



(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2019/0373077 A1**

Lepore et al. (43) **Pub. Date: Dec. 5, 2019**

(54) **SYSTEM FOR ASSOCIATION OF CUSTOMER INFORMATION ACROSS SUBSCRIBERS**

(52) **U.S. Cl.**
CPC **H04L 67/306** (2013.01); **H04L 63/12** (2013.01); **H04L 67/12** (2013.01)

(71) Applicant: **Cisco Technology, Inc.**, San Jose, CA (US)

(57) **ABSTRACT**

(72) Inventors: **Michael P. Lepore**, Marlborough, MA (US); **Stephen R. Quatrano**, Lexington, MA (US); **Tod Lloyd Famous**, Ayer, MA (US)

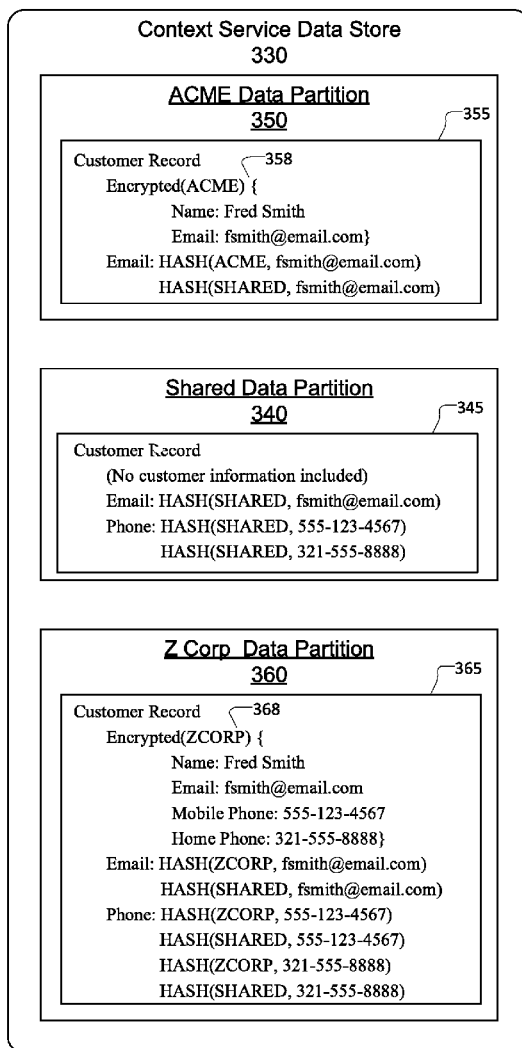
The disclosed technology relates to a context service system configured to receive, from a subscriber, a shared customer lookup request that includes a first customer data identifier and identify, in a shared data partition, a second customer data identifier associated with the first customer data identifier. The context service system is further configured to determine that the second customer data identifier is associated with customer information in a subscriber data partition and transmit, to the subscriber system, the customer information from the subscriber data partition.

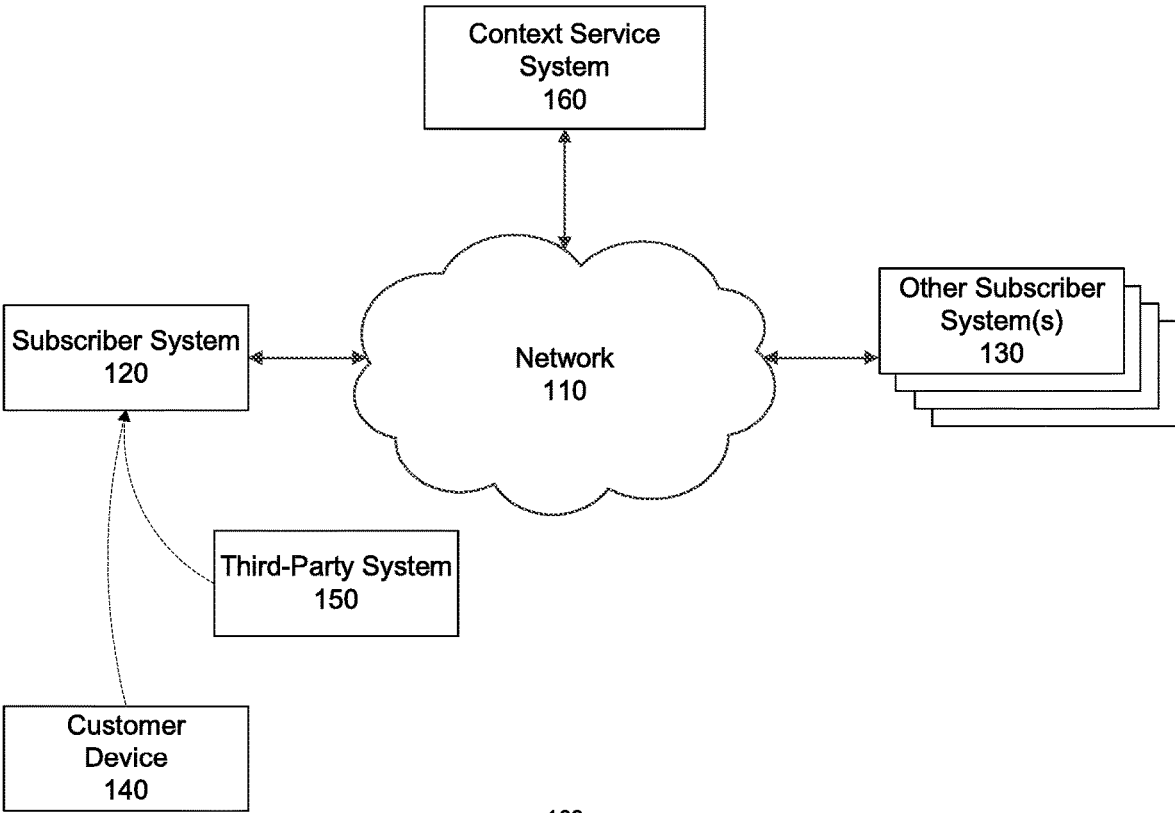
(21) Appl. No.: **15/992,048**

(22) Filed: **May 29, 2018**

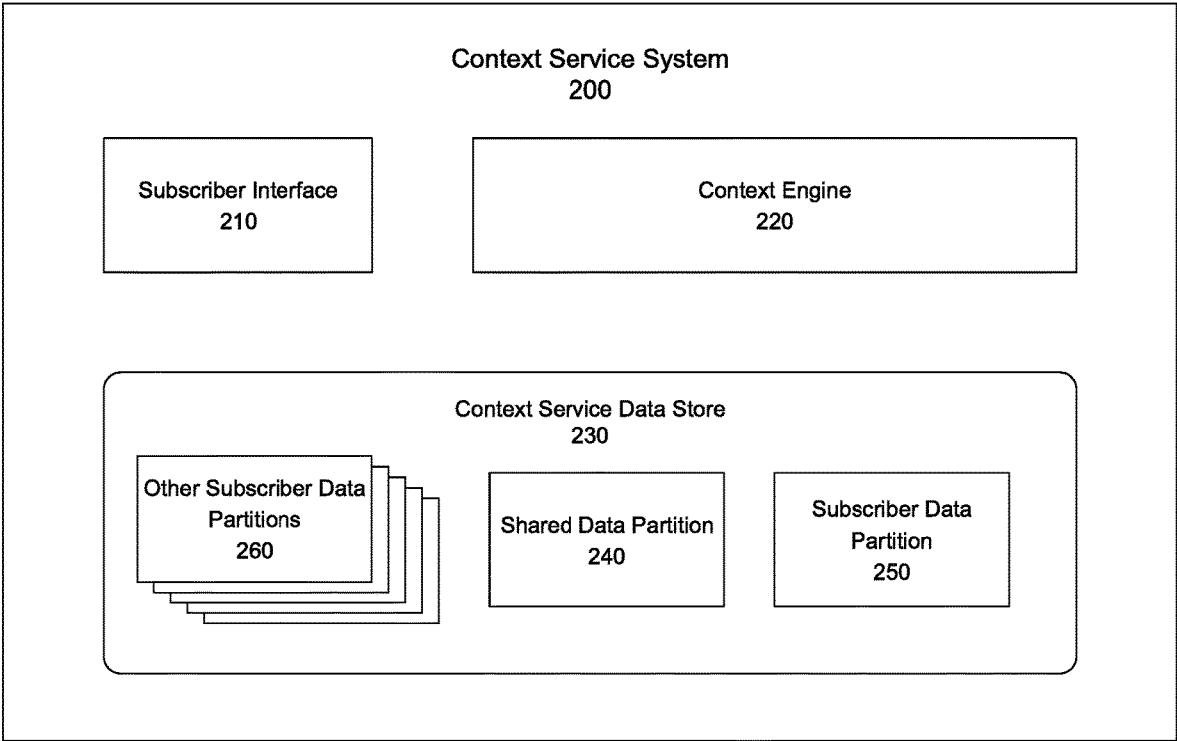
Publication Classification

(51) **Int. Cl.**
H04L 29/08 (2006.01)
H04L 29/06 (2006.01)



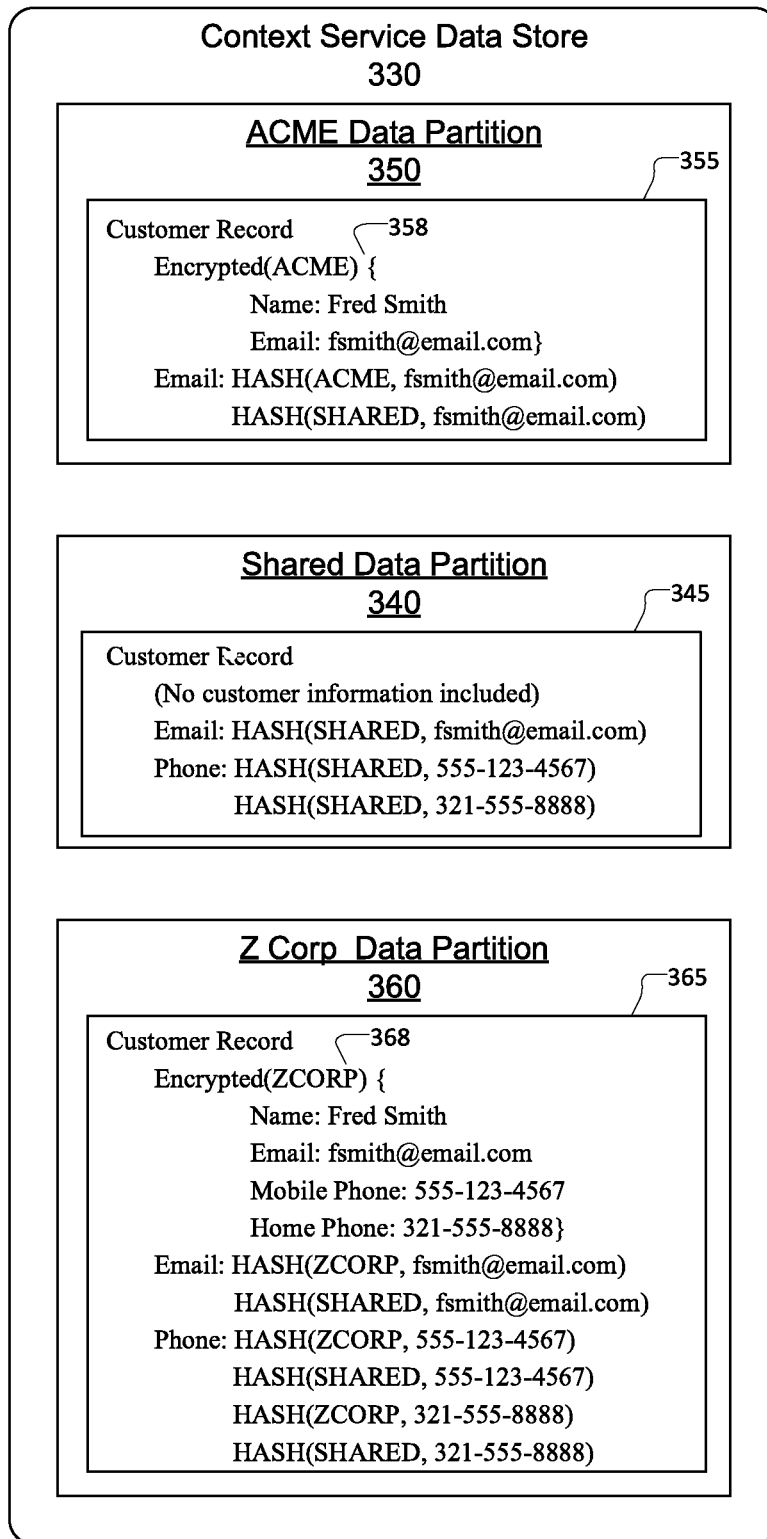


100
FIG. 1



200

FIG. 2



300

FIG. 3

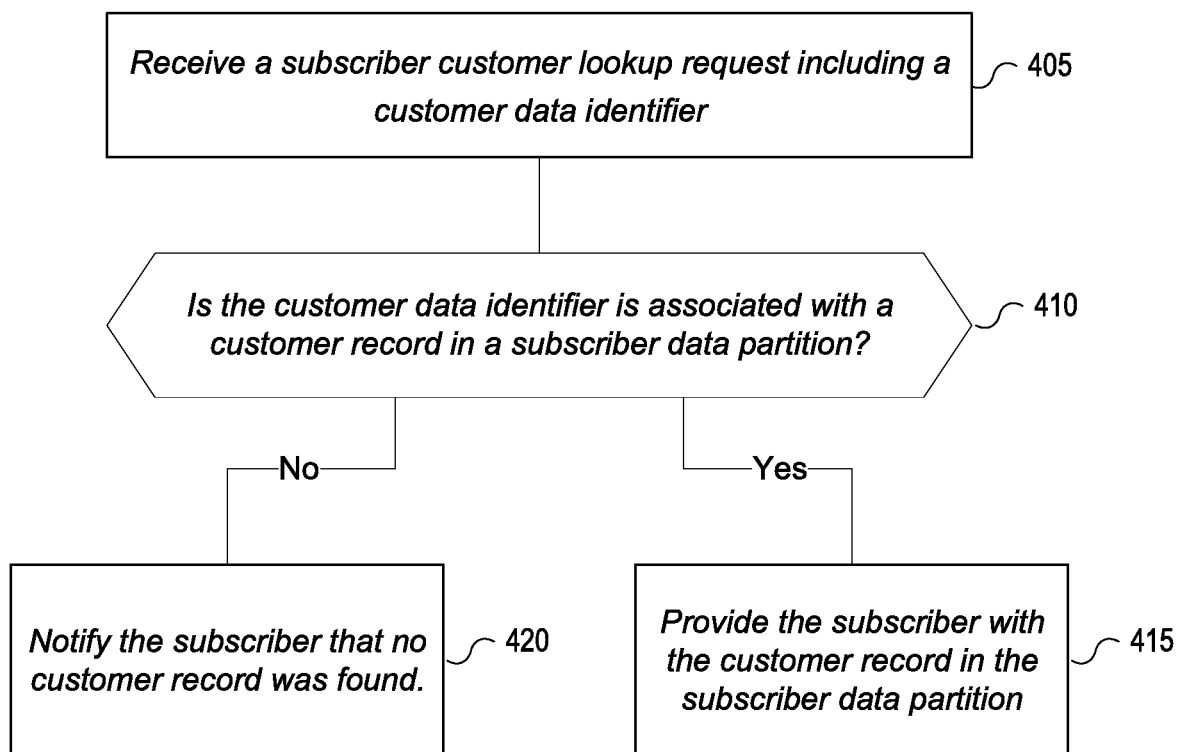
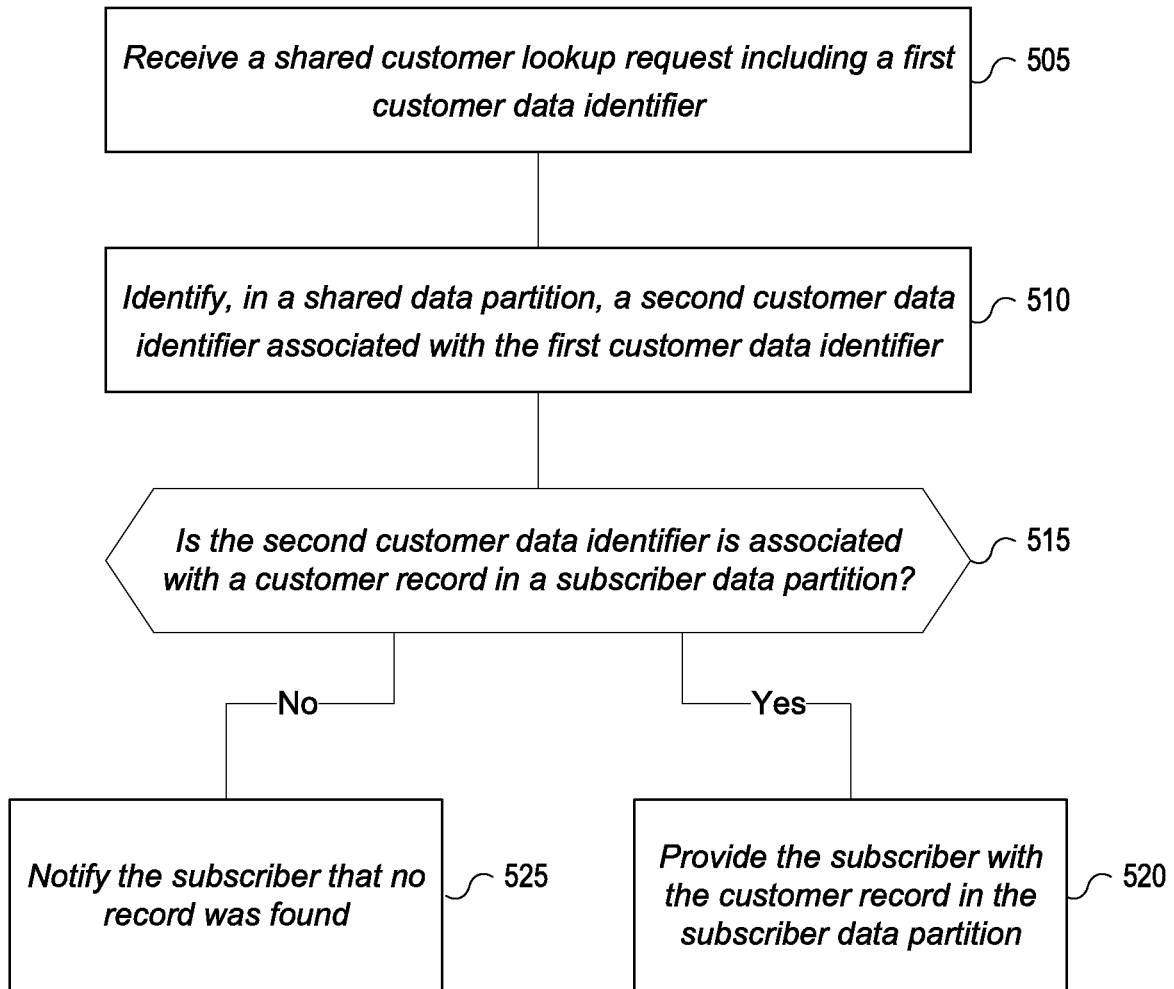
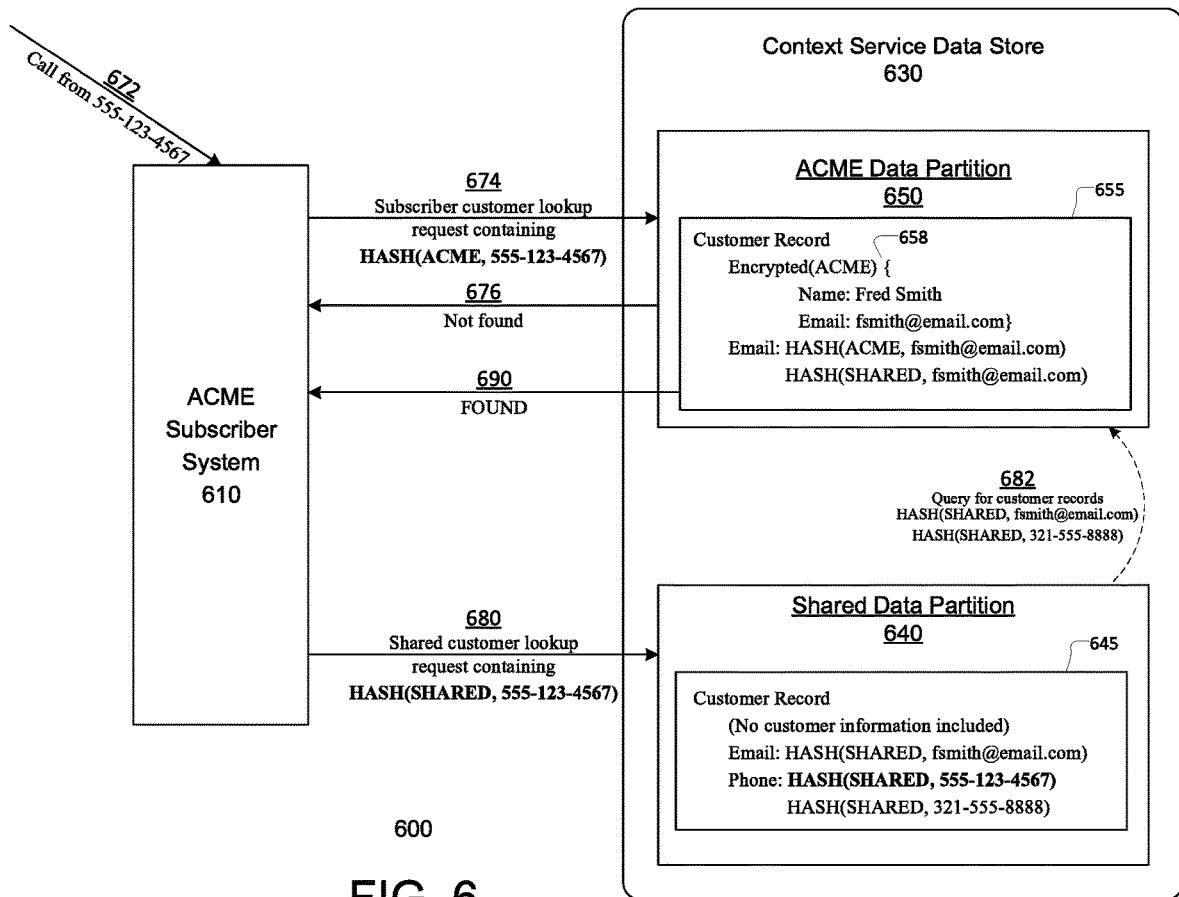


FIG. 4



500
FIG. 5






FIG. 7A




FIG. 7B

SYSTEM FOR ASSOCIATION OF CUSTOMER INFORMATION ACROSS SUBSCRIBERS

TECHNICAL FIELD

[0001] The subject matter of this disclosure relates in general to the field of computerized management of customer information, and more specifically to a platform for identifying associations between customer information.

BACKGROUND

[0002] Businesses and other organizations often store information associated with customers or other individuals or entities in order to better serve those customers or for some other purpose. For example, an e-commerce platform may store a customer's name and mailing address for shipping products to the customer, account name and password to enable a customer to login to an e-commerce website or for other services, age, and/or an order history. Some e-commerce platforms may use the customer information in order to provide better and more customized service such as suggesting other products that the customer may be interested in. Other types of platforms (e.g., social media platforms, financial services platforms, etc.) may store user profile information, activity information on the platform or other platforms, or other information.

[0003] However, in many cases, the information collected may be incomplete or fragmented. For example, the e-commerce platform in the example above may receive a call from customer from a phone number. Although the e-commerce platform has the customer's phone number, it is unable to retrieve other information associated with the customer (e.g., the customer name, mailing address, or order history) because the e-commerce platform does not have the customer's phone number in the customer record. Even though the e-commerce platform has access to the phone number and other customer information, the e-commerce platform is unable to leverage the customer information because the customer information is fragmented and there is no association between the phone number and the rest of the customer information.

[0004] The fragmentation of customer information may lead to a degraded level of service and/or an increase the time and effort needed to provide services to customers, which increases costs to the business. For example, instead of being able to automatically retrieve the customer record, a customer service representative of the e-commerce platform may need to request additional information from the customer (e.g., an account name, an order number, etc.) in order to identify the customer record and retrieve it.

BRIEF DESCRIPTION OF THE FIGURES

[0005] In order to describe the manner in which the above-recited and other advantages and features of the disclosure can be obtained, a more particular description of the principles briefly described above will be rendered by reference to specific embodiments that are illustrated in the appended drawings. Understanding that these drawings depict only embodiments of the disclosure and are not therefore to be considered to be limiting of its scope, the principles herein are described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0006] FIG. 1 is a conceptual block diagram illustrating an example network environment that includes a context service system, in accordance with various embodiments of the subject technology;

[0007] FIG. 2 is a conceptual block diagram illustrating an example context service system, in accordance with various embodiments of the subject technology;

[0008] FIG. 3 is a conceptual block diagram illustrating an example context service data store, in accordance with various embodiments of the subject technology;

[0009] FIG. 4 shows an example process for providing a subscriber system with a customer record from a data partition associated with the subscriber system, in accordance with various embodiments of the subject technology;

[0010] FIG. 5 shows an example process for providing a subscriber system with a customer record from a shared data partition, in accordance with various embodiments of the subject technology;

[0011] FIG. 6 is a conceptual block diagram illustrating an example scenario, in accordance with various embodiments of the subject technology; and

[0012] FIGS. 7A and 7B illustrate examples of systems in accordance with some embodiments.

DESCRIPTION OF EXAMPLE EMBODIMENTS

[0013] The detailed description set forth below is intended as a description of various configurations of embodiments and is not intended to represent the only configurations in which the subject matter of this disclosure can be practiced. The appended drawings are incorporated herein and constitute a part of the detailed description. The detailed description includes specific details for the purpose of providing a more thorough understanding of the subject matter of this disclosure. However, it will be clear and apparent that the subject matter of this disclosure is not limited to the specific details set forth herein and may be practiced without these details. In some instances, structures and components are shown in block diagram form in order to avoid obscuring the concepts of the subject matter of this disclosure.

[0014] Overview

[0015] The disclosed technology relates to a context service system configured to receive, from a subscriber, a shared customer lookup request that includes a first customer data identifier and identify, in a shared data partition, a second customer data identifier associated with the first customer data identifier. The context service system is further configured to determine that the second customer data identifier is associated with customer information in a subscriber data partition and transmit, to the subscriber system, the customer information from the subscriber data partition.

EXAMPLE EMBODIMENTS

[0016] As described above, fragmentation of customer information and/or the incompleteness of customer records may lead to various detrimental results including, for example, a degraded level of service, an increase the time and effort needed to provide services, or an increase in costs to provide services. In order to make their customer records more complete and robust, organizations may be tempted to share and exchange information about customers through partnerships or through third-party data brokers. However, the information collected may also include personally iden-

tifiable information (PII) or other private information that cannot (legally) be shared without the consent of the customer. In other cases, an organization may not wish to share collected information for various other business reasons. For example, sharing collected information may anger customers or the public. Furthermore, the information collected may also be a valuable resource that may serve as a competitive advantage over other businesses or organizations.

[0017] Aspects of the subject technology relate to a context service system configured to leverage the customer information shared across multiple subscribers and allow subscribers to share associations between different elements of customer information without exposing the customer information. As will be discussed in further detail, the context service system generates a shared partition of data that contains associations of different elements of customer information provided by a subscriber group. The context service system facilitates the querying of the shared partition of data by a subscriber to determine whether a first element of customer information provided by the subscriber is associated with another element of customer information that is already known to the subscriber.

[0018] By querying the context service system, the subscriber can determine whether an element of customer information is associated with a customer record already known to the subscriber. As an illustrative example, a call center system may receive a phone call and identify the phone number (e.g., using caller ID) associated with the phone call. However, the phone number may be previously unknown to the call center system. Accordingly, the call center system may not know any context for the call and may not be able to provide certain services or service levels because of the lack of context. In response to receiving the call, the call center system may transmit a customer lookup request that includes the unknown phone number to the context service system.

[0019] The context service system may determine whether the phone number is associated with any other items of customer information in a shared partition of data generated based on information from a number of subscribers. If an association is found with another element of customer information (e.g., an address or email address), the context service system may determine if the element of customer information is already known to the call center system. If the element of customer information is already known to the call center system, the context service system may inform the call center system that the unknown phone number is actually associated with the element of customer information (e.g., an address or email address) or provide the call center system with the call center system's record including the element of customer information.

[0020] In this way, the call center system may determine that the previously unknown phone number is actually associated with an existing customer record (e.g., the customer record with the address or email address), update the call center system's customer record, and/or provide additional services or benefits enabled by the new contextual information (e.g., the customer record) provided by the context service system. Furthermore, although the context service system leverages information shared across a number of subscribers in the shared partition of data, the context service system does not expose any elements of customer information that is unknown to a subscriber. For example, in

the example above, the call center system already knows of the phone number and the customer record for the caller, the context service system merely provided the association between the two elements of customer information. Accordingly, the context service system prevents and/or limits the dissemination of customer information, including personally identifiable information (PII) or other sensitive information.

[0021] For illustrative purposes, various embodiments described herein may refer to customer information or elements of customer information. However, these embodiments and others may also readily apply to other information stored by subscribers and not necessarily a "customer." For example, other types of information that is applicable include account information, user information, or profile information. Furthermore, the information may be associated with an individual, a company, an organization, or other entity.

[0022] Various embodiments of the disclosure are discussed in detail below. While specific implementations are discussed, it should be understood that this is done for illustrative purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without departing from the spirit and scope of the disclosure.

[0023] FIG. 1 is a conceptual block diagram illustrating an example network environment 100 that includes a context service system 160, in accordance with various embodiments of the subject technology. The network environment 100 described with respect to FIG. 1 includes the context service system 160, a subscriber system 120, a customer device 140, a third-party system 150, and other subscriber systems 130. However, in other embodiments, other network environment configurations may also be used.

[0024] The subscriber system 120 and other subscriber systems 130 are configured to communicate with the context service system 160 via a network 110. The network 110 may include, for example, any one or more of a cellular network, a satellite network, a personal area network (PAN), a local area network (LAN), a wide area network (WAN), a broadband network (BBN), the Internet, and the like. The network environment 100 can be a public network, a private network, or a combination thereof. The network environment 100 may be implemented using any number of communications links associated with one or more service providers, including one or more wired communication links, one or more wireless communication links, or any combination thereof. Additionally, the network 110 can be configured to support the transmission of data formatted using any number of protocols.

[0025] The subscriber system 120 and/or subscriber systems 130 may be associated with a financial services entity, a medical industry entity, a business entity, an e-commerce entity, or any other entity that deals with customer information, account information, user information, or profile information. The subscriber system 120 and/or subscriber systems 130 may also be associated with a particular function or purpose such as a call center, a contact center, a sales organization or group, a customer service center, a help desk, or other group that may deal with customer information, account information, user information, or profile information. The subscriber system 120 and/or subscriber systems 130 may each manage a separate database of customer information.

[0026] According to some embodiments, the subscriber system 120 may obtain an element of customer information such as a phone number, email address, mailing address, user name, or other identifier. The element of customer information may be generated by the subscriber system or received via an external party such as customer device 140 or third-party system 150. For example, the third-party system 150 may be configured to perform a task with the subscriber system 120 via an application programming interface (API) provided by the subscriber system 120 and the third-party system 150 may provide the subscriber system 120 with the element of customer information for a user, account, or other entity during the course of performing that task.

[0027] In another example, the customer device 140 may communicate with the subscriber system 120 and provide, either directly or indirectly (e.g., via the communication protocols used), the element of customer information. The communications may occur via, for example, a cellular network, a landline telephone network, an short message service network, a chat or messaging service, or any number of other communication technologies.

[0028] The context service system 160 enables the subscriber system 120 to determine whether the element of customer information is associated with other information that is already in the subscriber system's database of customer information by leveraging customer information associated with the other subscriber systems 130 without revealing any of the customer information associated with the other subscriber systems 130.

[0029] FIG. 2 is a conceptual block diagram illustrating an example context service system 200, in accordance with various embodiments of the subject technology. The context service system 200 described with respect to FIG. 2 includes a subscriber interface 210, a context engine 220, and a context service data store 230. The context service data store 230 is shown to include a shared data partition 240 and one or more subscriber data partitions. However, in other embodiments, configurations may also be used including additional components, fewer components, or alternative components.

[0030] The subscriber interface 210 is configured to communicate with one or more subscribers. For example, the subscriber interface 210 may receive customer lookup requests that include an element of customer information such as a customer data identifier and provide subscribers with customer records or other customer data identifiers associated with the element of customer information in the customer lookup request.

[0031] As will be discussed further below, the context engine 220 is configured to identify other customer information that is associated with the element of customer information in the customer lookup request and determine whether the subscriber that submitted the customer lookup request has access to the associated customer information.

[0032] The context service data store 230 may store customer information associated with the one or more subscribers of the context service system 200. According to some embodiments, the context service data store 230 may include a shared data partition 240 and a subscriber data partition for each subscriber of the context data store. For example, in FIG. 2, subscriber data partition 250 may be associated with the subscriber system 120 of FIG. 1 and the

other subscriber data partitions 260 may be associated with the other subscriber systems 130 of FIG. 1.

[0033] The subscriber data partitions 250 and 260 include sets of customer information (e.g., customer records) known by their corresponding subscriber. For example, subscriber data partition 250 may include sets of customer information or customer records that the subscriber system 120 of FIG. 1 is aware of. Subscriber data partition 250 may be provided by the subscriber and updated over time. In some embodiments, the subscriber data partition 250 is synchronized with a data store managed by the subscriber system (e.g., subscriber system 120 of FIG. 1). Although FIG. 2 is shown with the context service data store 230 including the subscriber data partitions 250 and 260, in other embodiments, one or more subscriber data partitions may be stored by a subscriber system and accessed by the context service system 200 via network communications.

[0034] The shared data partition 240 includes sets or records of related customer information that is generated based on the customer information in the subscriber data partitions 250 and 260 or customer information provided by subscriber systems. The information in the shared data partition 240 may include information that each subscriber specifically permits to be used by the context service system 200 to provide associations to other systems.

[0035] According to some embodiments, an end-to-end encryption scheme may be used to protect the customer information stored by the context service data store 230. In other words, the subscriber systems may encrypt some or all data transmitted and/or stored by the context service data store 230. For example, a subscriber system may encrypt the entire customer record for all customer records to be stored by the context service system 200 or a portion of the customer records (e.g., any personally identifiable information (PII), private information, or otherwise sensitive information). The encryption and decryption scheme and/or various parameters used for encryption and decryption may be known only to the subscriber system. In this way, a subscriber may be assured with an additional layer of protection that prevents other subscribers or even the context service system 200 from unpermitted access to the subscriber's encrypted customer information.

[0036] FIG. 3 is a conceptual block diagram illustrating an example context service data store 330, in accordance with various embodiments of the subject technology. For illustrative purposes, the context service data store 330 includes a shared data partition 340, an ACME data partition for a subscriber named "ACME," and a Z Corp data partition for a subscriber named "Z Corp." Each data partition may include a number of sets of associated customer information or customer records. To conserve space, each of the data partitions 340, 350, and 360 in FIG. 3 are shown with one customer record or set of associated customer information each. For example, ACME data partition 350 includes customer record 355, shared data partition 340 includes customer record 345, and Z Corp data partition 360 includes customer record 365.

[0037] The context service data store 330 is implemented using an encryption scheme where all or portions of the customer information is encrypted using a cryptographic function. The encryption and decryption scheme and/or various parameters used for encryption and decryption may be known only to the subscriber system or entities that the subscriber system entrusts in order to provide an additional

layer of protection that prevents other subscribers or even the context service system **200** from unpermitted access to the subscriber's encrypted customer information. For example, in customer record **355** in the ACME data partition **350**, the customer information **358** in the customer record is encrypted. Similarly, in customer record **365** in the Z Corp data partition **360**, the customer information **368** in the customer record is encrypted.

[0038] As will be explained in further detail below, in order to allow the context service system to determine whether an element of customer information is associated with the encrypted information that in the subscriber data partition, the subscriber system may hash each element of customer information using a subscriber specific cryptographic salt value and a hash function. As an added advantage, the hash value for the element of customer information enables the context service system to determine whether the element of customer information is associated with a customer record in the subscriber data partition without revealing the element of customer information to the context service system or any other party.

[0039] For example, for the email address "fsmith@email.com" in customer record **355**, the ACME subscriber system may use an "ACME" subscriber salt value to generate a customer data identifier (e.g., a hash value) for the email address. The ACME subscriber system may then transmit the customer data identifier to the context service system for storage in the ACME data partition **350**. Similarly, in customer record **365**, the Z Corp subscriber system may generate a customer data identifier (e.g., a hash value) for the email address, mobile phone number, or home phone number based on a "ZCORP" subscriber salt value. The Z Corp subscriber system may then transmit the customer data identifiers to the context service system for storage in the Z Corp data partition **360**.

[0040] Hash values generated based on a shared cryptographic salt value further enable the context service system to determine whether an element of customer information is associated with the encrypted information that in the subscriber data partition based on information in the shared data partition.

[0041] For example, for the email address "fsmith@email.com" in customer record **355**, the ACME subscriber system may use a shared salt value, which is represented by the "SHARED" string in the hash function shown in customer record **355**, to generate a customer data identifier (e.g., a hash value) for the email address. The ACME subscriber system may then transmit the customer data identifier to the context service system for storage in the ACME data partition **350**. Similarly, in customer record **365**, the Z Corp subscriber system may generate a customer data identifier (e.g., a hash value) for the email address, mobile phone number, or home phone number based on the shared salt value. The Z Corp subscriber system may then transmit the customer data identifiers to the context service system for storage in the Z Corp data partition **360**.

[0042] The context service system may identify the hashes of customer information in customer records across the various subscriber data partitions that have been hashed using the shared salt value and aggregate them into the shared data partition **340**. For example, customer record **345** includes all of the hash values of customer information that have been hashed using the shared salt value. Since the context service system does not store the actual customer

information in the context service data store **330** but instead stores customer data identifiers (e.g., the hash values) associated with the customer information, the customer information is prevented from being shared or accessed either intentionally or unintentionally. Furthermore, in some embodiments, the context service system may not know any of the salt values (e.g., the shared salt value or the subscriber salt values) used to encrypt the customer information. Accordingly, even the context service system may not access the customer information.

[0043] According to some embodiments, a key management server may be configured to disseminate salt values to the subscriber systems. For example, the key management server may provide a subscriber system with a shared salt value that is also provided to all other subscriber systems as well as a subscriber salt value that is only provided to the subscriber system. Accordingly, a subscriber system may use the various salt values to determine which elements of subscriber's customer information are to be used by the context service system to assist other subscribers.

[0044] Although FIG. 2 and FIG. 3 show only a single shared data partition **340** for illustrative purposes, in other embodiments, multiple shared data partitions may also be implemented. Multiple shared data partitions may allow for subscribers to have different sharing levels with different groups. For example, shared data partitions may be established for a group of business partners, a group of competitors, a group of customers, and/or a general public group and each group may be assigned a different shared salt value for the group. Each subscriber system belonging to the group may obtain a shared salt value for the group of subscribers. Accordingly, a subscriber system may receive multiple shared salt values for a multitude of groups and use the shared salt values to determine which elements of customer information to allow the context service system to use to assist other subscribers.

[0045] FIG. 4 shows an example process **400** for providing a subscriber system with a customer record from a data partition associated with the subscriber system, in accordance with various embodiments of the subject technology. It should be understood that, for any process discussed herein, there can be additional, fewer, or alternative steps performed in similar or alternative orders, or in parallel, within the scope of the various embodiments unless otherwise stated. The process **400** can be performed by a system such as, for example, the context service system **200** of FIG. 2 or a similar computing platform.

[0046] A subscriber system may receive an element of customer information during operation of the subscriber system. For example, a call center system may receive a call or text from a phone number. Alternatively, the call center may receive an email or message from a messaging service from an email address or user name that is not in the call center system's database. The subscriber system may attempt to determine if there is any additional contextual information associated with the element of customer information.

[0047] According to some embodiments, the subscriber system may maintain a local database of customer information and determine if the element of customer information is associated with other customer information (e.g., a customer account or record). However, in other embodiments, the subscriber system may not maintain a local database of customer information or may rely on the context service

system to store the subscriber system's customer information. Accordingly, the subscriber system may generate a subscriber customer lookup request that includes a customer data identifier associated with the element of customer information and transmit the request to the context service system. In some embodiments, the customer data identifier may be the element of customer information itself. However, in embodiments where encryption is used, the customer data identifier may be generated based on the application of a hashing function on the element of customer information using a subscriber salt value for the subscriber system.

[0048] At operation 405, the context service system receives the subscriber customer lookup request and determines whether the customer data identifier in the subscriber customer lookup request is associated with a customer record in the subscriber's data partition at operation 410. If a customer record is associated with the customer data identifier, the customer record may be returned to the subscriber system at operation 415.

[0049] If no customer record is found, at operation 420, the context service system may notify the subscriber system that no customer record was found in the subscriber data partition associated with the subscriber system. If no customer record was found in the subscriber data partition associated with the subscriber system, the subscriber system may wish to query the shared data partition for customer records associated with the element of customer information, which leads to process 500 of FIG. 5. Alternatively, even if a customer record is found in the subscriber data partition, the subscriber system may wish to determine whether other information in the subscriber data partition may also be associated with the element of customer information and the customer record that was found. This can also be accomplished through process 500 of FIG. 5.

[0050] FIG. 5 shows an example process 500 for providing a subscriber system with a customer record from a shared data partition, in accordance with various embodiments of the subject technology. It should be understood that, for any process discussed herein, there can be additional, fewer, or alternative steps performed in similar or alternative orders, or in parallel, within the scope of the various embodiments unless otherwise stated. The process 500 can be performed by a system such as, for example, the context service system 200 of FIG. 2 or a similar computing platform.

[0051] To query the shared data partition for customer records associated with the element of customer information, the subscriber system may generate a shared customer lookup request that includes a first customer data identifier associated with the element of customer information and transmit the request to the context service system. In some embodiments, the first customer data identifier may be the element of customer information itself. In other embodiments, the first customer data identifier may be generated by applying a hash function to the element of customer information using a shared salt value.

[0052] At operation 505, the context service system may receive the shared customer lookup request including the first customer data identifier. The context service system may query the shared data partition to determine if there is another customer data identifier (e.g., another hash value associated with different element of customer data) that is associated with the first customer data identifier included in the shared customer lookup request.

[0053] At operation 510, the context service system may identify a second customer data identifier in the shared data partition that is associated with the first customer data identifier. The context service system may then determine whether the second customer data identifier is associated with a customer record in a subscriber data partition at operation 515. If there is a customer record in the subscriber data partition, the context service system may provide the customer record to the subscriber system at operation 520. If there is no customer record in the subscriber data partition, the context service system notifies the subscriber system that no record was found at operation 525.

[0054] FIG. 6 is a conceptual block diagram illustrating an example scenario 600, in accordance with various embodiments of the subject technology. The scenario 600 includes an ACME subscriber system 610 and a context service data store 630 and may help to illustrate the processes 400 and 500 of FIGS. 4 and 5. The other components of the context service system associated with the context service data store 630 are not shown in FIG. 6 for illustrative purposes. The context service data store 630 is shown containing a subscriber partition for an "ACME" subscriber 650 and a shared data partition 640. However, the context service data store 630 may contain many other shared data partitions and/or subscriber partitions. For illustrative purposes, each data partition is shown containing a record for customer "Frank Smith." However, the data partitions may contain many more customer records or sets of associated data.

[0055] According to some embodiments, the context service system may enable a subscriber to retrieve a customer record in the subscriber's data partition that is associated with a customer data identifier (e.g., an element of customer information or a hash value for the element of customer information) provided in a customer lookup request. In some cases, the customer record may be located based on information in the subscriber's data partition. However, in other cases, the context service system may leverage the associations in the shared data partition to identify the customer record in the subscriber's data partition.

[0056] In the scenario 600 of FIG. 6, the ACME subscriber system 610 may come across an element of customer information. In particular, the ACME subscriber system 610 may receive a call 672 from a phone number 555-123-4567 that is unknown to the ACME subscriber system. The phone number of 555-123-4567 may be the element of customer information in this scenario 600, but in other scenarios, user names, email addresses, social media account identifiers, or any other identifiers may be used as elements of customer information and the subscriber system 610 may come across the element of customer information by other means.

[0057] To determine whether the ACME data partition 650 contains information associated with the received phone number, the ACME subscriber system 610 may generate a hash value for the received phone number using an ACME subscriber salt value, which is represented as "HASH(ACME, 555-123-4567)" in FIG. 6. The ACME subscriber system 610 inserts the hash value in a subscriber customer lookup request 674 and transmits the subscriber customer lookup request 674 to the context service system.

[0058] The context service system receives the subscriber customer lookup request 674 and determines whether the customer data identifier "HASH(ACME, 555-123-4567)" in the subscriber customer lookup request is associated with a customer record in the subscriber's data partition 650. In the

scenario 600 of FIG. 6, no customer record is found. Accordingly, the context service system may generate and transmit a notification 676 to the ACME subscriber system 610 informing the subscriber system that no customer record was found in the subscriber's data partition 650.

[0059] The ACME subscriber system 610 receives the notification 676 and may wish to leverage the information in the shared data partition 640 to determine whether the ACME data partition 650 contains information associated with the received phone number. Accordingly, the ACME subscriber system 610 generates a hash value for the received phone number using a shared subscriber salt value, which is represented as "HASH(SHARED, 555-123-4567)" in FIG. 6. The ACME subscriber system 610 inserts the hash value in a shared customer lookup request 680 and transmits the shared customer lookup request 680 to the context service system.

[0060] The context service system receives the shared customer lookup request 680 and determines whether the customer data identifier "HASH(SHARED, 555-123-4567)" in the shared customer lookup request is associated with a customer record in the shared data partition 640. In the scenario 600 of FIG. 6, customer record 645 is found. More specifically, customer record 645 also includes a matching customer data identifier with the same "HASH(SHARED, 555-123-4567)" value.

[0061] Based on the found customer record 645, the context service system identifies two additional customer data identifiers ("HASH(SHARED, fsmith@email.com)" and "HASH(SHARED, 321-555-8888)") that are associated with the customer data identifier ("HASH(SHARED, 555-123-4567)") included in the shared customer lookup request.

[0062] The context service system may then query 682 the ACME data partition 650 to determine whether either of the two identified customer data identifiers ("HASH(SHARED, fsmith@email.com)" and "HASH(SHARED, 321-555-8888)") are associated with a customer record in the ACME data partition 650. In the scenario 600 of FIG. 6, no corresponding customer record is found for "HASH(SHARED, 321-555-8888)." However, customer record 655 is found for "HASH(SHARED, fsmith@email.com)." Accordingly, the context service system provides the customer record 655 or the encrypted customer information 658 in the customer record 655 to the ACME subscriber system 610. This is represented by transmission 690 in FIG. 6.

[0063] After receiving the transmission 690, the ACME subscriber system 610 may decrypt the customer information 658 in the customer record 655 to reveal that the customer information elements of "Name: Fred Smith" and "Email: fsmith@email.com" are associated with the call 672 from the phone number 555-123-4567. In many cases, the additional information that is revealed may enable the subscriber system to provide improved services and/or additional services.

[0064] As the data sets in the context service data store increases in size and the number of subscriber data partitions grows there is an increased likelihood that multiple customer records may be identified in the subscriber data partition. According to various embodiments, the context service system may also be configured to generate scores (e.g., confidence scores) for the identified customer records and/or rank the identified customer records. The context service system may provide the subscriber system with a top ranked customer record or a set of top ranked customer records.

Additionally, the context service system may provide one or more or all of the customer records to the subscriber system along with their calculated scores to allow the subscriber system to select.

[0065] According to various embodiments, the confidence score for an identified customer record may be based on the number of other subscriber systems that have identified the association between the customer record and the element of customer information and/or how recent the association between the customer record and the element of customer information was made. For example, if a large number of subscribers have identified an association between the customer record and the element of customer information, it is more likely to be an accurate identification. Accordingly, the confidence score should be higher than for a customer with fewer subscribers having identified the association.

[0066] Furthermore, associations between customer records and the element of customer information may be timestamp according to when they were created or updated. An association between a customer record and the element of customer information with a recent timestamp may be more accurate than an association with an older timestamp because the association with the recent timestamp may reflect updated information that has not reached the data partition for another subscriber. For example, a customer record may include a residential address for a user. This information may be reflected in a customer record with a particular timestamp. However, the user may have moved to a new address. This new information may be reflected in a customer record with a more recent timestamp.

[0067] FIG. 7A and FIG. 7B illustrate systems in accordance with various embodiments. The more appropriate system will be apparent to those of ordinary skill in the art when practicing the various embodiments. Persons of ordinary skill in the art will also readily appreciate that other systems are possible.

[0068] FIG. 7A illustrates an example architecture for a bus computing system 700 wherein the components of the system are in electrical communication with each other using a bus 705. The computing system 700 can include a processing unit (CPU or processor) 710 and a system bus 705 that may couple various system components including the system memory 715, such as read only memory (ROM) in a storage device 720 and random access memory (RAM) 725, to the processor 710. The computing system 700 can include a cache 712 of high-speed memory connected directly with, in close proximity to, or integrated as part of the processor 710. The computing system 700 can copy data from the memory 715 and/or the storage device 730 to the cache 712 for quick access by the processor 710. In this way, the cache 712 can provide a performance boost that avoids processor delays while waiting for data. These and other modules can control or be configured to control the processor 710 to perform various actions. Other system memory 715 may be available for use as well. The memory 715 can include multiple different types of memory with different performance characteristics. The processor 710 can include any general purpose processor and a hardware module or software module, such as module 1 732, module 2 734, and module 3 736 stored in storage device 730, configured to control the processor 710 as well as a special-purpose processor where software instructions are incorporated into the actual processor design. The processor 710 may essentially be a completely self-contained computing system,

containing multiple cores or processors, a bus, memory controller, cache, etc. A multi-core processor may be symmetric or asymmetric.

[0069] To enable user interaction with the computing system **700**, an input device **745** can represent any number of input mechanisms, such as a microphone for speech, a touch-protected screen for gesture or graphical input, keyboard, mouse, motion input, speech and so forth. An output device **735** can also be one or more of a number of output mechanisms known to those of skill in the art. In some instances, multimodal systems can enable a user to provide multiple types of input to communicate with the computing system **700**. The communications interface **740** can govern and manage the user input and system output. There may be no restriction on operating on any particular hardware arrangement and therefore the basic features here may easily be substituted for improved hardware or firmware arrangements as they are developed.

[0070] Storage device **730** can be a non-volatile memory and can be a hard disk or other types of computer readable media which can store data that are accessible by a computer, such as magnetic cassettes, flash memory cards, solid state memory devices, digital versatile disks, cartridges, random access memories (RAMs) **725**, read only memory (ROM) **720**, and hybrids thereof.

[0071] The storage device **730** can include software modules **732**, **734**, **736** for controlling the processor **710**. Other hardware or software modules are contemplated. The storage device **730** can be connected to the system bus **705**. In one aspect, a hardware module that performs a particular function can include the software component stored in a computer-readable medium in connection with the necessary hardware components, such as the processor **710**, bus **705**, output device **735**, and so forth, to carry out the function.

[0072] FIG. 7B illustrates an example architecture for a chipset computing system **750** that can be used in accordance with an embodiment. The computing system **750** can include a processor **755**, representative of any number of physically and/or logically distinct resources capable of executing software, firmware, and hardware configured to perform identified computations. The processor **755** can communicate with a chipset **760** that can control input to and output from the processor **755**. In this example, the chipset **760** can output information to an output device **765**, such as a display, and can read and write information to storage device **770**, which can include magnetic media, and solid state media, for example. The chipset **760** can also read data from and write data to RAM **775**. A bridge **780** for interfacing with a variety of user interface components **785** can be provided for interfacing with the chipset **760**. The user interface components **785** can include a keyboard, a microphone, touch detection and processing circuitry, a pointing device, such as a mouse, and so on. Inputs to the computing system **750** can come from any of a variety of sources, machine generated and/or human generated.

[0073] The chipset **760** can also interface with one or more communication interfaces **790** that can have different physical interfaces. The communication interfaces **790** can include interfaces for wired and wireless LANs, for broadband wireless networks, as well as personal area networks. Some applications of the methods for generating, displaying, and using the GUI disclosed herein can include receiving ordered datasets over the physical interface or be generated by the machine itself by processor **755** analyzing data stored

in the storage device **770** or the RAM **775**. Further, the computing system **700** can receive inputs from a user via the user interface components **785** and execute appropriate functions, such as browsing functions by interpreting these inputs using the processor **755**.

[0074] It will be appreciated that computing systems **700** and **750** can have more than one processor **710** and **755**, respectively, or be part of a group or cluster of computing devices networked together to provide greater processing capability.

[0075] For clarity of explanation, in some instances the various embodiments may be presented as including individual functional blocks including functional blocks comprising devices, device components, steps or routines in a method embodied in software, or combinations of hardware and software.

[0076] In some embodiments the computer-readable storage devices, mediums, and memories can include a cable or wireless signal containing a bit stream and the like. However, when mentioned, non-transitory computer-readable storage media expressly exclude media such as energy, carrier signals, electromagnetic waves, and signals per se.

[0077] Methods according to the above-described examples can be implemented using computer-executable instructions that are stored or otherwise available from computer readable media. Such instructions can comprise, for example, instructions and data which cause or otherwise configure a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Portions of computer resources used can be accessible over a network. The computer executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, firmware, or source code. Examples of computer-readable media that may be used to store instructions, information used, and/or information created during methods according to described examples include magnetic or optical disks, flash memory, USB devices provided with non-volatile memory, networked storage devices, and so on.

[0078] Devices implementing methods according to these disclosures can comprise hardware, firmware, and/or software, and can take any of a variety of form factors. Examples of such form factors include laptops, smart phones, small form factor personal computers, personal digital assistants, rackmount devices, standalone devices, and so on. Functionality described herein also can be embodied in peripherals or add-in cards. Such functionality can also be implemented on a circuit board among different chips or different processes executing in a single device, by way of further example.

[0079] The instructions, media for conveying such instructions, computing resources for executing them, and other structures for supporting such computing resources are means for providing the functions described in these disclosures.

[0080] Although a variety of examples and other information was used to explain aspects within the scope of the appended claims, no limitation of the claims should be implied based on particular features or arrangements in such examples, as one of ordinary skill would be able to use these examples to derive a wide variety of implementations. Further and although some subject matter may have been described in language specific to examples of structural features and/or method steps, it is to be understood that the

subject matter defined in the appended claims is not necessarily limited to these described features or acts. For example, such functionality can be distributed differently or performed in components other than those identified herein. Rather, the described features and steps are disclosed as examples of components of systems and methods within the scope of the appended claims.

1. A computer-implemented method comprising:
 - receiving, from a subscriber system, a shared customer lookup request that includes a first customer data identifier;
 - identifying, in a shared data partition, a second customer data identifier associated with the first customer data identifier;
 - determining that the second customer data identifier is associated with a customer record in a subscriber data partition; and
 - providing to the subscriber system, in response to the second customer data identifier being associated with the customer record, customer information in the customer record from the subscriber data partition.
2. The computer-implemented method of claim 1, wherein the customer information is encrypted.
3. The computer-implemented method of claim 1, wherein the first customer data identifier is a first hashed value based on a shared salt value and a first element of customer information received by the subscriber system, and wherein the second customer data identifier is a second hashed value based on the shared salt value and a second element of customer information.
4. The computer-implemented method of claim 1, wherein the first customer data identifier is an element of customer information received by the subscriber system.
5. The computer-implemented method of claim 1, further comprising identifying a match between the second customer data identifier and a third customer data identifier associated with the customer record in the subscriber data partition.
6. The computer-implemented method of claim 1, further comprising:
 - receiving, from the subscriber system, a subscriber customer lookup request that includes a fourth customer data identifier;
 - determining whether the fourth customer data identifier is associated with a second customer record in the subscriber data partition; and
 - providing, to the subscriber system, second customer information in the second customer record from the subscriber data partition when the fourth customer data identifier is associated with the second customer record in the subscriber data partition.
7. The computer-implemented method of claim 6, wherein the fourth customer data identifier is a hashed value based on a subscriber salt value and an element of customer information received by the subscriber system.
8. The computer-implemented method of claim 1, further comprising:
 - calculating a confidence score for the customer record based on a timestamp associated with the customer record; and
 - providing the confidence score to the subscriber system.

9. The computer-implemented method of claim 1, wherein the first customer data identifier is based on one of a user name, a phone number, an email address, or an address.

10. The computer-implemented method of claim 1, wherein the subscriber data partition is associated with the subscriber system.

11. A non-transitory computer-readable medium comprising instructions, the instructions, when executed by a computing system, cause the computing system to:

- receive, from a subscriber, a shared customer lookup request that includes a first customer data identifier;
- identify, in a shared data partition, a second customer data identifier associated with the first customer data identifier;
- determine that the second customer data identifier is associated with customer information in a subscriber data partition; and
- transmit, to the subscriber, the customer information from the subscriber data partition.

12. The non-transitory computer-readable medium of claim 11, wherein the first customer data identifier is a first hashed value based on a shared salt value and a first element of customer information received by the subscriber, and wherein the second customer data identifier is a second hashed value based on the shared salt value and a second element of customer information.

13. The non-transitory computer-readable medium of claim 11, wherein the first customer data identifier is an element of customer information received by the subscriber.

14. The non-transitory computer-readable medium of claim 11, wherein the instructions further cause the computing system to identify a match between the second customer data identifier and a third customer data identifier associated with the customer information in the subscriber data partition.

15. A system comprising:
- a processor;
 - a subscriber data partition comprising customer records associated with a subscriber;
 - a shared data partition comprising customer records associated with a set of subscribers; and
 - a non-transitory computer-readable medium storing instructions that, when executed by the system, cause the system to:
 - receive, from the subscriber, a shared customer lookup request that includes a first customer data identifier;
 - identify, in the shared data partition, a second customer data identifier associated with the first customer data identifier;
 - determine that the second customer data identifier is associated with customer information in the subscriber data partition; and
 - transmit, to the subscriber, the customer information from the subscriber data partition.

16. The system of claim 15, wherein the subscriber is a member of the set of subscribers.

17. The system of claim 15, wherein the customer information is encrypted.

18. The system of claim 15, wherein the first customer data identifier is a first hashed value based on a shared salt value and a first element of customer information received by the subscriber, and wherein the second customer data

identifier is a second hashed value based on the shared salt value and a second element of customer information.

19. The system of claim **15**, wherein the first customer data identifier is an element of customer information received by the subscriber.

20. The system of claim **15**, wherein the instructions further cause the system to identify a match between the second customer data identifier and a third customer data identifier associated with the customer information in the subscriber data partition.

* * * * *