able to sell data 104 that is not (e.g., directly) associated with a particular end-user 116 and/or that does not identify an end-user 116, such as sensor data (e.g., temperature data, location data, etc.) collected at the platform 102.

[0052] A potential data consumer 204 may access the marketplace 112 and request information for purchase or non-purchase distribution. In some examples, the data consumer 204 may provide data set request(s) 208 to the marketplace 112, each data set request 208 specifying particular data that is needed, such as a certain number of records for end-users that match a particular demographic profile or particular characteristics. The data set request 208 may also specify an intended use for the data, a price range, and/or information regarding the data consumer 204. The marketplace 112 may search available data 104 and/or aggregate data and present the data consumer 204 with one or more possible data sets 210 that match the specified criteria and the intended use, and that comply with any restrictions 118. The marketplace 112 may also present the price of the data sets 210, e.g., as a sum of the prices specified by the individual end-users 116 whose data 104 is included the data set(s) 210 or has been aggregated to generate the data set(s) 210. The data consumer 204 may select one or more of the data sets 210 to purchase or otherwise receive access to.

[0053] In some examples, the data consumer 204 may provide one or more analysis requests 212. An analysis request 212 may describe an insight, or an answer to a question, sought by the data consumer 204. An analysis request 212 may be specified using natural language (e.g., "tell me about the buying habits that are unique to people in Bakersfield, California"). In response to receiving an analysis request 212, the marketplace 112 may perform a semantic analysis of the language used in the analysis request 212 and identify one or more data sets 210 that may be useful to the data consumer 204 (e.g., purchase history of users in Bakersfield), and such data set(s) 210 may be offered to the data consumer 204 for purchase. In some examples, in response to receiving an analysis request 212, the marketplace 112 may determine one or more analytic jobs that may be run on the data 104 to provide the desired answer or insight. For example, an analysis job may be a correlation or other type of statistical analysis to identify buying habits unique to the Bakersfield area. The marketplace 112 may indicate a price for the data set(s) 210 and/or the analytic job(s) that address the analysis request 212. In some examples, the data con-

sumer 204 may be charged for the analysis job itself. In such examples, the analysis job may be executed on the super-platform 108 and the data consumer 204 may be presented with the analysis result(s) 214 resulting from the execution of the analysis job. In some examples, the analysis job may be provided to the data consumer 204 for execution on the data consumer device 202. In such examples, data purchasers may be required to use an agent and/or daemon on the computing device (e.g., data appliance) performing the analytics. As part of this workflow, data purchasers may be required to specify the intended use for the analysis and subject their results to audit by the super-platform. In some implementations, machine learning is employed to automate the auditing of stated use compared to actual use, and may also be employed to audit for compliance with usage constraints from the portfolio of data disclosers. When a purchaser buys data, super-platform based and/or cloud-based analysis job(s) may be suggested at varying degrees of accuracy or fidelity. For example, a more accurate analysis job may run longer, provide a more accurate answer, and cost more, whereas a less accurate analysis job may provide a faster, cheaper, less accurate answer to the data consumer 204.

[0054] Implementations support various use cases for data analysis. In some instances, if data analysis is to be run "on premises" on the data consumer's data appliance, then the client-side daemon may be required to be running on the appliance to ensure compliance with restrictions. Alternatively, if data analysis is to be run in the cloud and/or by the super-platform, then an auditing capability may be included in the cloud server(s) and/or super-platform where the analysis is running. In such instances, the level of analysis scrutiny that is directed to the data consumer (e.g., purchaser) may be reduced. In some examples, the same or similar level of scrutiny may be applied in both use cases, but the levels of control in these two use cases may be different.

[0055] As described above, the marketplace 112 may perform aggregation or other types of analyses to generate aggregate data sets based on data from multiple individual platforms 102 or multiple end-users 116. The marketplace 112 may aggregate user data according to particular demographic characteristics, location, and/or other characteristics. For example, the marketplace 112 may collect user data from users in Iowa and aggregate the data to create a data set describing the buying habits, browsing habits, or other characteristics of users in a particular market segment defined by location, age, gender, or other criteria. In some examples, the aggregate data sets may be generated such that they do not include any PII. In some examples, the aggregate data sets may be generated based on particular requests from data consumers 204 such as data set request(s) 208 and/or analysis request(s) 212. The marketplace 112 may also generate aggregate data sets prior to, and in anticipation of, receiving any particular data set request(s) 208 and/or analysis request(s) 212.

[0056] In some implementations, an agent 206 may execute on the data consumer device(s) 202. The agent 206 may be act as an interface, filter, and/or search engine that enables the data consumer 204 to submit data set request(s) 208 and purchase data set(s) 210. The agent 206 may also enable the data consumer 204 to submit analysis request(s) 212, purchase analysis job(s) to be executed on the super-platform 108, receive the analysis result(s) 214, and/or view

the analysis result(s) 214. In examples where the super-platform 108 provides the analysis job(s) to the data consumer device(s) 202, the agent 206 may provide an execution container or otherwise manage the execution of the analysis job(s) on the data consumer device(s) 202.

[0057] The agent 206 may validate the data set(s) 210 and/or analysis result(s) 214 received at the data consumer device 202. In some examples, the agent 206 may perform digital rights management (DRM) operations to ensure that only authorized entities have access to the data set(s) 210, analysis result(s) 214, and/or analysis job(s). In some implementations, the permissions to retrieve data, analysis results, and/or analysis jobs may be incorporated into the block chain (or other data structure) and can be queried by the super-platform and/or agent (e.g., daemon) executing on the purchaser's device(s). When a purchaser attempts to retrieve results and/or associated data, the permissions (which may include data discloser preferences and/or licensing agreements) may be first verified through metadata in the block chain before being delivered to the requestor. The agent 206 may also monitor the use of any received information to ensure that it is only being used according to the originally stated purpose of the data consumer 204. For example, if the data consumer 204 purchased data under conditions that the data would not be transferred to a third party, the agent 206 may ensure that the data is not transferred. If the data consumer 204 purchased data with the stated purpose of analyzing the data to determine result X, the agent 206 may ensure that the data is not analyzed to determine a different result Y. The agent 206 may also perform decryption and/or decompression operations to enable access to the data, in examples where the data is provided in an encrypted and/or compressed form.

[0058] The marketplace 112 may facilitate two-way transactions involving information such as the data 104, aggregate data, and/or analysis result(s) 214. For example, data may be purchased from an end-user 116 and/or an individual platform 102, and the data may be sold to data consumer(s) 204. A data consumer 204 may also provide results of their analysis of the data back to the marketplace 112, and such results may be sold to other data consumers 204.

[0059] The super-platform 108 may create an adhesive experience for data consumers 204, such that data consumer (s) 204 are incentivized to return to the super-platform 108 for their future data needs. This adhesive experience may be achieved by providing data consumers 204 with data 104, aggregate data, analysis job(s), and/or data analysis result(s) 214 that the data consumer(s) 204 may not have access to elsewhere. By sharing or aggregating data 104 generated by multiple individual platforms 102, the super-platform 108 may provide greater value than that provided by the individual platform(s) 102 alone.

[0060] Various actions may be performed in instances where the ingested data 104 is later determined to be inaccurate or fraudulent, after the data 104 has been distributed to data consumer(s) 204 and/or after aggregate data or analysis result(s) 214 based on the data 104 have been distributed. In some examples, the data 104 may be associated with metadata 106 that indicates a source, version, or timestamp for the data 104. The super-platform 108 may provide an interface, e.g., through the agent 206, that enables the data consumer(s) 204 to check existing purchases and/or other data distributions for changes or alerts regarding the integrity of the data 104 generated by various platforms 102.

The super-platform 108 may notify the data consumer(s) 204 that previous distributed information may be inaccurate in some way, and may provide the data consumer(s) 204 with the option to received updated data that at least partly corrects the flaws of the previously distributed data. Such updates may be implemented using a content management system with version control. In some examples, the individual platform(s) 102 may inform the super-platform 108 that previously provided data 104 is inaccurate. In some examples, the super-platform 108 may be configured to determine that data 104 previously received at the super-platform 108 is inaccurate.

[0061] FIG. 4 depicts an example process for ingesting data 104 from multiple data generating platforms 102, in accordance with implementations of the present disclosure. Operations of the process may be performed by one or more of the gateway(s) 110, the marketplace 112, the end-user interface 114, the agent 206, or other software module(s) executing on the platform(s) 102, the super-platform 108, the data consumer device(s) 202, or elsewhere.

[0062] An indication may be received (402), e.g., at the super-platform 108, of data 104 that is available from individual platform(s) 102. A determination may be made (404) whether the data 104 is saleable. In some implementations, this determination may be based at least partly on whether the end-user 116 has indicated that the data 104 is saleable and/or indicated a price for the data 104.

[0063] If the data 104 is saleable, the super-platform 108 may request and receive (406) pricing information 120 from the end-user 116 and/or from the individual platform(s) 102. The process may proceed to 408. If the data 104 is not saleable, the process may proceed to 408.

[0064] In some examples, one or more restrictions 118 may be determined (408). Such restriction(s) 118 may be specified by the end-user 116, the platform(s) 102, or both the end-user 116 and the platform(s) 102. The data 104 may be ingested (410) or otherwise received by the super-platform 108 for storage, analysis, aggregation and/or distribution. In some implementations, the data 104 may be aggregated and/or analyzed (412) as described above. The data 104 and/or aggregate data may be stored (414) to be available for sale or other distribution through the marketplace 112.

[0065] FIG. 5 depicts an example process for providing data set(s) 210 to data consumer(s) 204 in response to data set request(s) 208, in accordance with implementations of the present disclosure. Operations of the process may be performed by one or more of the gateway(s) 110, the marketplace 112, the end-user interface 114, the agent 206, or other software module(s) executing on the platform(s) 102, the super-platform 108, the data consumer device(s) 202, or elsewhere.

[0066] A data set request 208 may be received (502) at the marketplace 112. In some examples, the data set request 208 may be provided by a data consumer 204 through the agent 206, and the agent may facilitate the interactions between the data consumer 204 and the marketplace 112. The data set request 208 may indicate the intended use of the data, identify the data consumer 204, and/or provide further information regarding the data consumer 204 (e.g., the type of organization or entity, name, address, location, etc.).

[0067] The marketplace 112 may analyze the data set request 208 and determine (504) one or more data set(s) 210 to satisfy the data set request 208. In some examples, the

data set(s) 210 may be determined to ensure compliance with any restriction(s) 118 on the use and/or distribution of the data in the data set(s) 210. For example, the data set(s) 210 may include data 104 that is not restricted from being used in the manner stated by the data consumer 204 in the data set request 208.

[0068] A determination may be made (506) whether the data is saleable. If so, the price of the data set(s) 210 may be determined (508) and payment may be collected (510) from the data consumer 204. Payment may include the transfer of funds (e.g., monetary value) or the transfer of any other type of value. If the data is not saleable, the process may proceed to 512.

[0069] The data set(s) 210 may be sent (512) to the data consumer 204. In some implementations, the data set(s) 210 may be communicated to the agent 206 and the agent 206 may present the data set(s) 210 to the data consumer 204. In some implementations, the use of the data set(s) 210 by the data consumer 204 may be monitored (514) by the agent 206 to ensure that the use complies with any restriction(s) 118 on the data.

[0070] FIG. 6 depicts an example process for providing analysis result(s) 214 of analysis job(s), and/or the analysis job(s) themselves, to data consumer(s) 204, in accordance with implementations of the present disclosure. Operations of the process may be performed by one or more of the gateway(s) 110, the marketplace 112, the end-user interface 114, the agent 206, or other software module(s) executing on the platform(s) 102, the super-platform 108, the data consumer device(s) 202, or elsewhere.

[0071] An analysis request 212 may be received (602) at the marketplace 112. In some examples, the analysis request 212 may be provided by a data consumer 204 through the agent 206. The analysis request 212 may indicate a question to which the data consumer 204 seeks an answer, and/or an insight sought by the data consumer 204. The analysis request 212 may indicate the intended use of the analysis result(s) 214, identify the data consumer 204, and/or provide further information regarding the data consumer 204 (e.g., the type of organization or entity, name, address, location, etc.).

[0072] The marketplace 112 may analyze the analysis request 212 and determine (604) one or more analysis jobs that may be performed to satisfy the analysis request 212. In some examples, the analysis job(s) may be determined to ensure compliance with any restriction(s) 118 on the use and/or distribution of the data in the data set(s) 210. For example, the analysis job(s) may be determined to analyze data 104 that is not restricted from being used in the manner stated by the data consumer 204 in the analysis request 212.

[0073] A determination may be made (606) whether the data is saleable. If so, the price of the data underlying the analysis job(s) may be determined (608) and payment may be collected (610) from the data consumer 204. If the data is not saleable, the process may proceed to 612.

[0074] In some implementations, the analysis job(s) may be executed on the super-platform 108 and the analysis result(s) 214 may be sent (612) to the data consumer 204. In some implementations, the analysis result(s) 214 may be communicated to the agent 206 and the agent 206 may present the analysis result(s) 214 to the data consumer 204. In some implementations, the analysis job(s) may be communicated (612) to the agent 206. The agent 206 may

execute the analysis job(s) and present the analysis result(s) 214 to the data consumer 204.

[0075] In some implementations, the use of the analysis result(s) 214 by the data consumer 204 may be monitored (614) by the agent 206 to ensure that the use complies with any restriction(s) 118 on the data. In examples where the analysis job(s) are sent to the agent 206, the agent 206 may ensure that the analysis job(s) are executed locally and the analysis result(s) 214 are used in accordance with the restriction(s) 118.

[0076] FIG. 7 depicts an example computing system 700 in accordance with implementations of the present disclosure. The system 700 may be used for any of the operations described with respect to the various implementations discussed herein. For example, the system 700 may be included, at least in part, in one or more of platforms 102, the super-platform 108, or the data consumer device(s) 202 described herein. The system 700 may include one or more processors 710, a memory 720, one or more storage devices 730, and one or more input/output (I/O) devices 750 controllable using one or more I/O interfaces 740. Two or more of the components 710, 720, 730, 740, or 750 may be interconnected using at least one system bus 760, which may enable the transfer of data between the various modules and components of the system 700.

[0077] The processor(s) 710 may be configured to process instructions for execution within the system 700. The processor(s) 710 may include single-threaded processor(s), multi-threaded processor(s), or both. The processor(s) 710 may be configured to process instructions stored in the memory 720 or on the storage device(s) 730. The processor(s) 710 may include hardware-based processor(s) each including one or more cores. The processor(s) 710 may include general purpose processor(s), special purpose processor(s), or both.

[0078] The memory 720 may store information within the system 700. In some implementations, the memory 720 includes one or more computer-readable media. The memory 720 may include any number of volatile memory units, any number of non-volatile memory units, or both volatile and non-volatile memory units. The memory 720 may include read-only memory, random access memory, or both. In some examples, the memory 720 may be employed as active or physical memory by one or more executing software modules.

[0079] The storage device(s) 730 may be configured to provide (e.g., persistent) mass storage for the system 700. In some implementations, the storage device(s) 730 may include one or more computer-readable media. For example, the storage device(s) 730 may include a floppy disk device, a hard disk device, an optical disk device, or a tape device. The storage device(s) 730 may include read-only memory, random access memory, or both. The storage device(s) 730 may include one or more of an internal hard drive, an external hard drive, or a removable drive.

[0080] One or both of the memory 720 or the storage device(s) 730 may include one or more computer-readable storage media (CRSM). The CRSM may include one or more of an electronic storage medium, a magnetic storage medium, an optical storage medium, a magneto-optical storage medium, a quantum storage medium, a mechanical computer storage medium, and so forth. The CRSM may provide storage of computer-readable instructions describing data structures, processes, applications, programs, other

modules, or other data for the operation of the system **700**. In some implementations, the CRSM may include a data store that provides storage of computer-readable instructions or other information in a non-transitory format. The CRSM may be incorporated into the system **700** or may be external with respect to the system **700**. The CRSM may include read-only memory, random access memory, or both. One or more CRSM suitable for tangibly embodying computer program instructions and data may include any type of non-volatile memory, including but not limited to: semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. In some examples, the processor(s) **710** and the memory **720** may be supplemented by, or incorporated into, one or more application-specific integrated circuits (ASICs).

[0081] The system **700** may include one or more I/O devices **750**. The I/O device(s) **750** may include one or more input devices such as a keyboard, a mouse, a pen, a game controller, a touch input device, an audio input device (e.g., a microphone), a gestural input device, a haptic input device, an image or video capture device (e.g., a camera), or other devices. In some examples, the I/O device(s) **750** may also include one or more output devices such as a display, LED(s), an audio output device (e.g., a speaker), a printer, a haptic output device, and so forth. The I/O device(s) **750** may be physically incorporated in one or more computing devices of the system **700**, or may be external with respect to one or more computing devices of the system **700**.

[0082] The system **700** may include one or more I/O interfaces **740** to enable components or modules of the system **700** to control, interface with, or otherwise communicate with the I/O device(s) **750**. The I/O interface(s) **740** may enable information to be transferred in or out of the system **700**, or between components of the system **700**, through serial communication, parallel communication, or other types of communication. For example, the I/O interface(s) **740** may comply with a version of the RS-232 standard for serial ports, or with a version of the IEEE 1284 standard for parallel ports. As another example, the I/O interface(s) **740** may be configured to provide a connection over Universal Serial Bus (USB) or Ethernet. In some examples, the I/O interface(s) **740** may be configured to provide a serial connection that is compliant with a version of the IEEE 1394 standard.

[0083] The I/O interface(s) **740** may also include one or more network interfaces that enable communications between computing devices in the system **700**, or between the system **700** and other network-connected computing systems. The network interface(s) may include one or more network interface controllers (NICs) or other types of transceiver devices configured to send and receive communications over one or more networks using any network protocol.

[0084] Computing devices of the system **700** may communicate with one another, or with other computing devices, using one or more networks. Such networks may include public networks such as the internet, private networks such as an institutional or personal intranet, or any combination of private and public networks. The networks may include any type of wired or wireless network, including but not limited to local area networks (LANs), wide area networks (WANs), wireless WANs (WWANs), wireless LANs (WLANs), mobile communications networks (e.g., 3G, 4G,

Edge, etc.), and so forth. In some implementations, the communications between computing devices may be encrypted or otherwise secured. For example, communications may employ one or more public or private cryptographic keys, ciphers, digital certificates, or other credentials supported by a security protocol, such as any version of the Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) protocol.

[0085] The system **700** may include any number of computing devices of any type. The computing device(s) may include, but are not limited to: a personal computer, a smartphone, a tablet computer, a wearable computer, an implanted computer, a mobile gaming device, an electronic book reader, an automotive computer, a desktop computer, a laptop computer, a notebook computer, a game console, a home entertainment device, a network computer, a server computer, a mainframe computer, a distributed computing device (e.g., a cloud computing device), a microcomputer, a system on a chip (SoC), a system in a package (SiP), and so forth. Although examples herein may describe computing device(s) as physical device(s), implementations are not so limited. In some examples, a computing device may include one or more of a virtual computing environment, a hypervisor, an emulation, or a virtual machine executing on one or more physical computing devices. In some examples, two or more computing devices may include a cluster, cloud, farm, or other grouping of multiple devices that coordinate operations to provide load balancing, failover support, parallel processing capabilities, shared storage resources, shared networking capabilities, or other aspects.

[0086] The features described may be implemented in digital electronic circuitry or in computer hardware, software, or any combinations of hardware and software. The features may be implemented in at least one computer program product that is tangibly embodied in an information carrier (e.g., a machine-readable storage device) for execution by a programmable processor. The method or process steps may be performed by a programmable processor executing instructions to perform functions of the described implementations by operating on input data and generating output. The described features may be implemented in one or more computer programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. A computer program may include a set of instructions that can be used, directly or indirectly, in a computer to perform a certain activity or bring about a certain result. A computer program may be written in any form of programming language, including compiled or interpreted languages, and it may be deployed in any form, including as a stand-alone program or as a module suitable for use in a computing environment.

[0087] The logic flows depicted in the figures do not require the particular order shown, or any particular sequential order, to achieve desirable results. In some implementations, other steps may be provided or steps may be eliminated from the described flows. Moreover, the steps may be performed in parallel or serially with respect to other steps. The systems depicted in the figures do not require the particular components, or the particular arrangement of components, shown in the figures. In some implementations, the various systems may include more or fewer components

than shown in the figures, and components may be arranged differently to achieve desirable results. Accordingly, implementations other than those explicitly depicted in the figures or described herein are within the scope of the following claims.

[0088] A number of implementations of the present disclosure have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the present disclosure. Accordingly, other implementations are within the scope of the following claims.

What is claimed is:

1. A computer-implemented method performed by at least one processor, the method comprising:

receiving, by the at least one processor, data generated at a plurality of individual platforms, the data associated with one or more end-users of the plurality of individual platforms;

receiving, by the at least one processor, a request for at least a portion of the data, the request provided by a data consumer through a marketplace; and

providing, by the at least one processor, at least the portion of the data to the data consumer at a price that is based at least partly on pricing information provided by the one or more end-users associated with the data.

2. The method of claim 1, further comprising:

aggregating, by the at least one processor, the data from multiple individual platforms to generate aggregate data; and

providing, by the at least one processor, the aggregate data to the data consumer at the price.

3. The method of claim 1, further comprising:

receiving, by the at least one processor, an analysis request provided by the data consumer through the marketplace;

executing, by the at least one processor, at least one analysis job that is determined to satisfy the analysis request; and

providing, by the at least one processor, at least one analysis result of the at least one analysis job to the data consumer.

4. The method of claim 1, further comprising:

receiving, by the at least one processor, an analysis request provided by the data consumer through the marketplace;

determining, by the at least one processor, at least one analysis job to satisfy the analysis request; and

providing, by the at least one processor, the at least one analysis job for execution on a data consumer device of the data consumer.

5. The method of claim 1, wherein the plurality of individual platforms include at least two of:

an e-commerce platform configured to provide an e-commerce service;

a retail platform configured to support retail operations of a physical store;

a media platform configured to provide an online media service;

a mobile computing device;

an Internet of Things (IoT) computing device; or

an in-vehicle computer.

6. The method of claim 1, further comprising:

receiving, by the at least one processor, one or more restrictions on the use of the data, the one or more restrictions provided by the one or more end-users;

receiving, by the at least one processor, from the data consumer, an indicated use of the data; and

determining, by the at least one processor, at least the portion of the data to be provided to the data consumer, such that the indicated use complies with the one or more restrictions on the use of at least the portion of the data.

7. The method of claim 6, further comprising:

monitoring, by the at least one processor, the use of at least the portion of the data provided to the data consumer to ensure compliance with the one or more restrictions.

8. The method of claim 1, further comprising:

receiving, by the at least one processor, one or more restrictions on data consumers to be provided the data, the one or more restrictions provided by the one or more end-users;

determining, by the at least one processor, at least one characteristic of the data consumer; and

determining, by the at least one processor, at least the portion of the data to be provided to the data consumer, such that the at least one characteristic complies with the one or more restrictions on data consumers to be provided the data.

9. A system comprising:

at least one processor; and

a memory communicatively coupled to the at least one processor, the memory storing instructions which, when executed by the at least one processor, cause the at least one processor to perform operations comprising:

receiving data generated at a plurality of individual platforms, the data associated with one or more end-users of the plurality of individual platforms;

receiving a request for at least a portion of the data, the request provided by a data consumer through a marketplace; and

providing at least the portion of the data to the data consumer at a price that is based at least partly on pricing information provided by the one or more end-users associated with the data.

10. The system of claim 9, the operations further comprising:

aggregating the data from multiple individual platforms to generate aggregate data; and

providing the aggregate data to the data consumer at the price.

11. The system of claim 9, the operations further comprising:

receiving an analysis request provided by the data consumer through the marketplace;

executing at least one analysis job that is determined to satisfy the analysis request; and

providing at least one analysis result of the at least one analysis job to the data consumer.

12. The system of claim 9, the operations further comprising:

receiving an analysis request provided by the data consumer through the marketplace;

determining at least one analysis job to satisfy the analysis request; and
 providing the at least one analysis job for execution on a data consumer device of the data consumer.

13. The system of claim **9**, wherein the plurality of individual platforms include at least two of:

- an e-commerce platform configured to provide an e-commerce service;
- a retail platform configured to support retail operations of a physical store;
- a media platform configured to provide an online media service;
- a mobile computing device;
- an Internet of Things (IoT) computing device; or
- an in-vehicle computer.

14. The system of claim **9**, the operations further comprising:

- receiving one or more restrictions on the use of the data, the one or more restrictions provided by the one or more end-users;
- receiving, from the data consumer, an indicated use of the data; and
- determining at least the portion of the data to be provided to the data consumer, such that the indicated use complies with the one or more restrictions on the use of at least the portion of the data.

15. The system of claim **14**, the operations further comprising:

- monitoring the use of at least the portion of the data provided to the data consumer to ensure compliance with the one or more restrictions.

16. One or more computer-readable storage media storing instructions which, when executed by at least one processor, cause the at least one processor to perform operations comprising:

- receiving data generated at a plurality of individual platforms, the data associated with one or more end-users of the plurality of individual platforms;
- receiving a request for at least a portion of the data, the request provided by a data consumer through a marketplace; and

- providing at least the portion of the data to the data consumer at a price that is based at least partly on pricing information provided by the one or more end-users associated with the data.

17. The one or more computer-readable storage media of claim **16**, the operations further comprising:

- aggregating the data from multiple individual platforms to generate aggregate data; and
- providing the aggregate data to the data consumer at the price.

18. The one or more computer-readable storage media of claim **16**, the operations further comprising:

- receiving an analysis request provided by the data consumer through the marketplace;
- executing at least one analysis job that is determined to satisfy the analysis request; and
- providing at least one analysis result of the at least one analysis job to the data consumer.

19. The one or more computer-readable storage media of claim **16**, the operations further comprising:

- receiving an analysis request provided by the data consumer through the marketplace;
- determining at least one analysis job to satisfy the analysis request; and
- providing the at least one analysis job for execution on a data consumer device of the data consumer.

20. The one or more computer-readable storage media of claim **16**, the operations further comprising:

- receiving one or more restrictions on data consumers to be provided the data, the one or more restrictions provided by the one or more end-users;
- determining at least one characteristic of the data consumer; and
- determining at least the portion of the data to be provided to the data consumer, such that the at least one characteristic complies with the one or more restrictions on data consumers to be provided the data.

* * * * *