



US 20200242708A1

(19) **United States**

(12) **Patent Application Publication**
Schneider

(10) **Pub. No.: US 2020/0242708 A1**

(43) **Pub. Date: Jul. 30, 2020**

(54) **COMPUTER IMPLEMENTED METHODS AND APPARATUS FOR INFORMING A USER OF SOCIAL NETWORK DATA WHEN THE DATA IS RELEVANT TO THE USER**

Publication Classification

(51) **Int. Cl.**
G06Q 50/00 (2012.01)
(52) **U.S. Cl.**
CPC **G06Q 50/01** (2013.01)

(71) Applicant: **salesforce.com, inc.**, San Francisco, CA (US)

(57) **ABSTRACT**

(72) Inventor: **Elena Schneider**, San Francisco, CA (US)

Disclosed are examples of apparatus, methods, and computer readable storage media for informing one or more users of social network data such as group data when the data is relevant to the user(s). For example, when one or more keywords are identified as being associated with both a group of an online social network and the user, a communication with information identifying the group data can be generated and provided to a display device accessible by the user. In another example, social network data is associated with an entity of the online social network, and the entity is not being followed by the user. Social network data identified as related to one or more keywords can be provided to a display device in a suitable presentation accessible by the user.

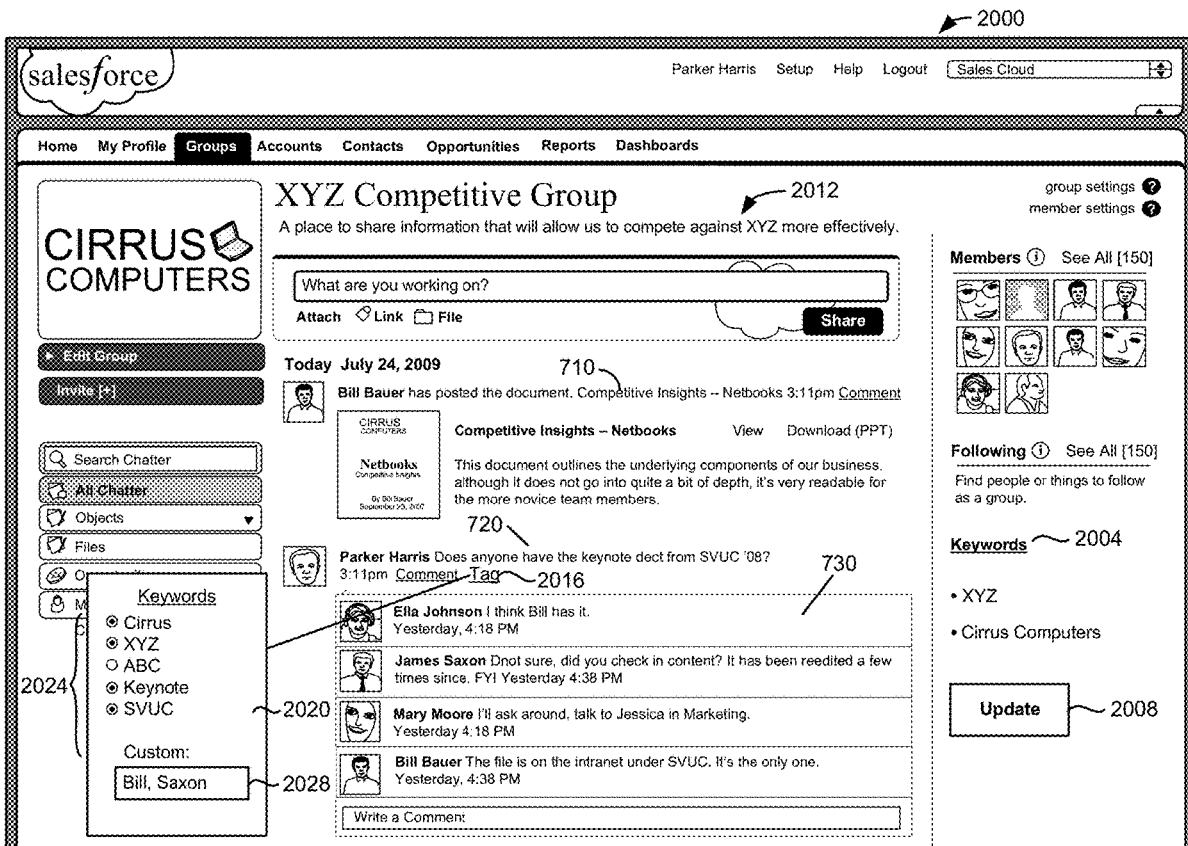
(21) Appl. No.: **16/850,127**

(22) Filed: **Apr. 16, 2020**

Related U.S. Application Data

(63) Continuation of application No. 15/640,000, filed on Jun. 30, 2017, now Pat. No. 10,664,928, which is a continuation of application No. 13/603,910, filed on Sep. 5, 2012, now Pat. No. 9,727,924.

(60) Provisional application No. 61/545,260, filed on Oct. 10, 2011.



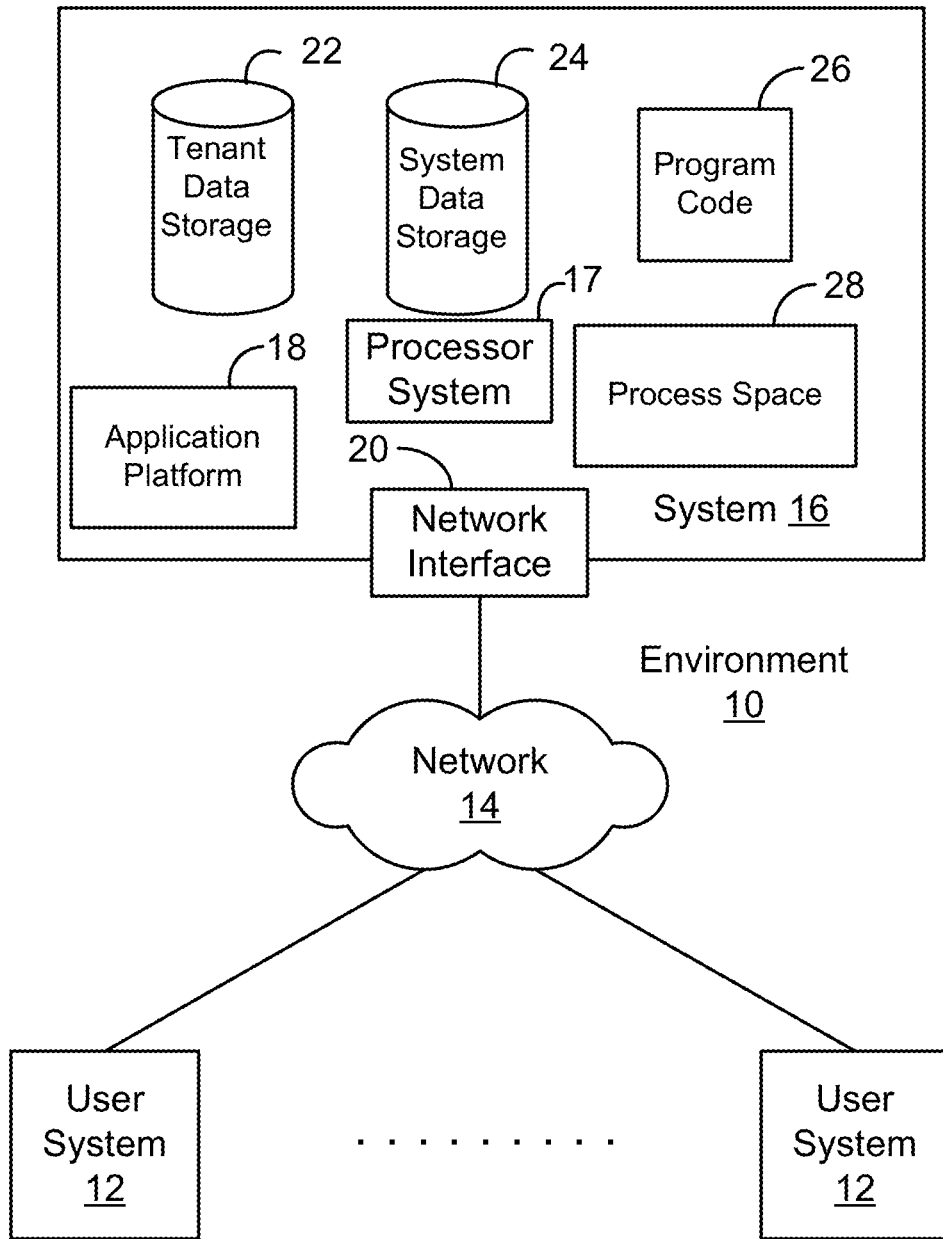


Figure 1A

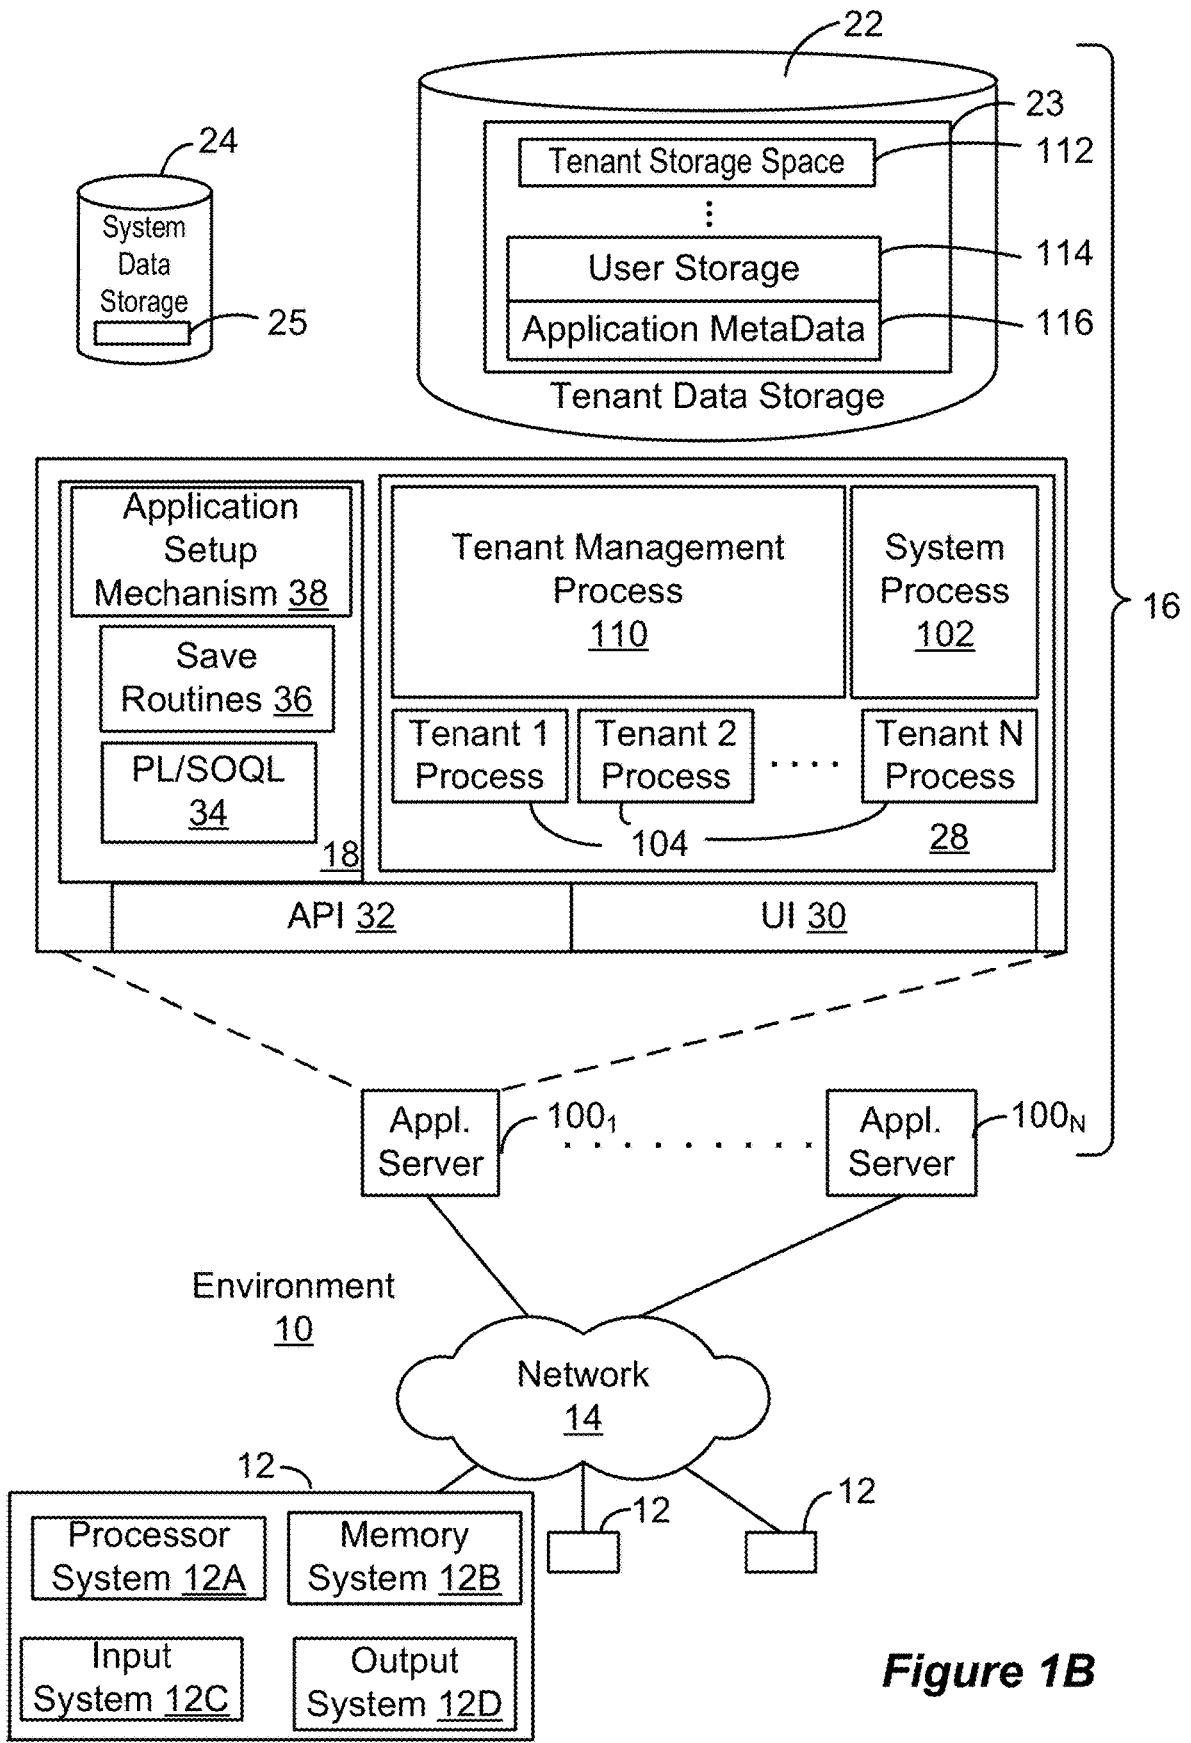


Figure 1B




Figure 2A

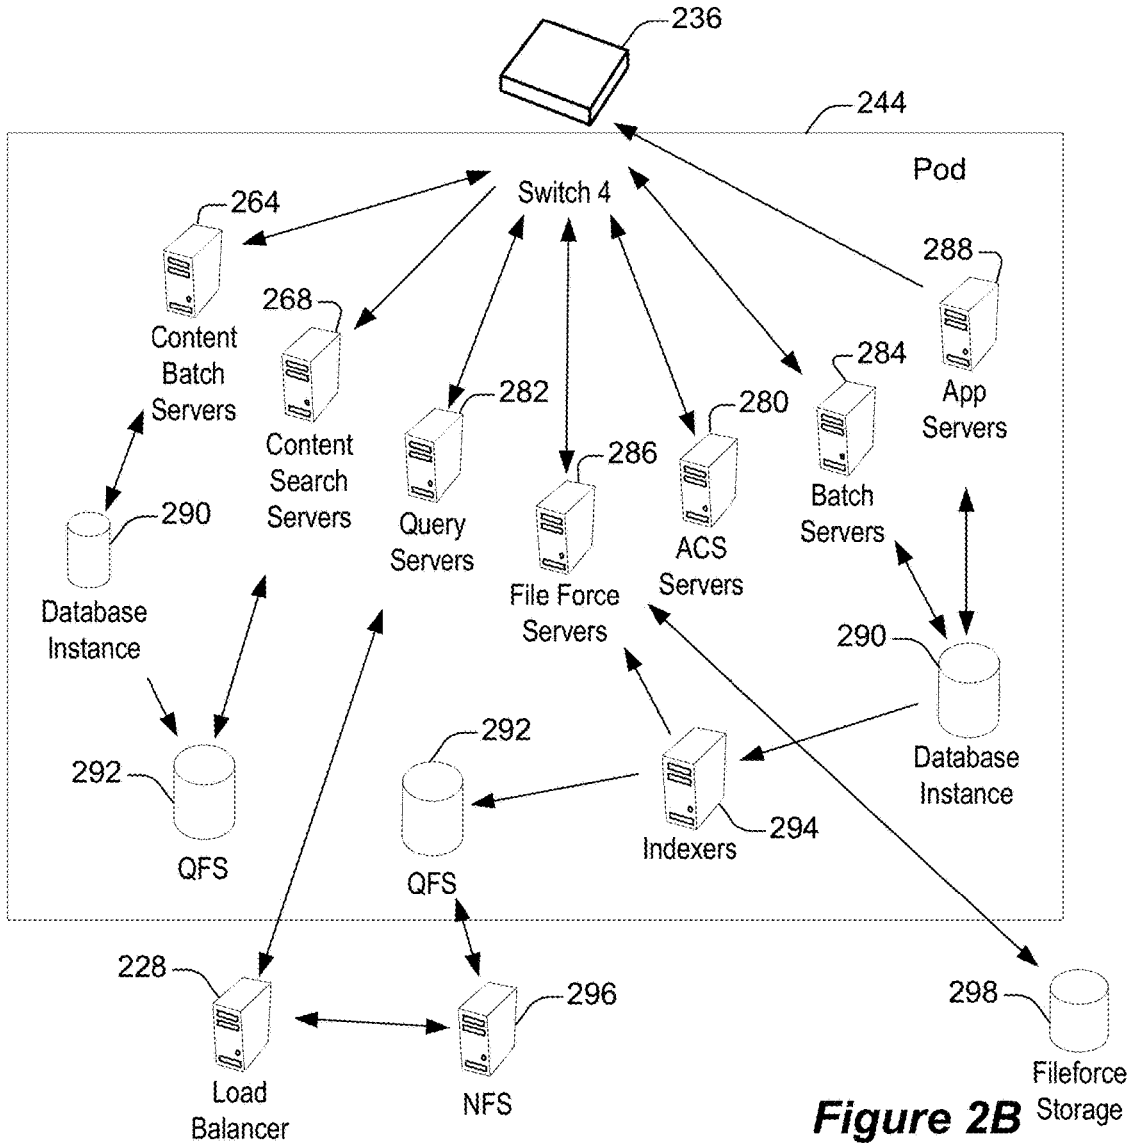


Figure 2B

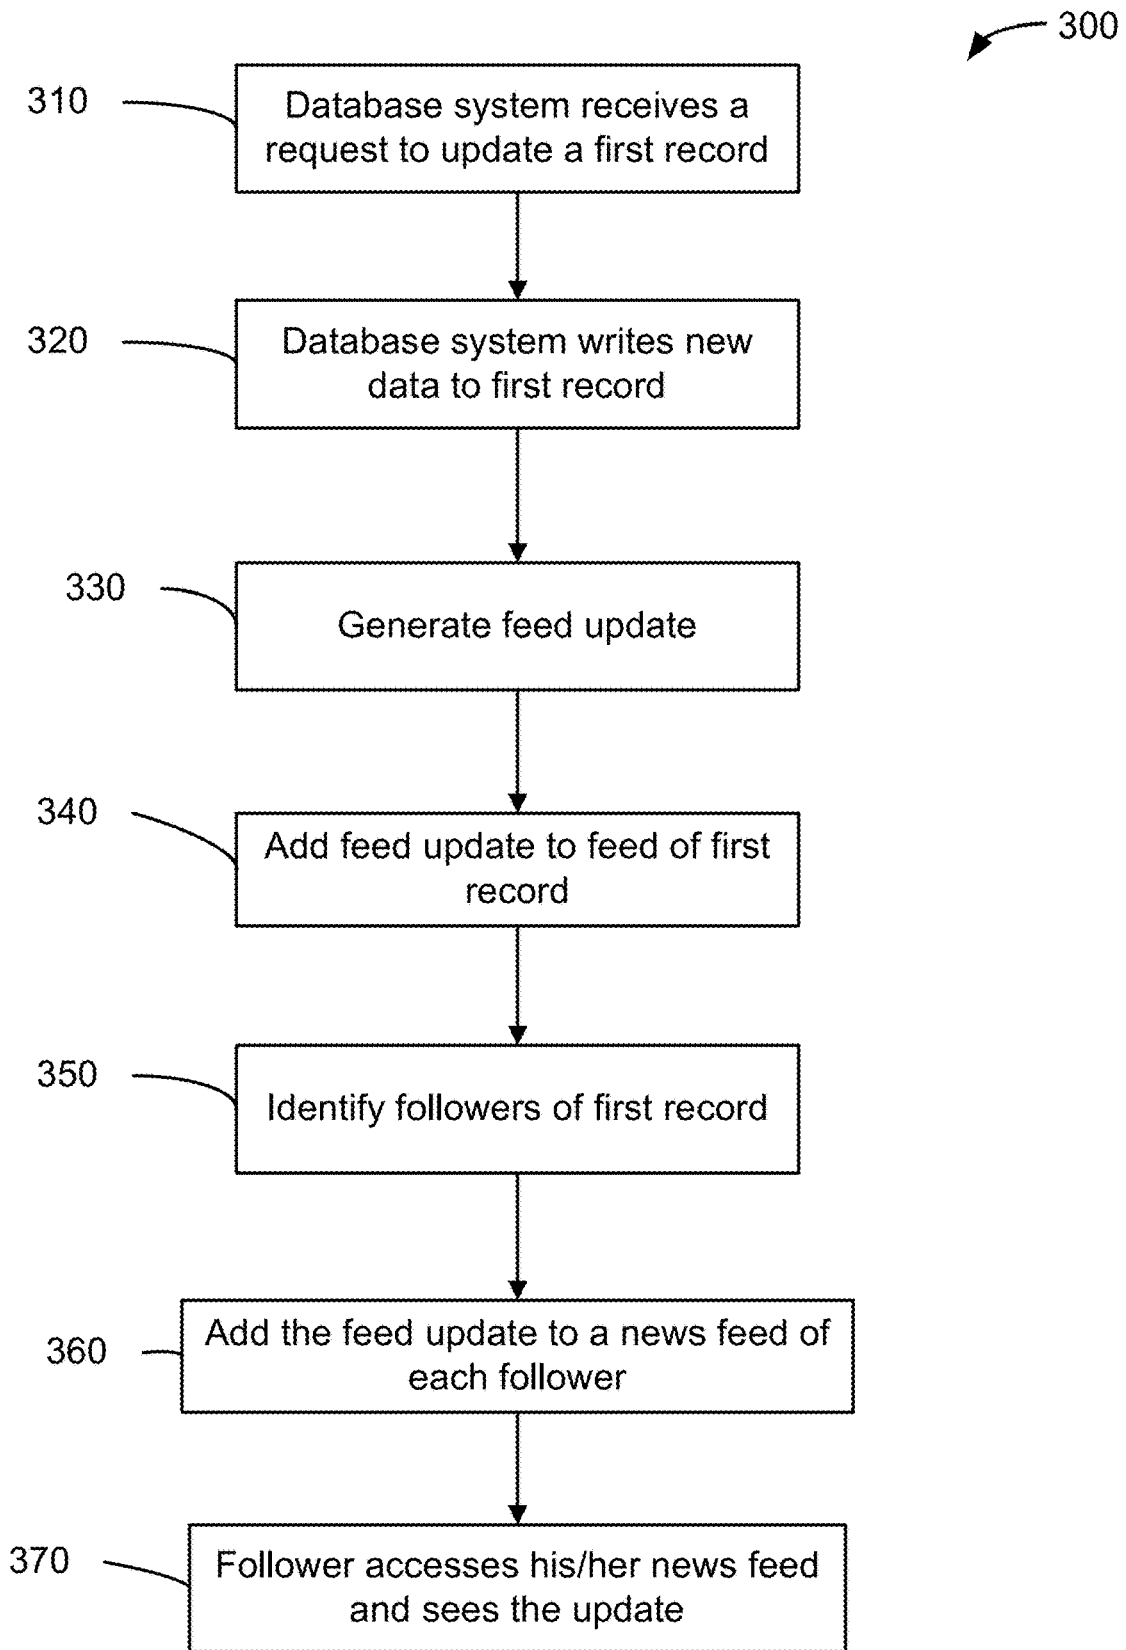


Figure 3




Figure 4

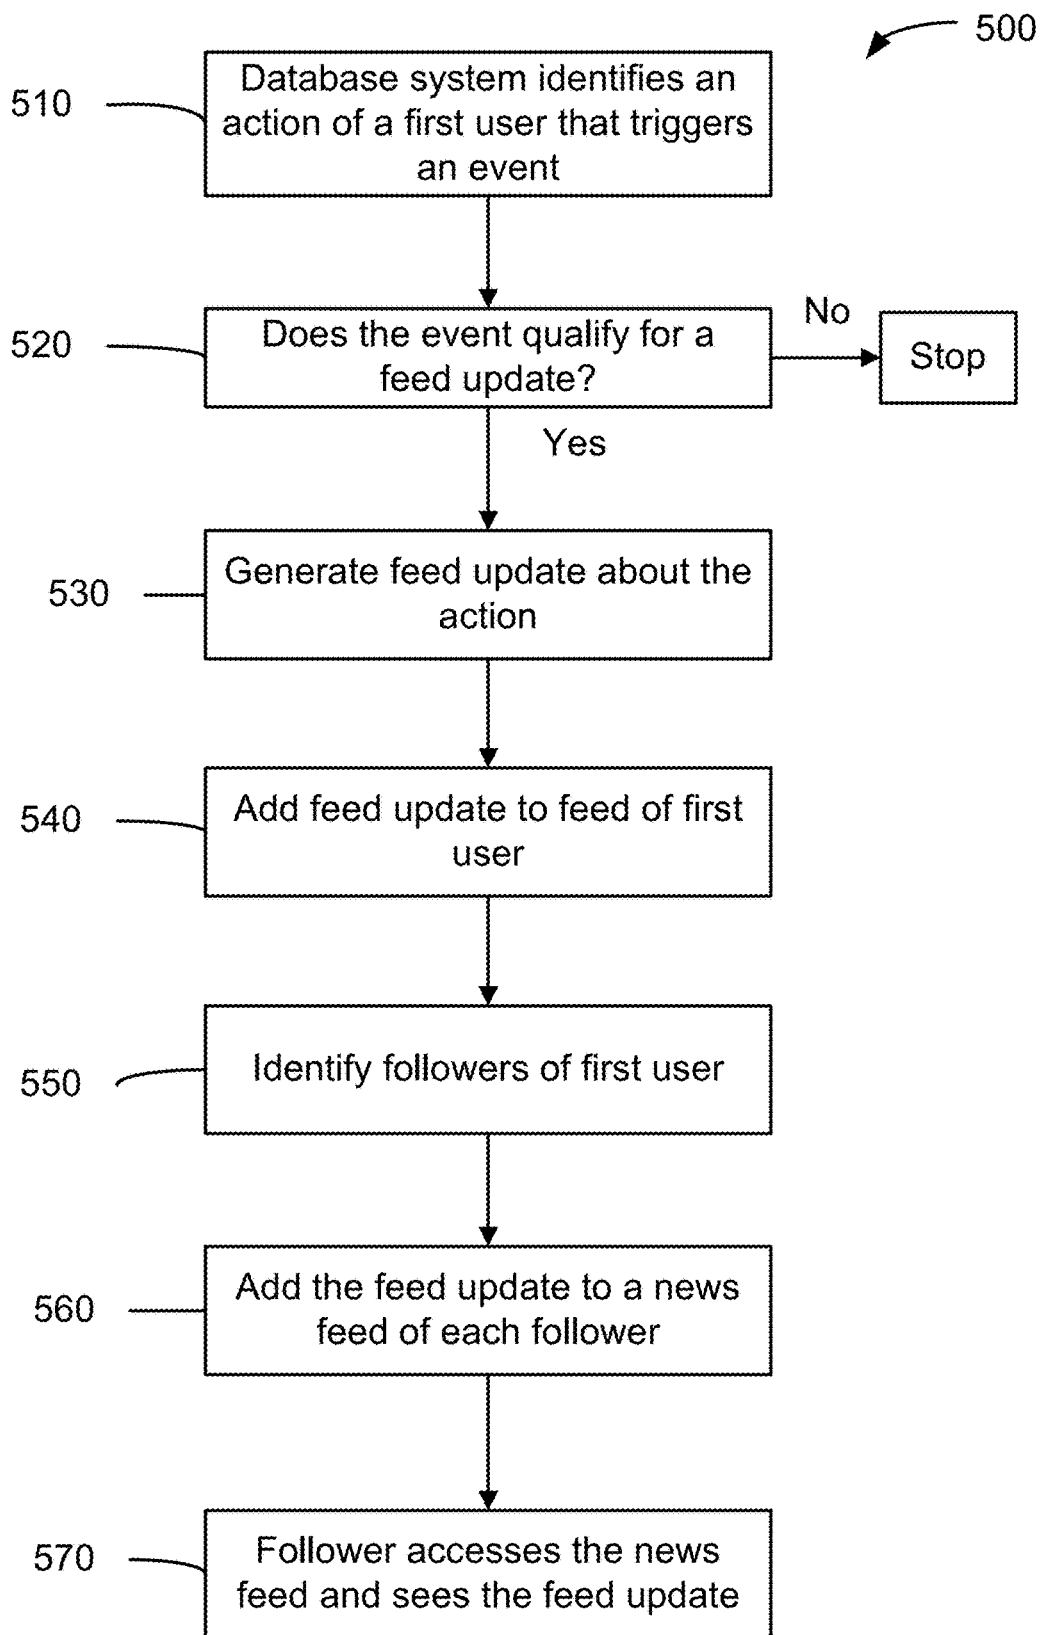


Figure 5

600

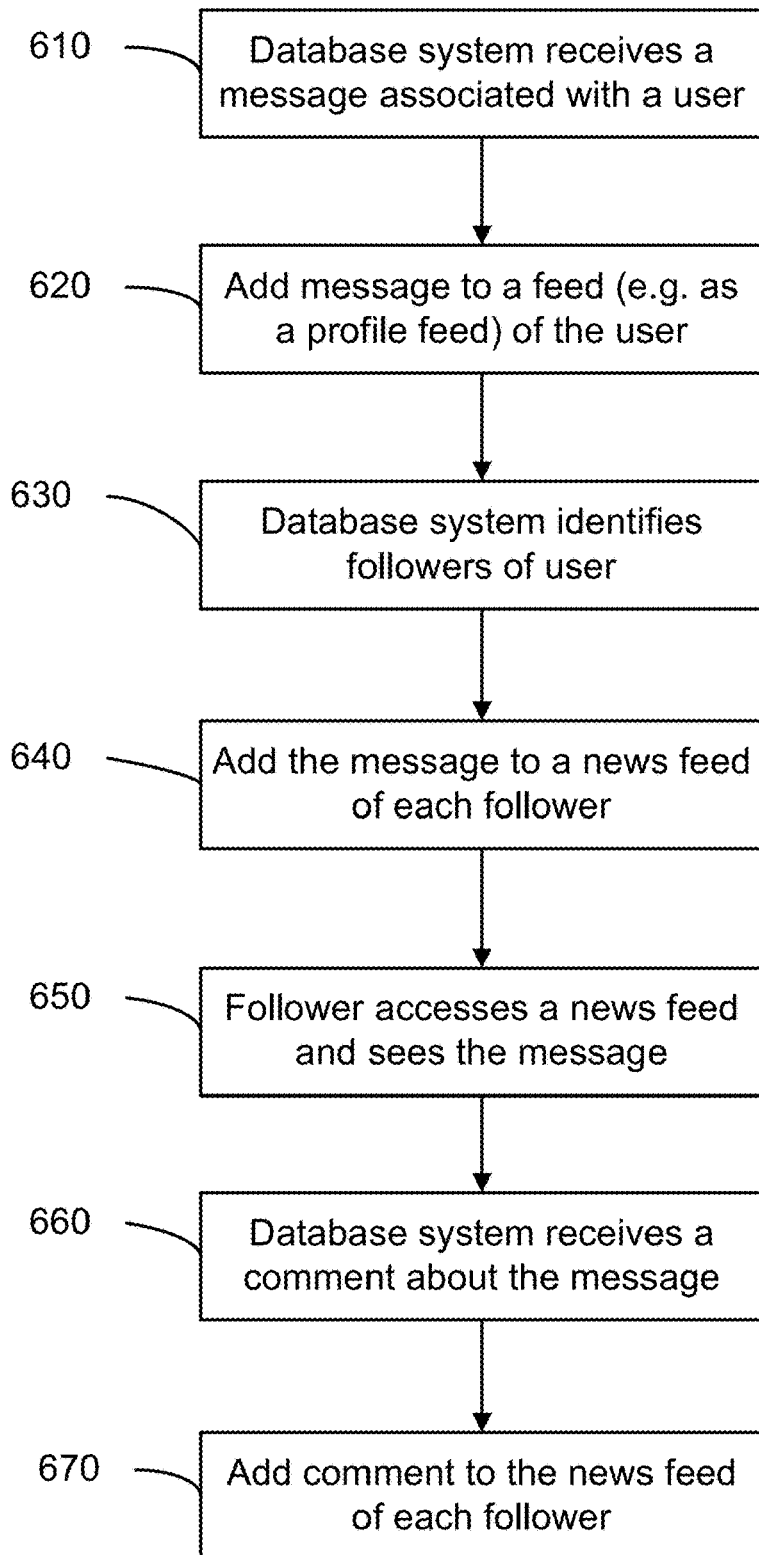


Figure 6

salesforce

[Parker Harris](#) [Setup](#) [Help](#) [Logout](#)

[Home](#) [My Profile](#) **Groups** [Accounts](#) [Contacts](#) [Opportunities](#) [Reports](#) [Dashboards](#)

XYZ Competitive Group

A place to share information that will allow us to compete against XYZ more effectively.

What are you working on?

Attach
Share

[group settings](#) [member settings](#)

Members 1 See All [150]

Following 1 See All [150]

Find people or things to follow as a group.

Twitter

XYZ Competitive Group
Keywords: Dell, computers, netbooks

sam_adams

I heard Dell is coming out with awesome deals this Spring. Can't wait to see the new **netbook**.

erin_martin123

Got my new **netbook** today!

sam_adams

I have too many **computers**, thinking about selling my desktop.

[edit](#)

Today July 24, 2009

Bill Bauer has posted the document. Competitive Insights -- Netbooks 3:11pm [Comment](#)

Netbooks
Competitive Insights

By Bill Bauer
September 23, 2007

710 [View](#) [Download \(PPT\)](#)

Competitive Insights -- Netbooks

This document outlines the underlying components of our business, although it does not go into quite a bit of depth, it's very readable for the more novice team members.

Parker Harris Does anyone have the keynote dect from SVUC '08? 3:11pm [Comment](#)

Ella Johnson I think Bill has it. Yesterday, 4:18 PM

James Saxon Dnot sure, did you check in coniant? It has been reeditied a few times since. FYI Yesterday 4:38 PM

Mary Moore I'll ask around, talk to Jessica in Marketing. Yesterday 4:18 PM

Bill Bauer The file is on the intranet under SVUC. It's the only one. Yesterday, 4:38 PM

Write a Comment

720 [View](#) [Download \(PPT\)](#)

730 [View](#) [Download \(PPT\)](#)

CIRRUS COMPUTERS

Edit Group

Invite [+]

Search Chatter

All Chatter

Objects

Files

Opportunities

My Team

Create List

What are you working on?

Attach
Share

Figure 7




Figure 8

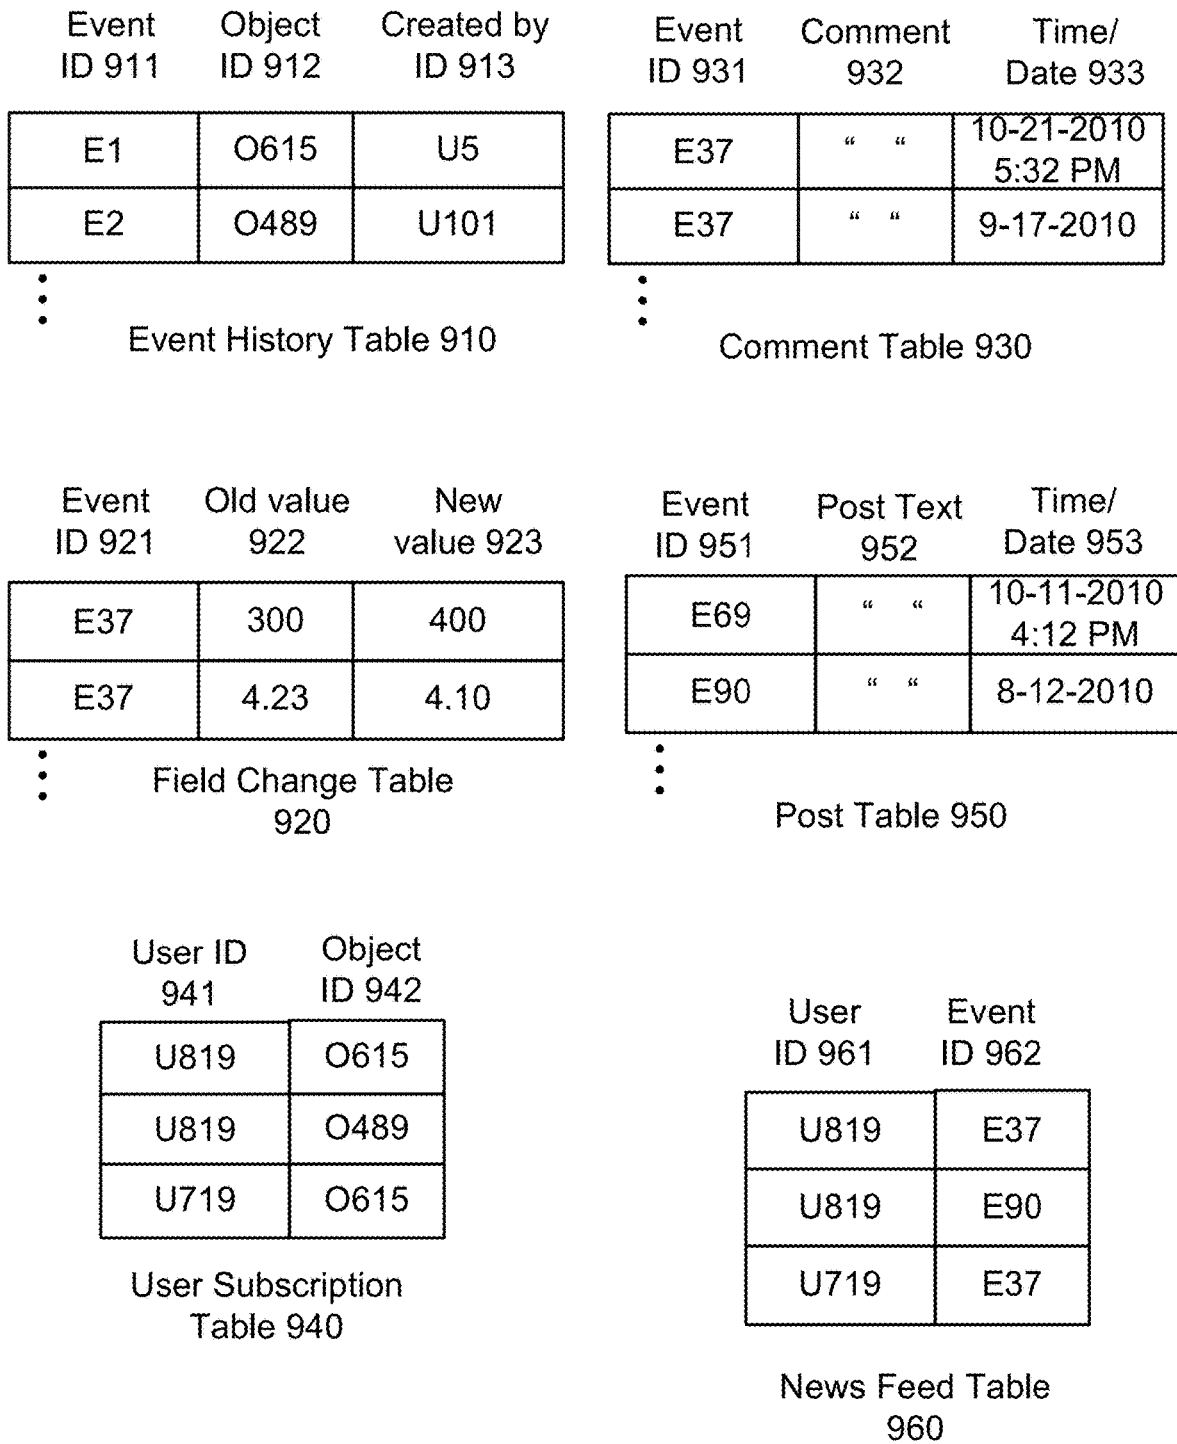


Figure 9A

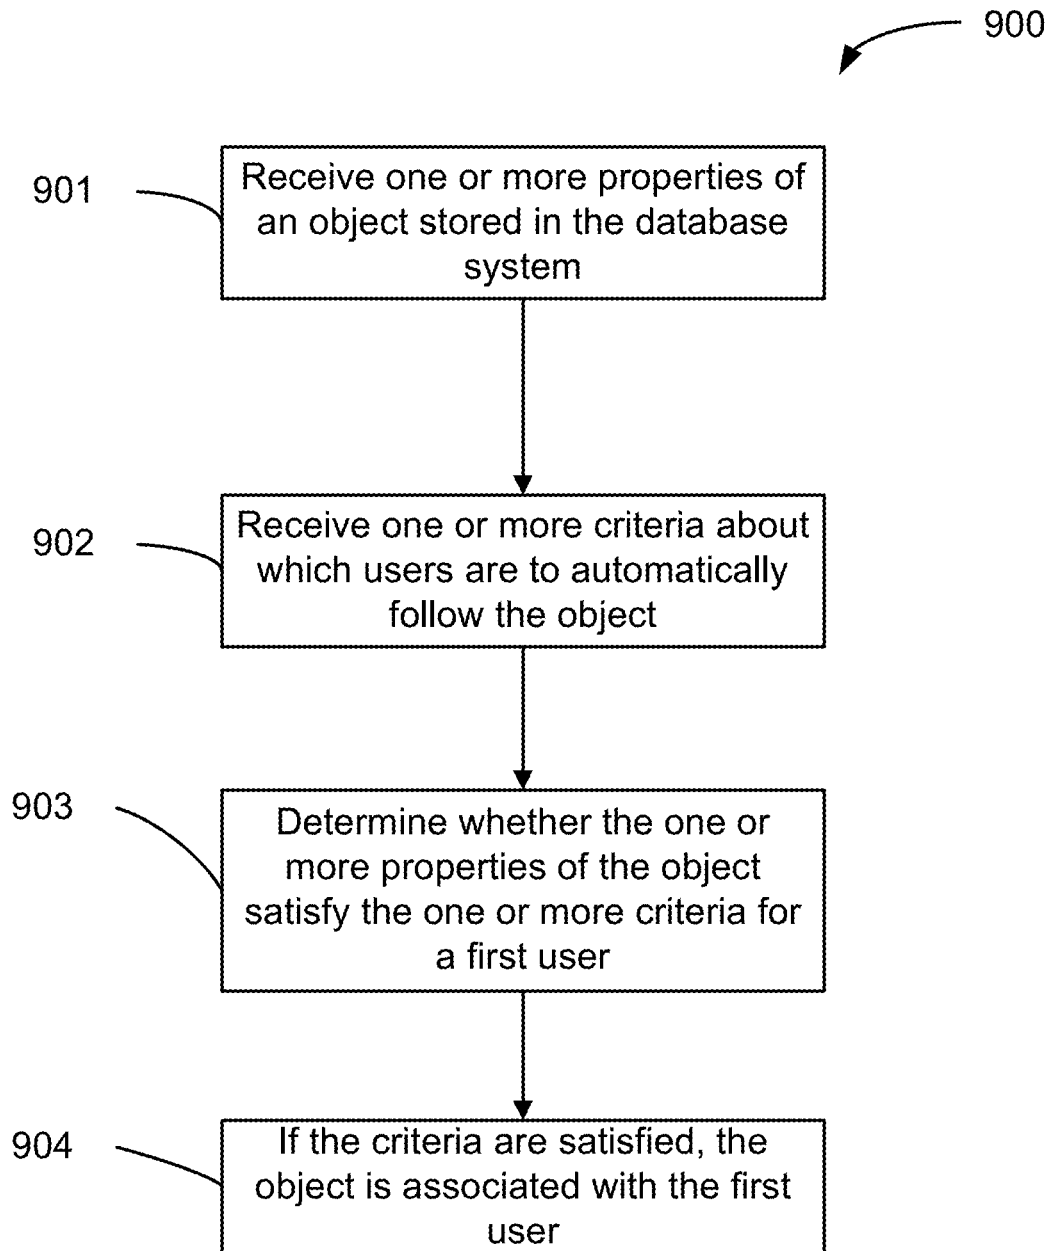


Figure 9B

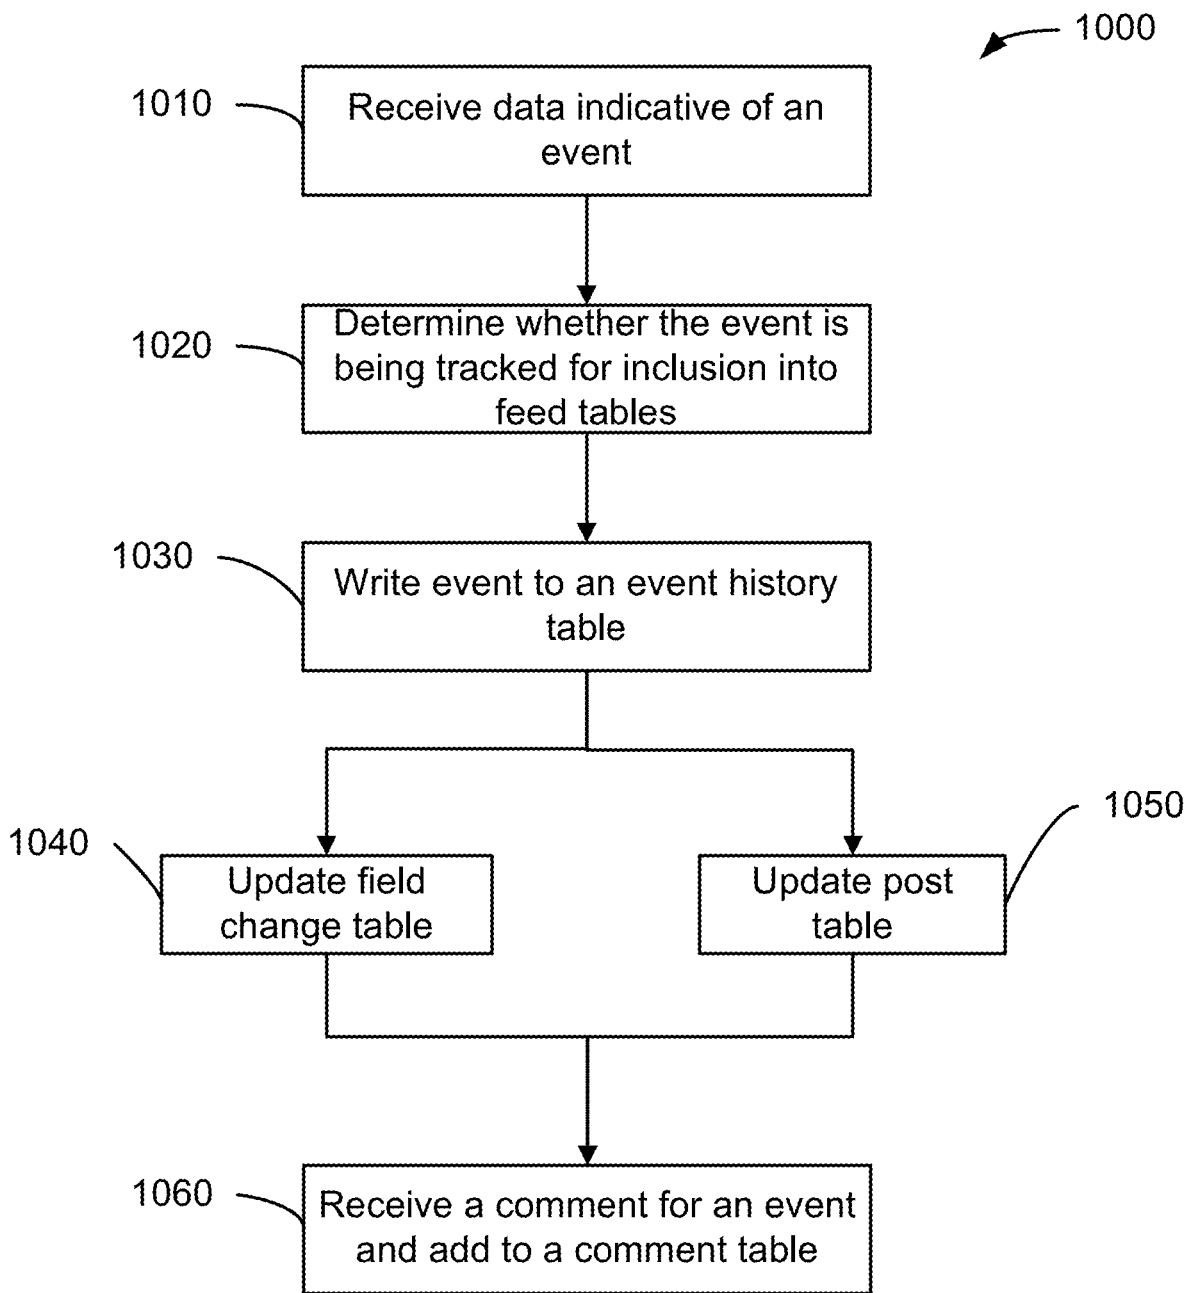


Figure 10

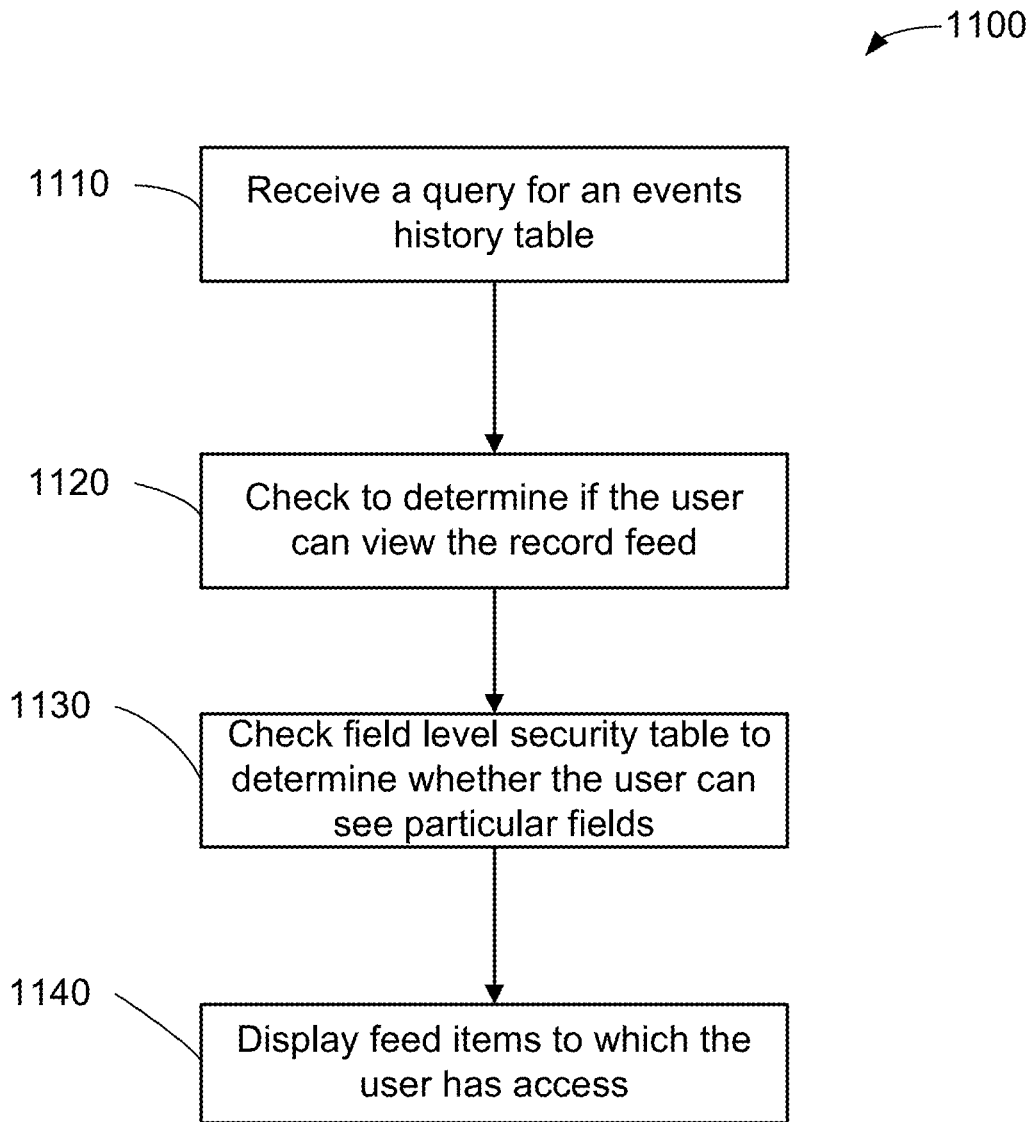


Figure 11

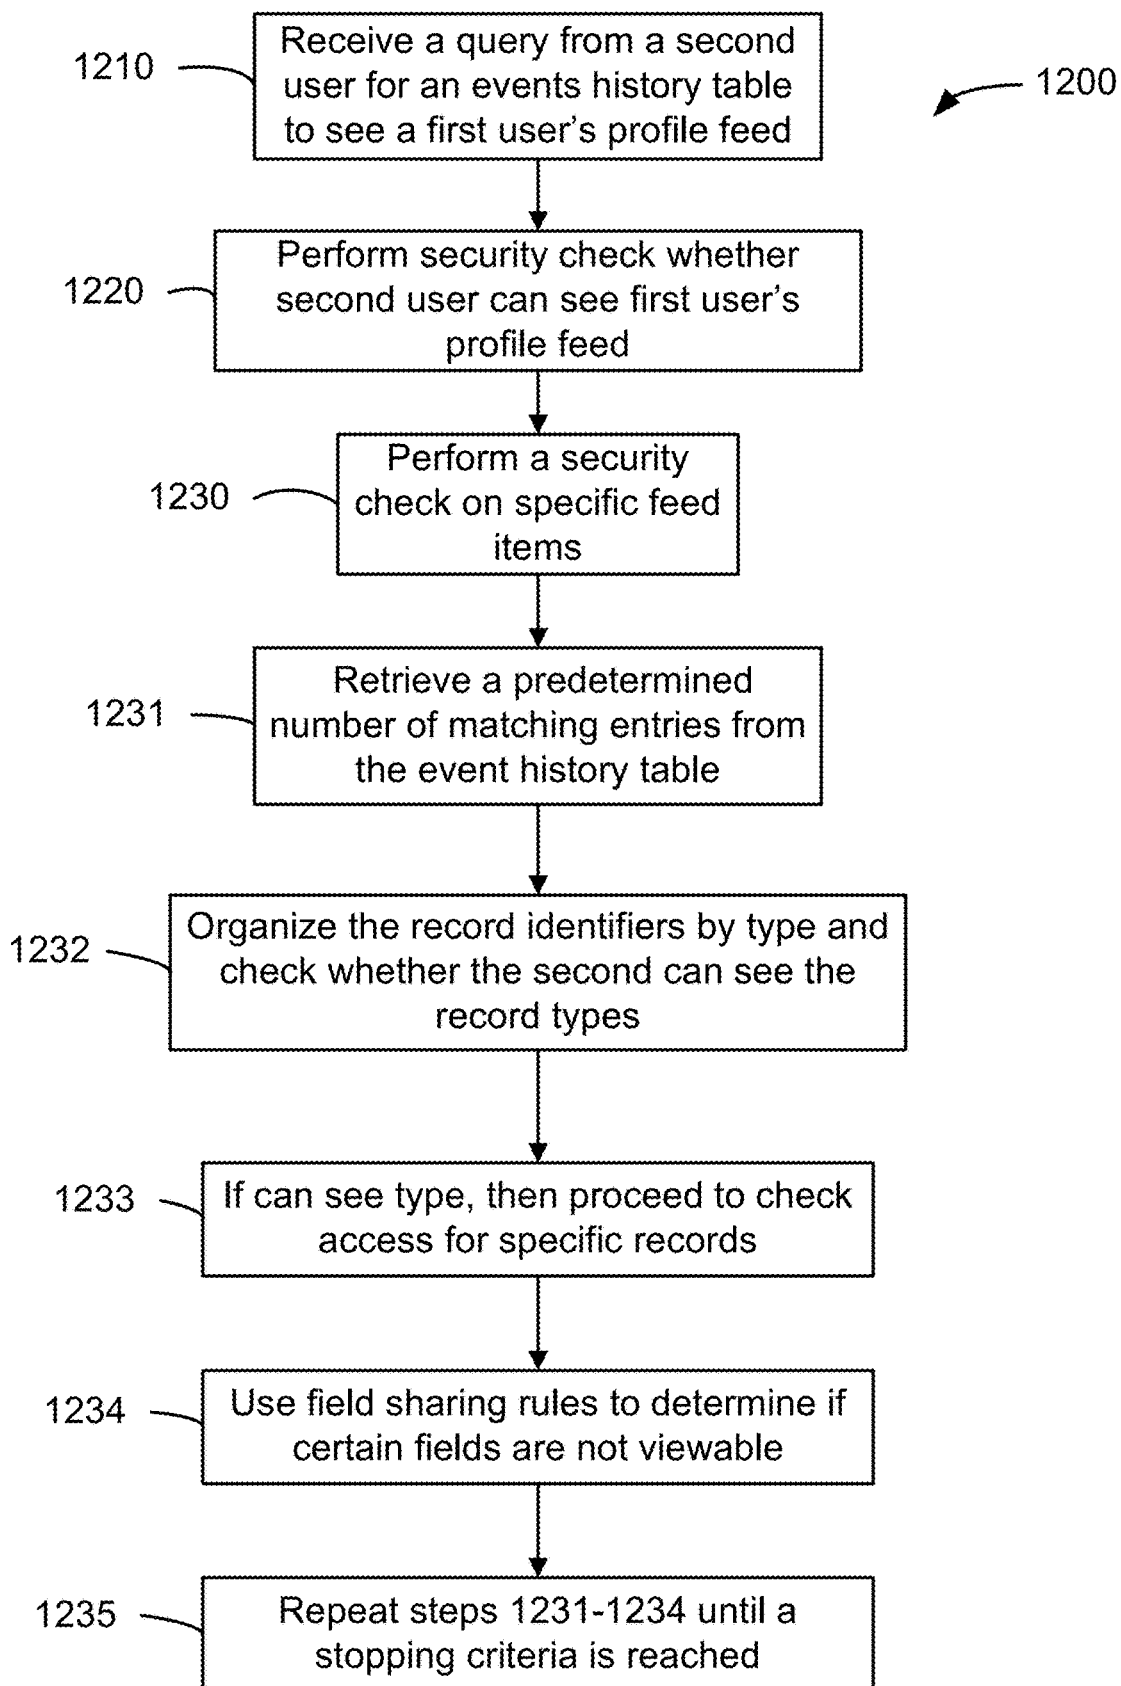


Figure 12

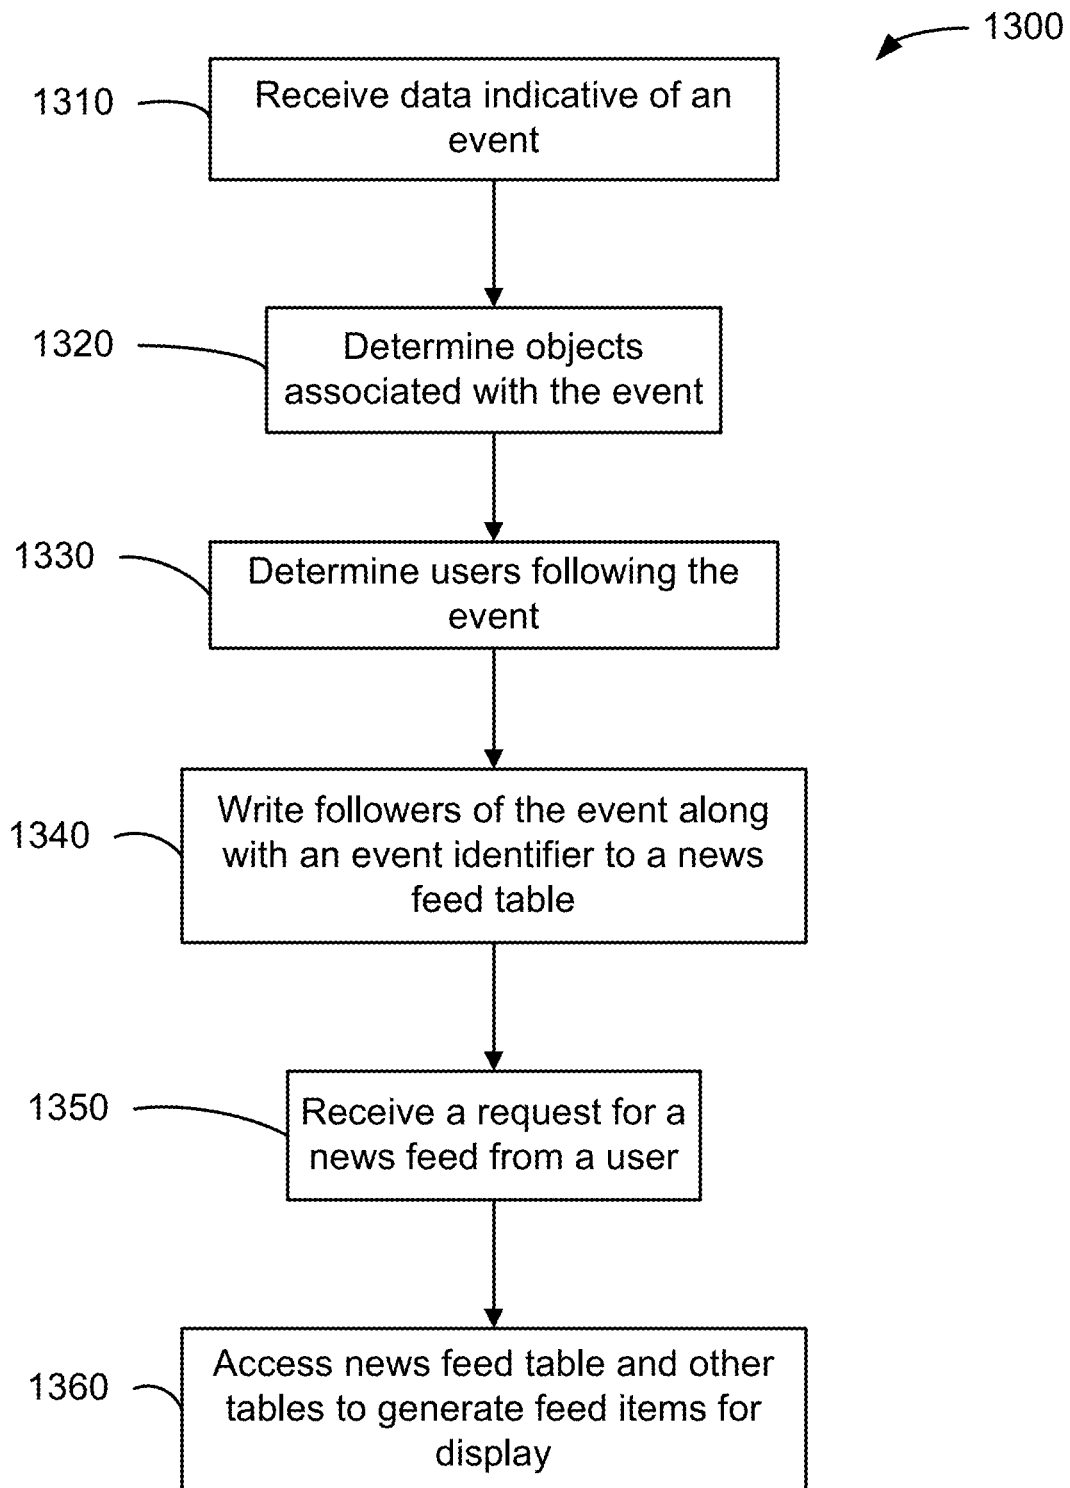


Figure 13

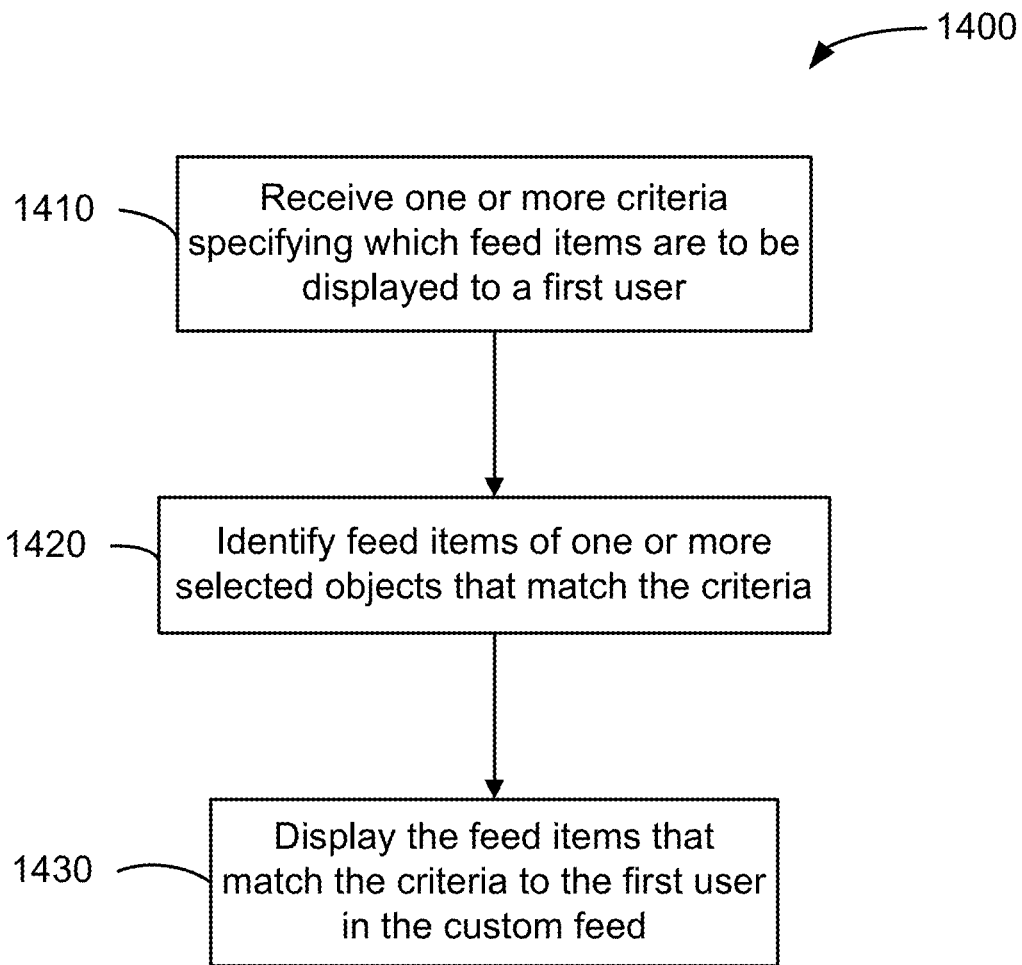


Figure 14

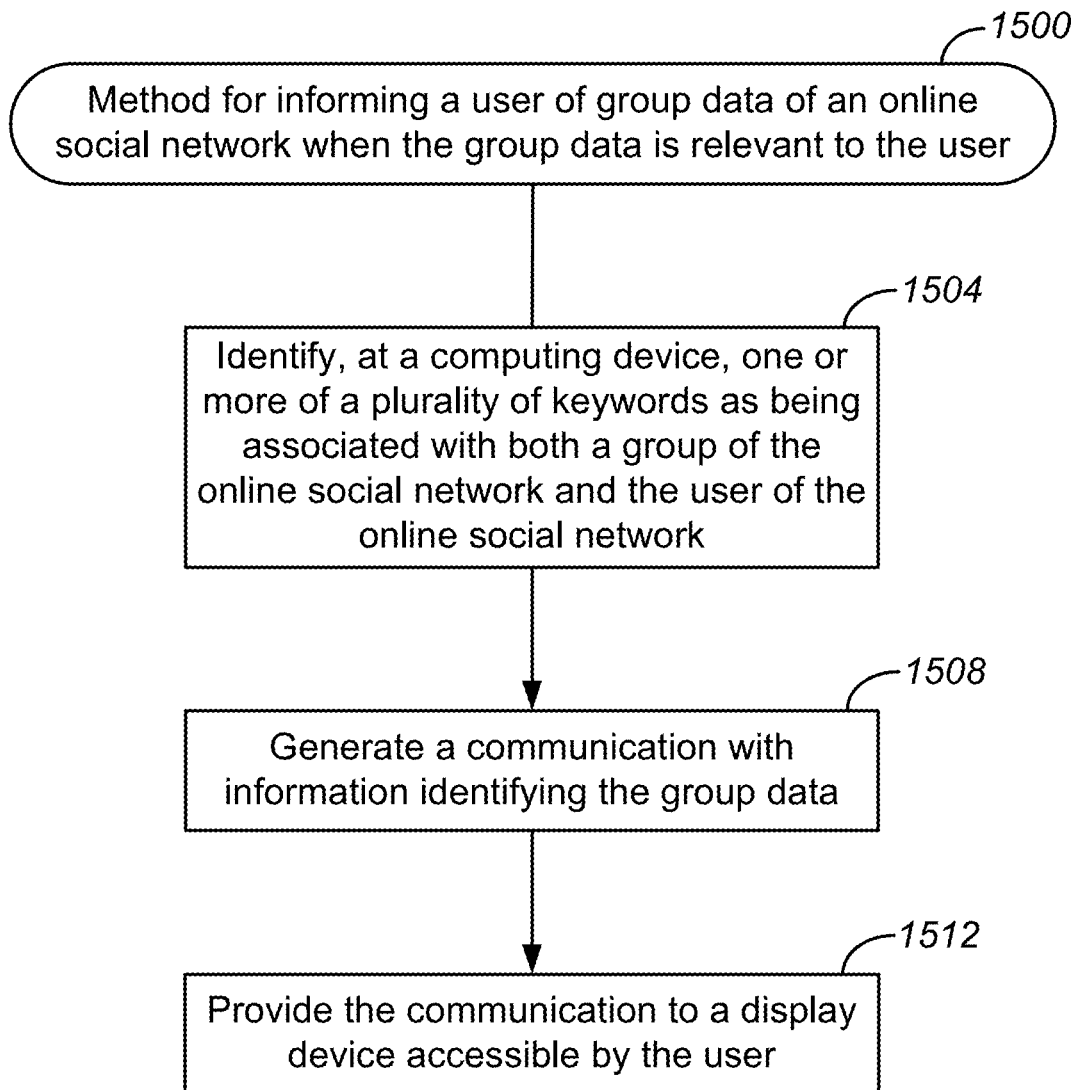


Figure 15

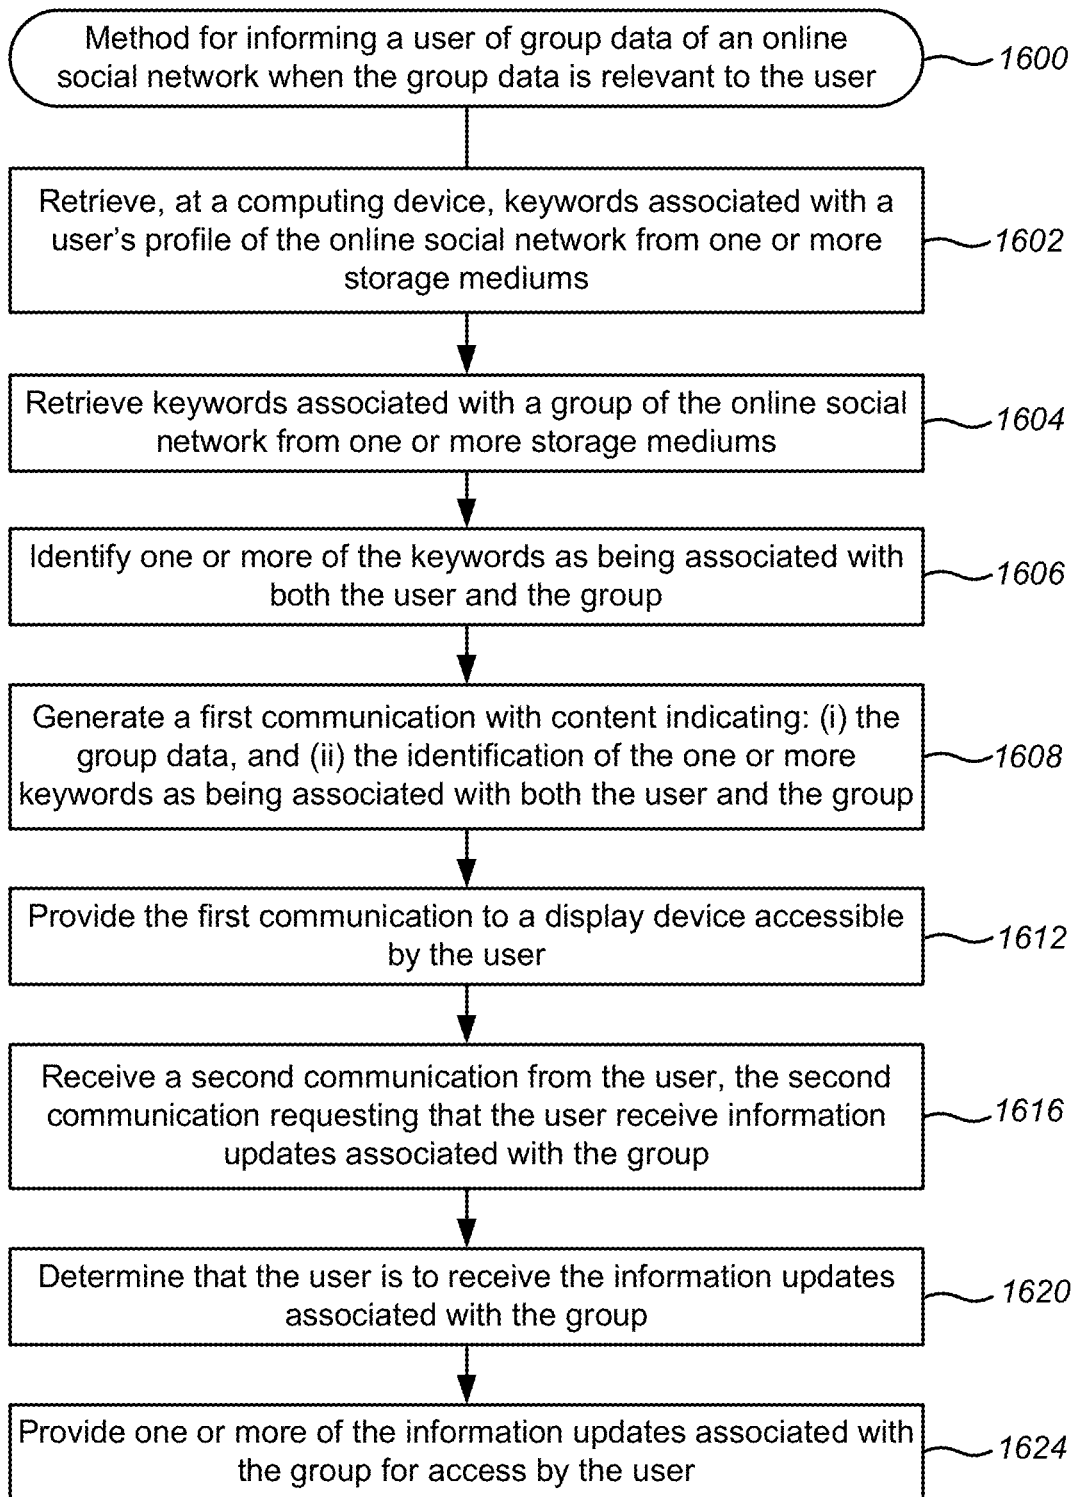


Figure 16

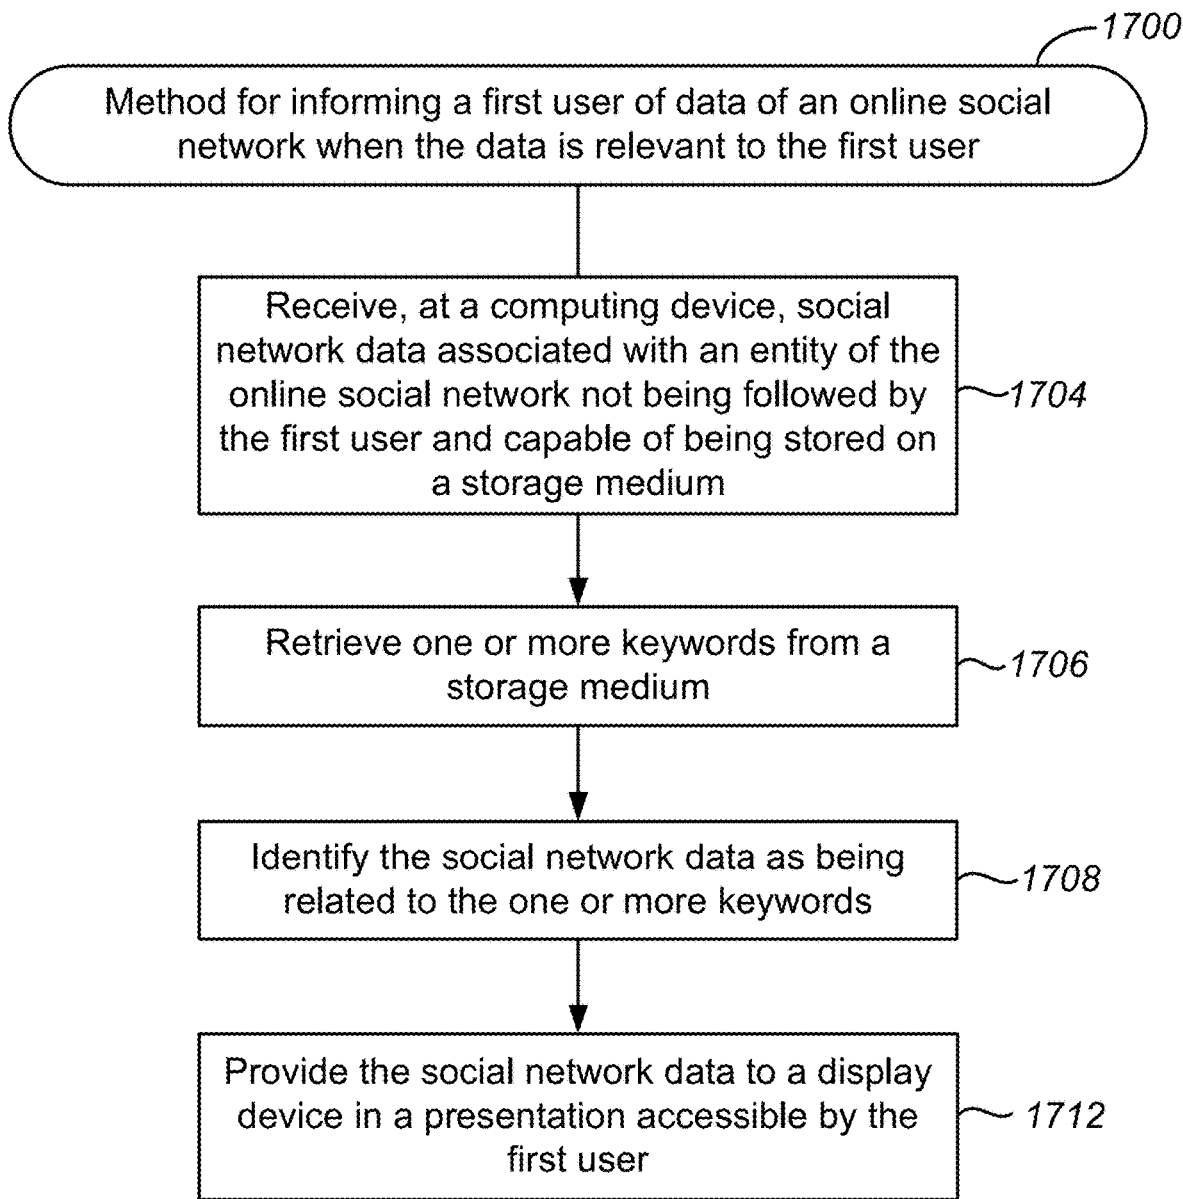


Figure 17

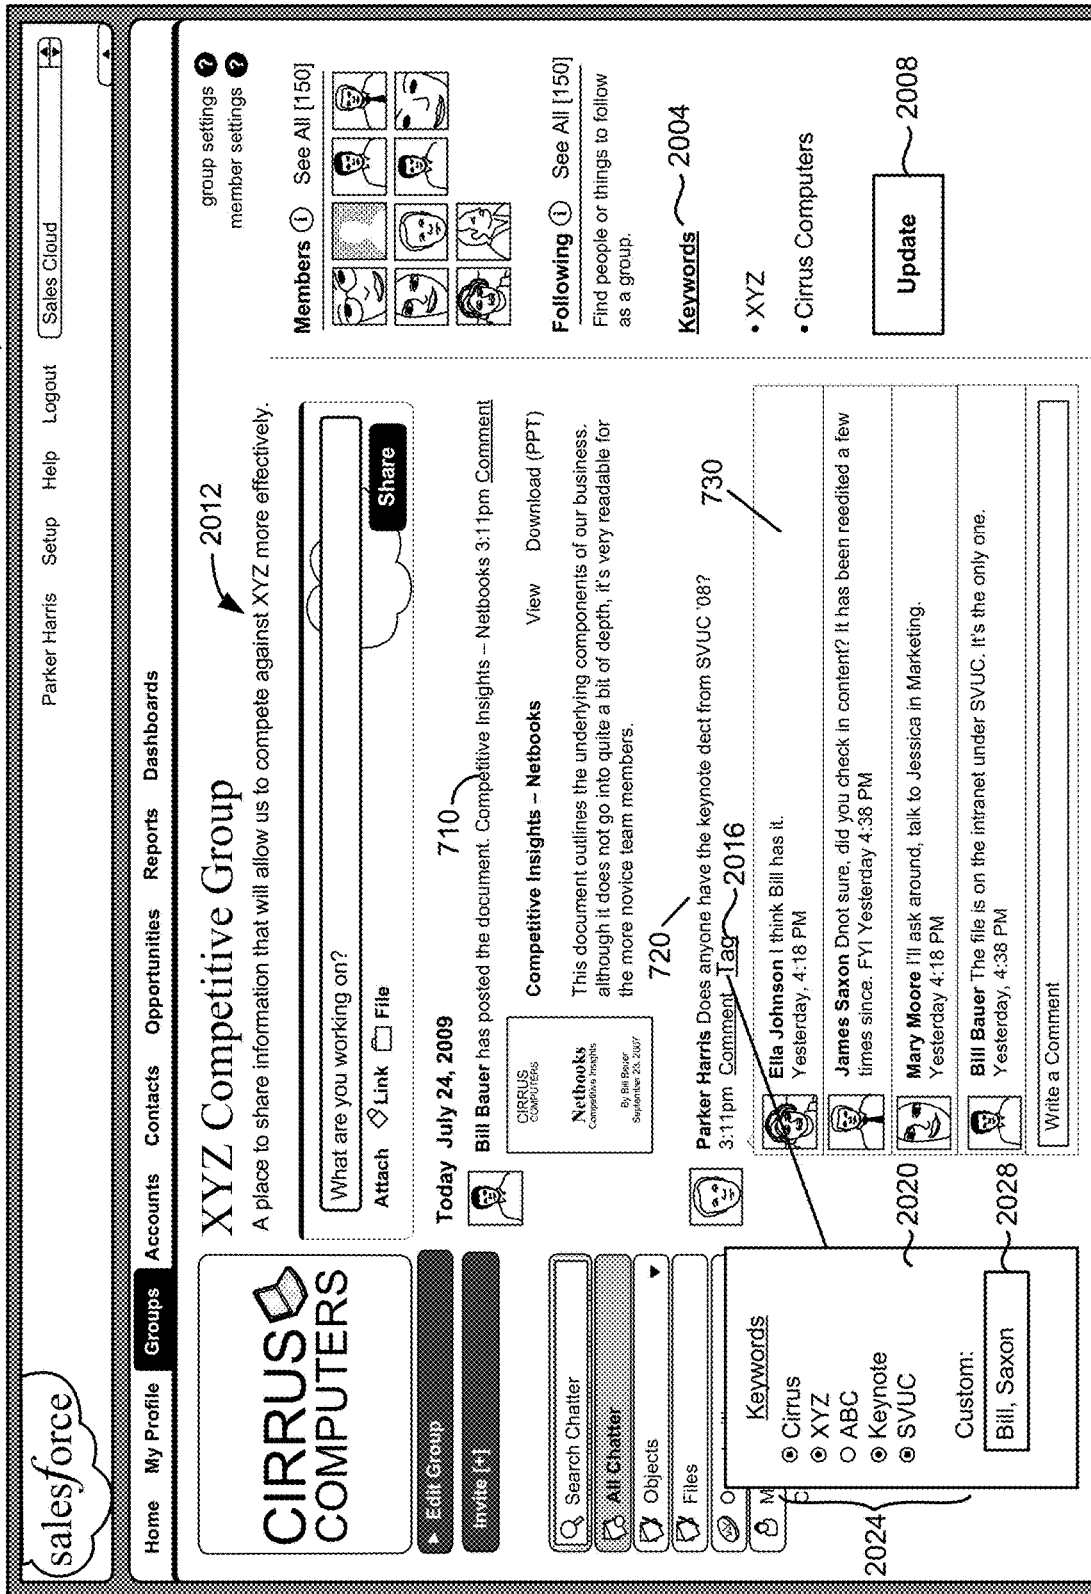


Figure 18

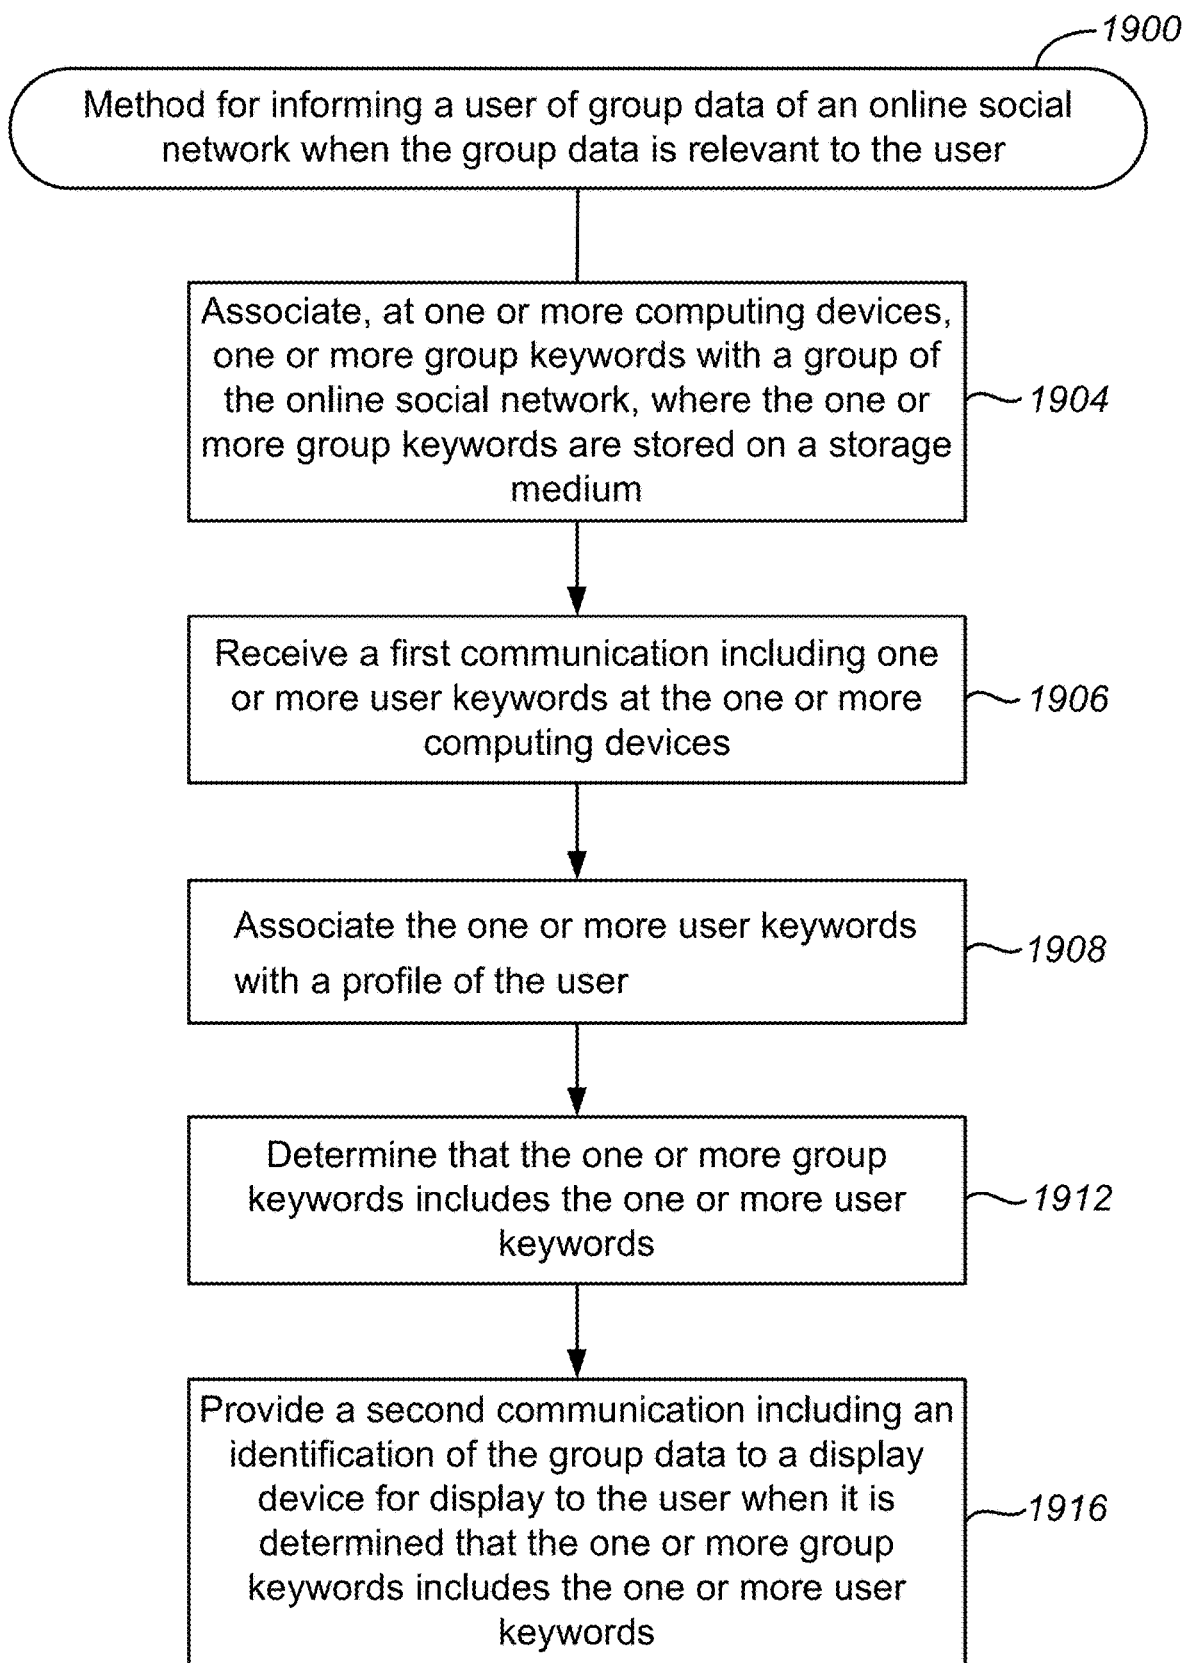


Figure 19

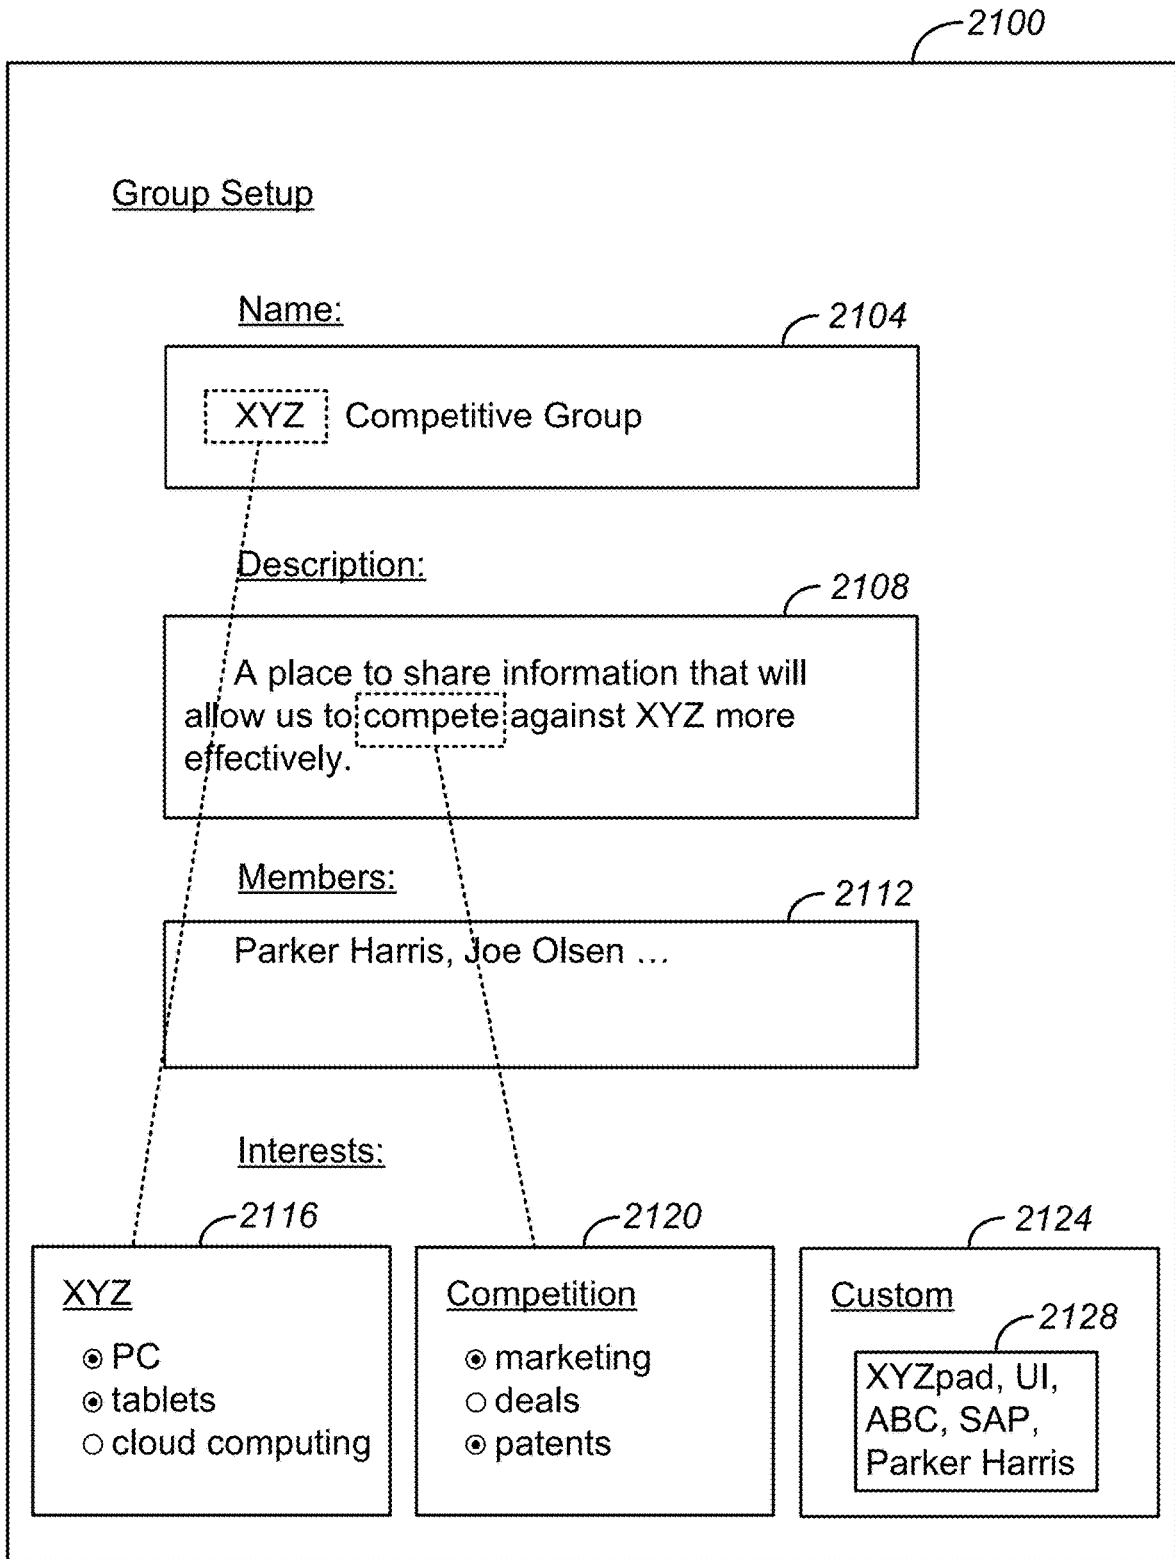


Figure 20

2200

New User Setup

Choose User ID: 2204

newuser12

Select Password: 2208

Enter Job Description: 2212

Associate Patent Counsel, IP
Department, Legal Division

Enter Bio: 2216

I finished law school at BYU in 2005. After 7
years at SAP, I am excited to start this new
position...

Interests:

2220 2224 2228

<p><u>Patent</u></p> <ul style="list-style-type: none"><input checked="" type="radio"/> enforcement<input type="radio"/> prosecution<input checked="" type="radio"/> defense<input checked="" type="radio"/> licensing	<p><u>Legal</u></p> <ul style="list-style-type: none"><input checked="" type="radio"/> transactions<input checked="" type="radio"/> litigation<input checked="" type="radio"/> counseling	<p><u>Custom</u> 2232</p> <p>Joseph Olsen, chatter, social networking, UI</p>
---	---	---

Figure 21

**COMPUTER IMPLEMENTED METHODS
AND APPARATUS FOR INFORMING A USER
OF SOCIAL NETWORK DATA WHEN THE
DATA IS RELEVANT TO THE USER**

COPYRIGHT NOTICE

[0001] A portion of the disclosure of this patent document contains material, which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

INCORPORATION BY REFERENCE

[0002] An Application Data Sheet is filed concurrently with this specification as part of the present application. Each application that the present application claims benefit of or priority to as identified in the concurrently filed Application Data Sheet is incorporated by reference herein in its entirety and for all purposes.

TECHNICAL FIELD

[0003] This patent document relates generally to providing on-demand services in an online social network using a database system and, more specifically, to techniques for informing users of social network data.

BACKGROUND

[0004] “Cloud computing” services provide shared resources, software, and information to computers and other devices upon request. In cloud computing environments, software can be accessible over the Internet rather than installed locally on in-house computer systems. Cloud computing typically involves over-the-Internet provision of dynamically scalable and often virtualized resources. Technological details can be abstracted from the users, who no longer have need for expertise in, or control over, the technology infrastructure “in the cloud” that supports them.

[0005] Database resources can be provided in a cloud computing context. However, using conventional database management techniques, it is difficult to know about the activity of other users of a database system in the cloud or other network. For example, the actions of a particular user, such as a salesperson, on a database resource may be important to the user’s boss. The user can create a report about what the user has done and send it to the boss, but such reports may be inefficient, not timely, and incomplete. Also, it may be difficult to identify other users who might benefit from the information in the report.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The included drawings are for illustrative purposes and serve only to provide examples of possible structures and operations for the disclosed inventive systems, apparatus, and methods for informing a user of social network data when the data is relevant to the user. These drawings in no way limit any changes in form and detail that may be made by one skilled in the art without departing from the spirit and scope of the disclosed implementations.

[0007] FIG. 1A shows a block diagram of an example of an environment 10 in which an on-demand database service can be used in accordance with some implementations.

[0008] FIG. 1B shows a block diagram of an example of some implementations of elements of FIG. 1A and various possible interconnections between these elements.

[0009] FIG. 2A shows a system diagram illustrating an example of architectural components of an on-demand database service environment 200 according to some implementations.

[0010] FIG. 2B shows a system diagram further illustrating an example of architectural components of an on-demand database service environment according to some implementations.

[0011] FIG. 3 shows a flowchart of an example of a method 300 for tracking updates to a record stored in a database system, performed in accordance with some implementations.

[0012] FIG. 4 shows a block diagram of an example of components of a database system configuration 400 performing a method for tracking an update to a record according to some implementations.

[0013] FIG. 5 shows a flowchart of an example of a method 500 for tracking actions of a user of a database system, performed in accordance with some implementations.

[0014] FIG. 6 shows a flowchart of an example of a method 600 for creating a news feed from messages created by a user about a record or another user, performed in accordance with some implementations.

[0015] FIG. 7 shows an example of a group feed on a group page according to some implementations.

[0016] FIG. 8 shows an example of a record feed containing a feed tracked update, post, and comments according to some implementations.

[0017] FIG. 9A shows an example of a plurality of tables that may be used in tracking events and creating feeds according to some implementations.

[0018] FIG. 9B shows a flowchart of an example of a method 900 for automatically subscribing a user to an object in a database system, performed in accordance with some implementations.

[0019] FIG. 10 shows a flowchart of an example of a method 1000 for saving information to feed tracking tables, performed in accordance with some implementations.

[0020] FIG. 11 shows a flowchart of an example of a method 1100 for reading a feed item as part of generating a feed for display, performed in accordance with some implementations.

[0021] FIG. 12 shows a flowchart of an example of a method 1200 for reading a feed item of a profile feed for display, performed in accordance with some implementations.

[0022] FIG. 13 shows a flowchart of an example of a method 1300 of storing event information for efficient generation of feed items to display in a feed, performed in accordance with some implementations.

[0023] FIG. 14 shows a flowchart of an example of a method 1400 for creating a custom feed for users of a database system using filtering criteria, performed in accordance with some implementations.

[0024] FIG. 15 shows a flowchart of an example of a method 1500 for informing a user of group data of an online

social network when the group data is relevant to the user, performed in accordance with some implementations.

[0025] FIG. 16 shows a flowchart of an example of a method 1600 for informing a user of group data of an online social network when the group data is relevant to the user, performed in accordance with some implementations.

[0026] FIG. 17 shows a flowchart of an example of a method 1700 for informing a first user of data of an online social network when the data is relevant to the first user, performed in accordance with some implementations.

[0027] FIG. 18 shows an example of a group page 2000 in the form of a graphical user interface (GUI) configured to provide group data and be accessible by various members of the group, according to some implementations.

[0028] FIG. 19 shows a flowchart of an example of a method 1900 for informing a user of group data of an online social network when the group data is relevant to the user, performed in accordance with some implementations.

[0029] FIG. 20 shows an example of a group setup window 2100 for generating and defining information about a particular group, according to some implementations.

[0030] FIG. 21 shows an example of a new user setup window 2200 as part of a registration process for a new user to obtain a user's profile in an online social network, according to some implementations.

DETAILED DESCRIPTION

[0031] Examples of systems, apparatus, and methods according to the disclosed implementations are described in this section. These examples are being provided solely to add context and aid in the understanding of the disclosed implementations. It will thus be apparent to one skilled in the art that implementations may be practiced without some or all of these specific details. In other instances, certain process/method operations, also referred to herein as "blocks," have not been described in detail in order to avoid unnecessarily obscuring implementations. Other applications are possible, such that the following examples should not be taken as definitive or limiting either in scope or setting.

[0032] In the following detailed description, references are made to the accompanying drawings, which form a part of the description and in which are shown, by way of illustration, specific implementations. Although these implementations are described in sufficient detail to enable one skilled in the art to practice the disclosed implementations, it is understood that these examples are not limiting, such that other implementations may be used and changes may be made without departing from their spirit and scope. For example, the blocks of methods shown and described herein are not necessarily performed in the order indicated. It should also be understood that the methods may include more or fewer blocks than are indicated. In some implementations, blocks described herein as separate blocks may be combined. Conversely, what may be described herein as a single block may be implemented in multiple blocks.

[0033] Various implementations described or referenced herein are directed to different methods, apparatus, systems, and computer-readable storage media for informing users of relevant data in an online social network, also referred to herein as a social networking system. One example of an online social network is Chatter®, provided by salesforce.com, inc. of San Francisco, Calif. Online social networks are increasingly becoming a common way to facilitate commu-

nication among people and groups of people, any of whom can be recognized as users of a social networking system. Some online social networks can be implemented in various settings, including organizations, e.g., enterprises such as companies or business partnerships, academic institutions, or groups within such an organization. For instance, Chatter® can be used by employee users in a division of a business organization to share data, communicate, and collaborate with each other for various purposes.

[0034] In some online social networks, users can access one or more information feeds, which include information updates presented as items or entries in the feed. Such a feed item can include a single information update or a collection of individual information updates. A feed item can include various types of data including character-based data, audio data, image data and/or video data. An information feed can be displayed in a graphical user interface (GUI) on a display device such as the display of a computing device as described below. The information updates can include various social network data from various sources and can be stored in an on-demand database service environment. In some implementations, the disclosed methods, apparatus, systems, and computer-readable storage media may be configured or designed for use in a multi-tenant database environment.

[0035] In some implementations, an online social network may allow a user to follow data objects in the form of records such as cases, accounts, or opportunities, in addition to following individual users and groups of users. The "following" of a record stored in a database, as described in greater detail below, allows a user to track the progress of that record. Updates to the record, also referred to herein as changes to the record, are one type of information update that can occur and be noted on an information feed such as a record feed or a news feed of a user subscribed to the record. Examples of record updates include field changes in the record, updates to the status of a record, as well as the creation of the record itself. Some records are publicly accessible, such that any user can follow the record, while other records are private, for which appropriate security clearance/permissions are a prerequisite to a user following the record.

[0036] Information updates can include various types of updates, which may or may not be linked with a particular record. For example, information updates can be user-submitted messages or can otherwise be generated in response to user actions or in response to events. Examples of messages include: posts, comments, indications of a user's personal preferences such as "likes" and "dislikes", updates to a user's status, uploaded files, and hyperlinks to social network data or other network data such as various documents and/or web pages on the Internet. Posts can include alpha-numeric or other character-based user inputs such as words, phrases, statements, questions, emotional expressions, and/or symbols. Comments generally refer to responses to posts, such as words, phrases, statements, answers, questions, and reactionary emotional expressions and/or symbols. Multimedia data can be included in, linked with, or attached to a post or comment. For example, a post can include textual statements in combination with a JPEG image or animated image. A like or dislike can be submitted in response to a particular post or comment. Examples of uploaded files include presentations, documents, multimedia files, and the like.

[0037] Users can follow a record by subscribing to the record, as mentioned above. Users can also follow other entities such as other types of data objects, other users, and groups of users. Feed tracked updates regarding such entities are one type of information update that can be received and included in the user's news feed. Any number of users can follow a particular entity and thus view information updates pertaining to that entity on the users' respective news feeds. In some social networks, users may follow each other by establishing connections with each other, sometimes referred to as "friending" one another. By establishing such a connection, one user may be able to see information generated by, generated about, or otherwise associated with another user. For instance, a first user may be able to see information posted by a second user to the second user's personal social network page. One implementation of such a personal social network page is a user's profile page, for example, in the form of a web page representing the user's profile. In one example, when the first user is following the second user, the first user's news feed can receive a post from the second user submitted to the second user's profile feed, also referred to herein as the user's "wall," which is one example of an information feed displayed on the user's profile page.

[0038] In some implementations, an information feed may be specific to a group of users of an online social network. For instance, a group of users may publish a news feed. Members of the group may view and post to this group feed in accordance with a permissions configuration for the feed and the group. Information updates in a group context can also include changes to group status information.

[0039] In some implementations, when data such as posts or comments input from one or more users are submitted to an information feed for a particular user, group, object, or other construct within an online social network, an email notification or other type of network communication may be transmitted to all users following the user, group, or object in addition to the inclusion of the data as a feed item in one or more feeds, such as a user's profile feed, a news feed, or a record feed. In some online social networks, the occurrence of such a notification is limited to the first instance of a published input, which may form part of a larger conversation. For instance, a notification may be transmitted for an initial post, but not for comments on the post. In some other implementations, a separate notification is transmitted for each such information update.

[0040] Some implementations of the disclosed systems, apparatus, methods, and computer readable storage media are configured to inform one or more users of any social network data identified as relevant to the user(s). In some implementations, a user can be identified as interested in one or more topics and, without having to follow any entities, the user can receive social network data relevant to the user's identified topics. For example, one or more computing devices can communicate relevant social network data to the user that identifies or is otherwise associated with a group or other entity of the online social network, such as a record, another user's profile, and other objects. Often, the user may not be aware of the group or other entity, so the user may not be following the entity or attempting to access the entity's data even when the data is of possible interest to the user.

[0041] In some implementations, the social network data can be identified as related to one or more keywords associated with the user, for instance, when such keywords

are linked with or stored in the user's profile in a database system. When social network data is determined to be related to user keywords, the data can be transmitted to the user. Thus, the social networking system can be configured to automatically "push" relevant social network content to a user having a certain profile or certain profile attributes in the form of keywords. Thus, rather than rely on user action, relevant posts and other information updates can be emailed to a user or caused to automatically appear on the user's profile, by way of example.

[0042] In one example, when an entity of possible interest to the user is a group, one or more keywords representing the group's purpose or topics of conversation can be associated with the group. Various group data such as the name of the group, names of members of the group, information updates in the group feed, group records, and other group-related data can be identified as relevant to the user when one or more of the user keywords match one or more of the group keywords. Various criteria can be applied to score data of possible interest to a given user and filter out less relevant data from being communicated to the user.

[0043] In some online social networks, a user can find it difficult to keep up with all of the different groups in existence at any moment, thereby hindering the user's ability to stay abreast of conversations and developments with respect to topics that the user would be interested in reading. Also, in some online social networks, there can be thousands or even millions of users, and any of the users can create a group. Thus, there can be thousands or millions of groups. A group leader may directly invite a few users to join when the group is created, and others having personal interactions with the initial group members may learn about the group and request to join. However, the majority of users of larger social networks may not know the group members and thus may never become aware of the group. In addition, it can be difficult for a user to ascertain groups of interest to the user that the user may wish to join, even when the user knows of all of the available groups. For instance, in some social networks, a list of all of the groups is maintained. It could take a user several hours, days, weeks, etc. to scan through hundreds or thousands of group names and descriptions and still not understand all of the groups' purposes and/or what topics the groups have been discussing. A user may give up and decide it is not worth the effort to find a group corresponding to the user's interests. It can similarly be difficult for the group to identify users who might be interested in participating in the group. These disconnects can have a significant detrimental impact on the underlying purpose of most social networks to connect and facilitate fruitful dialogue and collaboration among users sharing the same interests.

[0044] By way of example, Zach is a user interface (UI) developer and has recently been hired by Cloud Programming, Inc. as a new employee in one of Cloud Programming's software development teams. Shortly after Zach establishes a user profile on Cloud Programming, Inc.'s social network, Chatter®, Zach wants to find and join Chatter® UI-related groups and follow groups, records, users, and other entities having anything to do with UI's. Zach has basically no idea what interesting data may exist in Cloud Programming's implementation of Chatter® that could help him get up-to-speed on various UI's as a new member of the software development team. Zach may already be subscribed to one or more groups; however, there

may also be a number of relevant UI groups that are unknown to him. As a result, Zach may not learn any otherwise helpful information from these undiscovered groups. For instance, if another user publishes a post to a feed that is relevant to Zach as a UI developer, Zach may never learn of the post unless the user specifically mentions Zach within the post or proactively posts to Zach's profile feed.

[0045] With some implementations of the disclosed techniques, groups having information related to UI's can have one or more keywords such as "UI", "user interface", "UI developer", etc., attached to the group. By the same token, a user such as Zach may "tag" his profile with one or more of the same or similar keywords representing the user's interests. In some instances, the keywords are user-selected or user-defined and attached to the user's profile at the user's discretion. In some other instances, attachment of keywords is automated. For instance, joining a group may cause relevant keywords associated with the group to be associated with the user's profile. When Zach has one or more of the same keywords as the group, such as "UI", Zach can be notified of the name of the group, selected posts to the group feed pertinent to UI, UI-related group records, and/or other data of interest to a user having the keyword "UI", "user interface", or other similar keyword.

[0046] In some instances, group data determined to be related and of possible interest to Zach, by virtue of a shared keyword or keywords, can automatically be communicated to Zach, for instance, by an email to Zach's email address or a post to Zach's profile feed. In some implementations, other users who publish posts to various feeds in the social network have the capability to tag posts with keywords to which the posts may be relevant, and selected posts can then be pushed to users having those keywords. In some instances, the communication to Zach can include or be followed by an invitation to join or follow the group. Some implementations of the disclosed techniques can thus push any available and relevant social network data of possible interest to Zach as opposed to Zach having to spend valuable time and effort to search for the data.

[0047] In some instances, keywords can be associated with Zach's user profile as part of his initial registration as a new user of Cloud Programming's implementation of Chatter®. For instance, keywords could be selected from Zach's job description in his user profile to start automatically pushing data of possible interest to Zach from groups having one or more of the same keywords. The names of any groups followed by Zach's boss or Zach's direct reports could also be pushed to Zach. As explained in greater detail below, Zach could also personally select and enter custom keywords to identify himself as being interested in certain topics. Various social network data having any of Zach's keywords can automatically be communicated to Zach, using some of the disclosed implementations.

[0048] These and other implementations may be embodied in various types of hardware, software, firmware, and combinations thereof. For example, some techniques disclosed herein may be implemented, at least in part, by computer-readable media that include program instructions, state information, etc., for performing various services and operations described herein. Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher-level code that may be executed by a computing device such as a server or other

data processing apparatus using an interpreter. Examples of computer-readable media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media; and hardware devices that are specially configured to store program instructions, such as read-only memory ("ROM") devices and random access memory ("RAM") devices. These and other features of the disclosed implementations will be described in more detail below with reference to the associated drawings.

[0049] The term "multi-tenant database system" can refer to those systems in which various elements of hardware and software of a database system may be shared by one or more customers. For example, a given application server may simultaneously process requests for a great number of customers, and a given database table may store rows of data such as feed items for a potentially much greater number of customers. The term "query plan" generally refers to one or more operations used to access information in a database system.

[0050] A "user profile" or "user's profile" is generally configured to store and maintain data about a given user of the database system. The data can include general information, such as name, title, phone number, a photo, a biographical summary, and a status, e.g., text describing what the user is currently doing. As mentioned below, the data can include messages created by other users. Where there are multiple tenants, a user is typically associated with a particular tenant. For example, a user could be a salesperson of a company, which is a tenant of the database system that provides a database service.

[0051] The term "record" generally refers to a data entity, such as an instance of a data object created by a user of the database service, for example, about a particular (actual or potential) business relationship or project. The data object can have a data structure defined by the database service (a standard object) or defined by a user (custom object). For example, a record can be for a business partner or potential business partner (e.g., a client, vendor, distributor, etc.) of the user, and can include information describing an entire company, subsidiaries, or contacts at the company. As another example, a record can be a project that the user is working on, such as an opportunity (e.g., a possible sale) with an existing partner, or a project that the user is trying to get. In one implementation of a multi-tenant database system, each record for the tenants has a unique identifier stored in a common table. A record has data fields that are defined by the structure of the object (e.g., fields of certain data types and purposes). A record can also have custom fields defined by a user. A field can be another record or include links thereto, thereby providing a parent-child relationship between the records.

[0052] The terms "information feed" and "feed" are used interchangeably herein and generally refer to a combination (e.g., a list) of feed items or entries with various types of information and data. Such feed items can be stored and maintained in one or more database tables, e.g., as rows in the table(s), that can be accessed to retrieve relevant information to be presented as part of a displayed feed. The term "feed item" (or feed element) refers to an item of information, which can be presented in the feed such as a post submitted by a user. Feed items of information about a user can be presented in a user's profile feed of the database, while feed items of information about a record can be

presented in a record feed in the database, by way of example. A profile feed and a record feed are examples of different information feeds. A second user following a first user and a record can receive the feed items associated with the first user and the record for display in the second user's news feed, which is another type of information feed. In some implementations, the feed items from any number of followed users and records can be combined into a single information feed of a particular user.

[0053] As examples, a feed item can be a message, such as a user-generated post of text data, and a feed tracked update to a record or profile, such as a change to a field of the record. Feed tracked updates are described in greater detail below. A feed can be a combination of messages and feed tracked updates. Messages include text created by a user, and may include other data as well. Examples of messages include posts, user status updates, and comments. Messages can be created for a user's profile or for a record. Posts can be created by various users, potentially any user, although some restrictions can be applied. As an example, posts can be made to a wall section of a user's profile page (which can include a number of recent posts) or a section of a record that includes multiple posts. The posts can be organized in chronological order when displayed in a graphical user interface (GUI), for instance, on the user's profile page, as part of the user's profile feed. In contrast to a post, a user status update changes a status of a user and can be made by that user or an administrator. A record can also have a status, the update of which can be provided by an owner of the record or other users having suitable write access permissions to the record. The owner can be a single user, multiple users, or a group. In one implementation, there is only one status for a record.

[0054] In some implementations, a comment can be made on any feed item. In some implementations, comments are organized as a list explicitly tied to a particular feed tracked update, post, or status update. In some implementations, comments may not be listed in the first layer (in a hierarchal sense) of feed items, but listed as a second layer branching from a particular first layer feed item.

[0055] A "feed tracked update," also referred to herein as a "feed update," is one type of information update and generally refers to data representing an event. A feed tracked update can include text generated by the database system in response to the event, to be provided as one or more feed items for possible inclusion in one or more feeds. In one implementation, the data can initially be stored, and then the database system can later use the data to create text for describing the event. Both the data and/or the text can be a feed tracked update, as used herein. In various implementations, an event can be an update of a record and/or can be triggered by a specific action by a user. Which actions trigger an event can be configurable. Which events have feed tracked updates created and which feed updates are sent to which users can also be configurable. Messages and feed updates can be stored as a field or child object of the record. For example, the feed can be stored as a child object of the record.

[0056] A "group" is generally a collection of users. In some implementations, the group may be defined as users with a same or similar attribute, or by membership. In some implementations, a "group feed", also referred to herein as a "group news feed", includes one or more feed items about any user in the group. In some implementations, the group

feed also includes information updates and other feed items that are about the group as a whole, the group's purpose, the group's description, and group records and other objects stored in association with the group. Threads of information updates including group record updates and messages, such as posts, comments, likes, etc., can define group conversations and change over time.

[0057] An "entity feed" or "record feed" generally refers to a feed of feed items about a particular record in the database, such as feed tracked updates about changes to the record and posts made by users about the record. An entity feed can be composed of any type of feed item. Such a feed can be displayed on a page such as a web page associated with the record, e.g., a home page of the record. As used herein, a "profile feed" or "user's profile feed" is a feed of feed items about a particular user. In one example, the feed items for a profile feed include posts and comments that other users make about or send to the particular user, and status updates made by the particular user. Such a profile feed can be displayed on a page associated with the particular user. In another example, feed items in a profile feed could include posts made by the particular user and feed tracked updates initiated based on actions of the particular user.

[0058] I. General Overview

[0059] Systems, apparatus, and methods are provided for implementing enterprise level social and business information networking. Such implementations can provide more efficient use of a database system. For instance, a user of a database system may not easily know when important information in the database has changed, e.g., about a project or client. Implementations can provide feed tracked updates about such changes and other events, thereby keeping users informed.

[0060] By way of example, a user can update a record, e.g., an opportunity such as a possible sale of 1000 computers. Once the record update has been made, a feed tracked update about the record update can then automatically be provided, e.g., in a feed, to anyone subscribing to the opportunity or to the user. Thus, the user does not need to contact a manager regarding the change in the opportunity, since the feed tracked update about the update is sent via a feed right to the manager's feed page or other page.

[0061] Next, mechanisms and methods for providing systems implementing enterprise level social and business information networking will be described with reference to several implementations. First, an overview of an example of a database system is described, and then examples of tracking events for a record, actions of a user, and messages about a user or record are described. Various implementations about the data structure of feeds, customizing feeds, user selection of records and users to follow, generating feeds, and displaying feeds are also described.

[0062] II. System Overview

[0063] FIG. 1A shows a block diagram of an example of an environment 10 in which an on-demand database service can be used in accordance with some implementations. Environment 10 may include user systems 12, network 14, database system 16, processor system 17, application platform 18, network interface 20, tenant data storage 22, system data storage 24, program code 26, and process space 28. In other implementations, environment 10 may not have all of these components and/or may have other components instead of, or in addition to, those listed above.

[0064] Environment 10 is an environment in which an on-demand database service exists. User system 12 may be implemented as any computing device(s) or other data processing apparatus such as a machine or system that is used by a user to access a database system 16. For example, any of user systems 12 can be a handheld computing device, a mobile phone, a laptop computer, a work station, and/or a network of such computing devices. As illustrated in FIG. 1A (and in more detail in FIG. 1B) user systems 12 might interact via a network 14 with an on-demand database service, which is implemented in the example of FIG. 1A as database system 16.

[0065] An on-demand database service, implemented using system 16 by way of example, is a service that is made available to outside users, who do not need to necessarily be concerned with building and/or maintaining the database system. Instead, the database system may be available for their use when the users need the database system, i.e., on the demand of the users. Some on-demand database services may store information from one or more tenants into tables of a common database image to form a multi-tenant database system (MTS). A database image may include one or more database objects. A relational database management system (RDBMS) or the equivalent may execute storage and retrieval of information against the database object(s). Application platform 18 may be a framework that allows the applications of system 16 to run, such as the hardware and/or software, e.g., the operating system. In some implementations, application platform 18 enables creation, managing and executing one or more applications developed by the provider of the on-demand database service, users accessing the on-demand database service via user systems 12, or third party application developers accessing the on-demand database service via user systems 12.

[0066] The users of user systems 12 may differ in their respective capacities, and the capacity of a particular user system 12 might be entirely determined by permissions (permission levels) for the current user. For example, where a salesperson is using a particular user system 12 to interact with system 16, that user system has the capacities allotted to that salesperson. However, while an administrator is using that user system to interact with system 16, that user system has the capacities allotted to that administrator. In systems with a hierarchical role model, users at one permission level may have access to applications, data, and database information accessible by a lower permission level user, but may not have access to certain applications, database information, and data accessible by a user at a higher permission level. Thus, different users will have different capabilities with regard to accessing and modifying application and database information, depending on a user's security or permission level, also called authorization.

[0067] Network 14 is any network or combination of networks of devices that communicate with one another. For example, network 14 can be any one or any combination of a LAN (local area network), WAN (wide area network), telephone network, wireless network, point-to-point network, star network, token ring network, hub network, or other appropriate configuration. Network 14 can include a TCP/IP (Transfer Control Protocol and Internet Protocol) network, such as the global internetwork of networks often referred to as the "Internet" with a capital "I." The Internet will be used in many of the examples herein. However, it should be understood that the networks that the present

implementations might use are not so limited, although TCP/IP is a frequently implemented protocol.

[0068] User systems 12 might communicate with system 16 using TCP/IP and, at a higher network level, use other common Internet protocols to communicate, such as HTTP, FTP, AFS, WAP, etc. In an example where HTTP is used, user system 12 might include an HTTP client commonly referred to as a "browser" for sending and receiving HTTP signals to and from an HTTP server at system 16. Such an HTTP server might be implemented as the sole network interface 20 between system 16 and network 14, but other techniques might be used as well or instead. In some implementations, the network interface 20 between system 16 and network 14 includes load sharing functionality, such as round-robin HTTP request distributors to balance loads and distribute incoming HTTP requests evenly over a plurality of servers. At least for users accessing system 16, each of the plurality of servers has access to the MTS' data; however, other alternative configurations may be used instead.

[0069] In one implementation, system 16, shown in FIG. 1A, implements a web-based customer relationship management (CRM) system. For example, in one implementation, system 16 includes application servers configured to implement and execute CRM software applications as well as provide related data, code, forms, web pages and other information to and from user systems 12 and to store to, and retrieve from, a database system related data, objects, and Webpage content. With a multi-tenant system, data for multiple tenants may be stored in the same physical database object in tenant data storage 22, however, tenant data typically is arranged in the storage medium(s) of tenant data storage 22 so that data of one tenant is kept logically separate from that of other tenants so that one tenant does not have access to another tenant's data, unless such data is expressly shared. In certain implementations, system 16 implements applications other than, or in addition to, a CRM application. For example, system 16 may provide tenant access to multiple hosted (standard and custom) applications, including a CRM application. User (or third party developer) applications, which may or may not include CRM, may be supported by the application platform 18, which manages creation, storage of the applications into one or more database objects and executing of the applications in a virtual machine in the process space of the system 16.

[0070] One arrangement for elements of system 16 is shown in FIGS. 1A and 1B, including a network interface 20, application platform 18, tenant data storage 22 for tenant data 23, system data storage 24 for system data 25 accessible to system 16 and possibly multiple tenants, program code 26 for implementing various functions of system 16, and a process space 28 for executing MTS system processes and tenant-specific processes, such as running applications as part of an application hosting service. Additional processes that may execute on system 16 include database indexing processes.

[0071] Several elements in the system shown in FIG. 1A include conventional, well-known elements that are explained only briefly here. For example, each user system 12 could include a desktop personal computer, workstation, laptop, PDA, cell phone, or any wireless access protocol (WAP) enabled device or any other computing device capable of interfacing directly or indirectly to the Internet or other network connection. The term "computing device" is

also referred to herein simply as a “computer”. User system **12** typically runs an HTTP client, e.g., a browsing program, such as Microsoft’s Internet Explorer browser, Netscape’s Navigator browser, Opera’s browser, or a WAP-enabled browser in the case of a cell phone, PDA or other wireless device, or the like, allowing a user (e.g., subscriber of the multi-tenant database system) of user system **12** to access, process and view information, pages and applications available to it from system **16** over network **14**. Each user system **12** also typically includes one or more user input devices, such as a keyboard, a mouse, trackball, touch pad, touch screen, pen or the like, for interacting with a graphical user interface (GUI) provided by the browser on a display (e.g., a monitor screen, LCD display, etc.) of the computing device in conjunction with pages, forms, applications and other information provided by system **16** or other systems or servers. For example, the user interface device can be used to access data and applications hosted by system **16**, and to perform searches on stored data, and otherwise allow a user to interact with various GUI pages that may be presented to a user. As discussed above, implementations are suitable for use with the Internet, although other networks can be used instead of or in addition to the Internet, such as an intranet, an extranet, a virtual private network (VPN), a non-TCP/IP based network, any LAN or WAN or the like.

[0072] According to one implementation, each user system **12** and all of its components are operator configurable using applications, such as a browser, including computer code run using a central processing unit such as an Intel Pentium® processor or the like. Similarly, system **16** (and additional instances of an MTS, where more than one is present) and all of its components might be operator configurable using application(s) including computer code to run using processor system **17**, which may be implemented to include a central processing unit, which may include an Intel Pentium® processor or the like, and/or multiple processor units. Non-transitory computer-readable media can have instructions stored thereon/in, that can be executed by or used to program a computing device to perform any of the methods of the implementations described herein. Computer program code **26** implementing instructions for operating and configuring system **16** to intercommunicate and to process web pages, applications and other data and media content as described herein is preferably downloadable and stored on a hard disk, but the entire program code, or portions thereof, may also be stored in any other volatile or non-volatile memory medium or device as is well known, such as a ROM or RAM, or provided on any media capable of storing program code, such as any type of rotating media including floppy disks, optical discs, digital versatile disk (DVD), compact disk (CD), microdrive, and magneto-optical disks, and magnetic or optical cards, nanosystems (including molecular memory ICs), or any other type of computer-readable medium or device suitable for storing instructions and/or data. Additionally, the entire program code, or portions thereof, may be transmitted and downloaded from a software source over a transmission medium, e.g., over the Internet, or from another server, as is well known, or transmitted over any other conventional network connection as is well known (e.g., extranet, VPN, LAN, etc.) using any communication medium and protocols (e.g., TCP/IP, HTTP, HTTPS, Ethernet, etc.) as are well known. It will also be appreciated that computer code for the disclosed implementations can be realized in any programming lan-

guage that can be executed on a client system and/or server or server system such as, for example, C, C++, HTML, any other markup language, Java™, JavaScript, ActiveX, any other scripting language, such as VBScript, and many other programming languages as are well known may be used. (Java™ is a trademark of Sun Microsystems, Inc.).

[0073] According to some implementations, each system **16** is configured to provide web pages, forms, applications, data and media content to user (client) systems **12** to support the access by user systems **12** as tenants of system **16**. As such, system **16** provides security mechanisms to keep each tenant’s data separate unless the data is shared. If more than one MTS is used, they may be located in close proximity to one another (e.g., in a server farm located in a single building or campus), or they may be distributed at locations remote from one another (e.g., one or more servers located in city A and one or more servers located in city B). As used herein, each MTS could include one or more logically and/or physically connected servers distributed locally or across one or more geographic locations. Additionally, the term “server” is meant to refer to a computing device or system, including processing hardware and process space(s), an associated storage medium such as a memory device or database, and, in some instances, a database application (e.g., OODBMS or RDBMS) as is well known in the art. It should also be understood that “server system” and “server” are often used interchangeably herein. Similarly, the database objects described herein can be implemented as single databases, a distributed database, a collection of distributed databases, a database with redundant online or offline backups or other redundancies, etc., and might include a distributed database or storage network and associated processing intelligence.

[0074] FIG. 1B shows a block diagram of an example of some implementations of elements of FIG. 1A and various possible interconnections between these elements. That is, FIG. 1B also illustrates environment **10**. However, in FIG. 1B elements of system **16** and various interconnections in some implementations are further illustrated. FIG. 1B shows that user system **12** may include processor system **12A**, memory system **12B**, input system **12C**, and output system **12D**. FIG. 1B shows network **14** and system **16**. FIG. 1B also shows that system **16** may include tenant data storage **22**, tenant data **23**, system data storage **24**, system data **25**, User Interface (UI) **30**, Application Program Interface (API) **32**, PL/SOQL **34**, save routines **36**, application setup mechanism **38**, applications servers **1001-100N**, system process space **102**, tenant process spaces **104**, tenant management process space **110**, tenant storage space **112**, user storage **114**, and application metadata **116**. In other implementations, environment **10** may not have the same elements as those listed above and/or may have other elements instead of, or in addition to, those listed above.

[0075] User system **12**, network **14**, system **16**, tenant data storage **22**, and system data storage **24** were discussed above in FIG. 1A. Regarding user system **12**, processor system **12A** may be any combination of one or more processors. Memory system **12B** may be any combination of one or more memory devices, short term, and/or long term memory. Input system **12C** may be any combination of input devices, such as one or more keyboards, mice, trackballs, scanners, cameras, and/or interfaces to networks. Output system **12D** may be any combination of output devices, such as one or more monitors, printers, and/or interfaces to networks. As

shown by FIG. 1B, system 16 may include a network interface 20 (of FIG. 1A) implemented as a set of HTTP application servers 100, an application platform 18, tenant data storage 22, and system data storage 24. Also shown is system process space 102, including individual tenant process spaces 104 and a tenant management process space 110. Each application server 100 may be configured to communicate with tenant data storage 22 and the tenant data 23 therein, and system data storage 24 and the system data 25 therein to serve requests of user systems 12. The tenant data 23 might be divided into individual tenant storage spaces 112, which can be either a physical arrangement and/or a logical arrangement of data. Within each tenant storage space 112, user storage 114 and application metadata 116 might be similarly allocated for each user. For example, a copy of a user's most recently used (MRU) items might be stored to user storage 114. Similarly, a copy of MRU items for an entire organization that is a tenant might be stored to tenant storage space 112. A UI 30 provides a user interface and an API 32 provides an application programmer interface to system 16 resident processes to users and/or developers at user systems 12. The tenant data and the system data may be stored in various databases, such as one or more Oracle □ databases.

[0076] Application platform 18 includes an application setup mechanism 38 that supports application developers' creation and management of applications, which may be saved as metadata into tenant data storage 22 by save routines 36 for execution by subscribers as one or more tenant process spaces 104 managed by tenant management process 110 for example. Invocations to such applications may be coded using PL/SOQL 34 that provides a programming language style interface extension to API 32. A detailed description of some PL/SOQL language implementations is discussed in commonly assigned U.S. Pat. No. 7,730,478, titled METHOD AND SYSTEM FOR ALLOWING ACCESS TO DEVELOPED APPLICATIONS VIA A MULTI-TENANT ON-DEMAND DATABASE SERVICE, by Craig Weissman, issued on Jun. 1, 2010, and hereby incorporated by reference in its entirety and for all purposes. Invocations to applications may be detected by one or more system processes, which manage retrieving application metadata 116 for the subscriber making the invocation and executing the metadata as an application in a virtual machine.

[0077] Each application server 100 may be communicably coupled to database systems, e.g., having access to system data 25 and tenant data 23, via a different network connection. For example, one application server 1001 might be coupled via the network 14 (e.g., the Internet), another application server 100N-1 might be coupled via a direct network link, and another application server 100N might be coupled by yet a different network connection. Transfer Control Protocol and Internet Protocol (TCP/IP) are typical protocols for communicating between application servers 100 and the database system. However, it will be apparent to one skilled in the art that other transport protocols may be used to optimize the system depending on the network interconnect used.

[0078] In certain implementations, each application server 100 is configured to handle requests for any user associated with any organization that is a tenant. Because it is desirable to be able to add and remove application servers from the server pool at any time for any reason, there is preferably no

server affinity for a user and/or organization to a specific application server 100. In one implementation, therefore, an interface system implementing a load balancing function (e.g., an F5 Big-IP load balancer) is communicably coupled between the application servers 100 and the user systems 12 to distribute requests to the application servers 100. In one implementation, the load balancer uses a least connections algorithm to route user requests to the application servers 100. Other examples of load balancing algorithms, such as round robin and observed response time, also can be used. For example, in certain implementations, three consecutive requests from the same user could hit three different application servers 100, and three requests from different users could hit the same application server 100. In this manner, by way of example, system 16 is multi-tenant, wherein system 16 handles storage of, and access to, different objects, data and applications across disparate users and organizations.

[0079] As an example of storage, one tenant might be a company that employs a sales force where each salesperson uses system 16 to manage their sales process. Thus, a user might maintain contact data, leads data, customer follow-up data, performance data, goals and progress data, etc., all applicable to that user's personal sales process (e.g., in tenant data storage 22). In an example of a MTS arrangement, since all of the data and the applications to access, view, modify, report, transmit, calculate, etc., can be maintained and accessed by a user system having nothing more than network access, the user can manage his or her sales efforts and cycles from any of many different user systems. For example, if a salesperson is visiting a customer and the customer has Internet access in their lobby, the salesperson can obtain critical updates as to that customer while waiting for the customer to arrive in the lobby.

[0080] While each user's data might be separate from other users' data regardless of the employers of each user, some data might be organization-wide data shared or accessible by a plurality of users or all of the users for a given organization that is a tenant. Thus, there might be some data structures managed by system 16 that are allocated at the tenant level while other data structures might be managed at the user level. Because an MTS might support multiple tenants including possible competitors, the MTS should have security protocols that keep data, applications, and application use separate. Also, because many tenants may opt for access to an MTS rather than maintain their own system, redundancy, up-time, and backup are additional functions that may be implemented in the MTS. In addition to user-specific data and tenant-specific data, system 16 might also maintain system level data usable by multiple tenants or other data. Such system level data might include industry reports, news, postings, and the like that are sharable among tenants.

[0081] In certain implementations, user systems 12 (which may be client systems) communicate with application servers 100 to request and update system-level and tenant-level data from system 16 that may involve sending one or more queries to tenant data storage 22 and/or system data storage 24. System 16 (e.g., an application server 100 in system 16) automatically generates one or more SQL statements (e.g., one or more SQL queries) that are designed to access the desired information. System data storage 24 may generate query plans to access the requested data from the database.

[0082] Each database can generally be viewed as a collection of objects, such as a set of logical tables, containing

data fitted into predefined categories. A “table” is one representation of a data object, and may be used herein to simplify the conceptual description of objects and custom objects according to some implementations. It should be understood that “table” and “object” may be used interchangeably herein. Each table generally contains one or more data categories logically arranged as columns or fields in a viewable schema. Each row or record of a table contains an instance of data for each category defined by the fields. For example, a CRM database may include a table that describes a customer with fields for basic contact information such as name, address, phone number, fax number, etc. Another table might describe a purchase order, including fields for information such as customer, product, sale price, date, etc. In some multi-tenant database systems, standard entity tables might be provided for use by all tenants. For CRM database applications, such standard entities might include tables for case, account, contact, lead, and opportunity data objects, each containing pre-defined fields. It should be understood that the word “entity” may also be used interchangeably herein with “object” and “table”.

[0083] In some multi-tenant database systems, tenants may be allowed to create and store custom objects, or they may be allowed to customize standard entities or objects, for example by creating custom fields for standard objects, including custom index fields. Commonly assigned U.S. Pat. No. 7,779,039, titled CUSTOM ENTITIES AND FIELDS IN A MULTI-TENANT DATABASE SYSTEM, by Weissman et al., issued on Aug. 17, 2010, and hereby incorporated by reference in its entirety and for all purposes, teaches systems and methods for creating custom objects as well as customizing standard objects in a multi-tenant database system. In certain implementations, for example, all custom entity data rows are stored in a single multi-tenant physical table, which may contain multiple logical tables per organization. It is transparent to customers that their multiple “tables” are in fact stored in one large table or that their data may be stored in the same table as the data of other customers.

[0084] FIG. 2A shows a system diagram illustrating an example of architectural components of an on-demand database service environment **200** according to some implementations. A client machine located in the cloud **204**, generally referring to one or more networks in combination, as described herein, may communicate with the on-demand database service environment via one or more edge routers **208** and **212**. A client machine can be any of the examples of user systems **12** described above. The edge routers may communicate with one or more core switches **220** and **224** via firewall **216**. The core switches may communicate with a load balancer **228**, which may distribute server load over different pods, such as the pods **240** and **244**. The pods **240** and **244**, which may each include one or more servers and/or other computing resources, may perform data processing and other operations used to provide on-demand services. Communication with the pods may be conducted via pod switches **232** and **236**. Components of the on-demand database service environment may communicate with a database storage **256** via a database firewall **248** and a database switch **252**.

[0085] As shown in FIGS. 2A and 2B, accessing an on-demand database service environment may involve communications transmitted among a variety of different hardware and/or software components. Further, the on-demand

database service environment **200** is a simplified representation of an actual on-demand database service environment. For example, while only one or two devices of each type are shown in FIGS. 2A and 2B, some implementations of an on-demand database service environment may include anywhere from one to many devices of each type. Also, the on-demand database service environment need not include each device shown in FIGS. 2A and 2B, or may include additional devices not shown in FIGS. 2A and 2B.

[0086] Moreover, one or more of the devices in the on-demand database service environment **200** may be implemented on the same physical device or on different hardware. Some devices may be implemented using hardware or a combination of hardware and software. Thus, terms such as “data processing apparatus,” “machine,” “server” and “device” as used herein are not limited to a single hardware device, but rather include any hardware and software configured to provide the described functionality.

[0087] The cloud **204** is intended to refer to a data network or plurality of data networks, often including the Internet. Client machines located in the cloud **204** may communicate with the on-demand database service environment to access services provided by the on-demand database service environment. For example, client machines may access the on-demand database service environment to retrieve, store, edit, and/or process information.

[0088] In some implementations, the edge routers **208** and **212** route packets between the cloud **204** and other components of the on-demand database service environment **200**. The edge routers **208** and **212** may employ the Border Gateway Protocol (BGP). The BGP is the core routing protocol of the Internet. The edge routers **208** and **212** may maintain a table of IP networks or ‘prefixes’, which designate network reachability among autonomous systems on the Internet.

[0089] In one or more implementations, the firewall **216** may protect the inner components of the on-demand database service environment **200** from Internet traffic. The firewall **216** may block, permit, or deny access to the inner components of the on-demand database service environment **200** based upon a set of rules and other criteria. The firewall **216** may act as one or more of a packet filter, an application gateway, a stateful filter, a proxy server, or any other type of firewall.

[0090] In some implementations, the core switches **220** and **224** are high-capacity switches that transfer packets within the on-demand database service environment **200**. The core switches **220** and **224** may be configured as network bridges that quickly route data between different components within the on-demand database service environment. In some implementations, the use of two or more core switches **220** and **224** may provide redundancy and/or reduced latency.

[0091] In some implementations, the pods **240** and **244** may perform the core data processing and service functions provided by the on-demand database service environment. Each pod may include various types of hardware and/or software computing resources. An example of the pod architecture is discussed in greater detail with reference to FIG. 2B.

[0092] In some implementations, communication between the pods **240** and **244** may be conducted via the pod switches **232** and **236**. The pod switches **232** and **236** may facilitate communication between the pods **240** and **244** and client

machines located in the cloud **204**, for example via core switches **220** and **224**. Also, the pod switches **232** and **236** may facilitate communication between the pods **240** and **244** and the database storage **256**.

[0093] In some implementations, the load balancer **228** may distribute workload between the pods **240** and **244**. Balancing the on-demand service requests between the pods may assist in improving the use of resources, increasing throughput, reducing response times, and/or reducing overhead. The load balancer **228** may include multilayer switches to analyze and forward traffic.

[0094] In some implementations, access to the database storage **256** may be guarded by a database firewall **248**. The database firewall **248** may act as a computer application firewall operating at the database application layer of a protocol stack. The database firewall **248** may protect the database storage **256** from application attacks such as structure query language (SQL) injection, database rootkits, and unauthorized information disclosure.

[0095] In some implementations, the database firewall **248** may include a host using one or more forms of reverse proxy services to proxy traffic before passing it to a gateway router. The database firewall **248** may inspect the contents of database traffic and block certain content or database requests. The database firewall **248** may work on the SQL application level atop the TCP/IP stack, managing applications' connection to the database or SQL management interfaces as well as intercepting and enforcing packets traveling to or from a database network or application interface.

[0096] In some implementations, communication with the database storage **256** may be conducted via the database switch **252**. The multi-tenant database storage **256** may include more than one hardware and/or software components for handling database queries. Accordingly, the database switch **252** may direct database queries transmitted by other components of the on-demand database service environment (e.g., the pods **240** and **244**) to the correct components within the database storage **256**.

[0097] In some implementations, the database storage **256** is an on-demand database system shared by many different organizations. The on-demand database system may employ a multi-tenant approach, a virtualized approach, or any other type of database approach. An on-demand database system is discussed in greater detail with reference to FIGS. 1A and 1B.

[0098] FIG. 2B shows a system diagram further illustrating an example of architectural components of an on-demand database service environment according to some implementations. The pod **244** may be used to render services to a user of the on-demand database service environment **200**. In some implementations, each pod may include a variety of servers and/or other systems. The pod **244** includes one or more content batch servers **264**, content search servers **268**, query servers **282**, file force servers **286**, access control system (ACS) servers **280**, batch servers **284**, and app servers **288**. Also, the pod **244** includes database instances **290**, quick file systems (QFS) **292**, and indexers **294**. In one or more implementations, some or all communication between the servers in the pod **244** may be transmitted via the switch **236**.

[0099] In some implementations, the app servers **288** may include a hardware and/or software framework dedicated to the execution of procedures (e.g., programs, routines,

scripts) for supporting the construction of applications provided by the on-demand database service environment **200** via the pod **244**. In some implementations, the hardware and/or software framework of an app server **288** is configured to execute operations of the services described herein, including performance of the blocks of methods described with reference to FIGS. 15-22. In alternative implementations, two or more app servers **288** may be included and cooperate to perform such methods, or one or more other servers described herein can be configured to perform the disclosed methods.

[0100] The content batch servers **264** may handle requests internal to the pod. These requests may be long-running and/or not tied to a particular customer. For example, the content batch servers **264** may handle requests related to log mining, cleanup work, and maintenance tasks.

[0101] The content search servers **268** may provide query and indexer functions. For example, the functions provided by the content search servers **268** may allow users to search through content stored in the on-demand database service environment.

[0102] The file force servers **286** may manage requests for information stored in the Fileforce storage **298**. The Fileforce storage **298** may store information such as documents, images, and basic large objects (BLOBs). By managing requests for information using the file force servers **286**, the image footprint on the database may be reduced.

[0103] The query servers **282** may be used to retrieve information from one or more file systems. For example, the query system **282** may receive requests for information from the app servers **288** and then transmit information queries to the NFS **296** located outside the pod.

[0104] The pod **244** may share a database instance **290** configured as a multi-tenant environment in which different organizations share access to the same database. Additionally, services rendered by the pod **244** may call upon various hardware and/or software resources. In some implementations, the ACS servers **280** may control access to data, hardware resources, or software resources.

[0105] In some implementations, the batch servers **284** may process batch jobs, which are used to run tasks at specified times. Thus, the batch servers **284** may transmit instructions to other servers, such as the app servers **288**, to trigger the batch jobs.

[0106] In some implementations, the QFS **292** may be an open source file system available from Sun Microsystems® of Santa Clara, Calif. The QFS may serve as a rapid-access file system for storing and accessing information available within the pod **244**. The QFS **292** may support some volume management capabilities, allowing many disks to be grouped together into a file system. File system metadata can be kept on a separate set of disks, which may be useful for streaming applications where long disk seeks cannot be tolerated. Thus, the QFS system may communicate with one or more content search servers **268** and/or indexers **294** to identify, retrieve, move, and/or update data stored in the network file systems **296** and/or other storage systems.

[0107] In some implementations, one or more query servers **282** may communicate with the NFS **296** to retrieve and/or update information stored outside of the pod **244**. The NFS **296** may allow servers located in the pod **244** to access information to access files over a network in a manner similar to how local storage is accessed.

[0108] In some implementations, queries from the query servers **222** may be transmitted to the NFS **296** via the load balancer **228**, which may distribute resource requests over various resources available in the on-demand database service environment. The NFS **296** may also communicate with the QFS **292** to update the information stored on the NFS **296** and/or to provide information to the QFS **292** for use by servers located within the pod **244**.

[0109] In some implementations, the pod may include one or more database instances **290**. The database instance **290** may transmit information to the QFS **292**. When information is transmitted to the QFS, it may be available for use by servers within the pod **244** without using an additional database call.

[0110] In some implementations, database information may be transmitted to the indexer **294**. Indexer **294** may provide an index of information available in the database **290** and/or QFS **292**. The index information may be provided to file force servers **286** and/or the QFS **292**.

[0111] III. Tracking Updates to a Record Stored in a Database

[0112] As multiple users might be able to change the data of a record, it can be useful for certain users to be notified when a record is updated. Also, even if a user does not have authority to change a record, the user still might want to know when there is an update to the record. For example, a vendor may negotiate a new price with a salesperson of company X, where the salesperson is a user associated with tenant Y. As part of creating a new invoice or for accounting purposes, the salesperson can change the price saved in the database. It may be important for co-workers to know that the price has changed. The salesperson could send an email to certain people, but this is onerous and the salesperson might not email all of the people who need to know or want to know. Accordingly, some implementations of the disclosed techniques can inform others (e.g., co-workers) who want to know about an update to a record automatically.

[0113] FIG. 3 shows a flowchart of an example of a method **300** for tracking updates to a record stored in a database system, performed in accordance with some implementations. Method **300** (and other methods described herein) may be implemented at least partially with multi-tenant database system **16**, e.g., by one or more processors configured to receive or retrieve information, process the information, store results, and transmit the results. In other implementations, method **300** may be implemented at least partially with a single tenant database system. In various implementations, blocks may be omitted, combined, or split into additional blocks for method **300**, as well as for other methods described herein.

[0114] In block **310**, the database system receives a request to update a first record. In one implementation, the request is received from a first user. For example, a user may be accessing a page associated with the first record, and may change a displayed field and hit save. In another implementation, the database system can automatically create the request. For instance, the database system can create the request in response to another event, e.g., a request to change a field could be sent periodically at a particular date and/or time of day, or a change to another field or object. The database system can obtain a new value based on other fields of a record and/or based on parameters in the system.

[0115] The request for the update of a field of a record is an example of an event associated with the first record for

which a feed tracked update may be created. In other implementations, the database system can identify other events besides updates to fields of a record. For example, an event can be a submission of approval to change a field. Such an event can also have an associated field (e.g., a field showing a status of whether a change has been submitted). Other examples of events can include creation of a record, deletion of a record, converting a record from one type to another (e.g., converting a lead to an opportunity), closing a record (e.g., a case type record), and potentially any other state change of a record—any of which could include a field change associated with the state change. Any of these events update the record whether by changing a field of the record, a state of the record, or some other characteristic or property of the record. In one implementation, a list of supported events for creating a feed tracked update can be maintained within the database system, e.g., at a server or in a database.

[0116] In block **320**, the database system writes new data to the first record. In one implementation, the new data may include a new value that replaces old data. For example, a field is updated with a new value. In another implementation, the new data can be a value for a field that did not contain data before. In yet another implementation, the new data could be a flag, e.g., for a status of the record, which can be stored as a field of the record.

[0117] In some implementations, a “field” can also include records, which are child objects of the first record in a parent-child hierarchy. A field can alternatively include a pointer to a child record. A child object itself can include further fields. Thus, if a field of a child object is updated with a new value, the parent record also can be considered to have a field changed. In one example, a field could be a list of related child objects, also called a related list.

[0118] In block **330**, a feed tracked update is generated about the update to the record. In one implementation, the feed tracked update is created in parts for assembling later into a display version. For example, event entries can be created and tracked in a first table, and changed field entries can be tracked in another table that is cross-referenced with the first table. More specifics of such implementations are provided later, e.g., with respect to FIG. 9A. In another implementation, the feed tracked update is automatically generated by the database system. The feed tracked update can convey in words that the first record has been updated and provide details about what was updated in the record and who performed the update. In some implementations, a feed tracked update is generated for only certain types of event and/or updates associated with the first record.

[0119] In one implementation, a tenant (e.g., through an administrator) can configure the database system to create (enable) feed tracked updates only for certain types of records. For example, an administrator can specify that records of designated types such as accounts and opportunities are enabled. When an update (or other event) is received for the enabled record type, then a feed tracked update would be generated. In another implementation, a tenant can also specify the fields of a record whose changes are to be tracked, and for which feed tracked updates are created. In one aspect, a maximum number of fields can be specified for tracking, and may include custom fields. In one implementation, the type of change can also be specified, for example, that the value change of a field is to be larger than a threshold (e.g., an absolute amount or a percentage change). In yet another implementation, a tenant can specify

which events are to cause a generation of a feed tracked update. Also, in one implementation, individual users can specify configurations specific to them, which can create custom feeds as described in more detail below.

[0120] In one implementation, changes to fields of a child object are not tracked to create feed tracked updates for the parent record. In another implementation, the changes to fields of a child object can be tracked to create feed tracked updates for the parent record. For example, a child object of the parent type can be specified for tracking, and certain fields of the child object can be specified for tracking. As another example, if the child object is of a type specified for tracking, then a tracked change for the child object is propagated to parent records of the child object.

[0121] In block 340, the feed tracked update is added to a feed for the first record. In one implementation, adding the feed tracked update to a feed can include adding events to a table (which may be specific to a record or be for all or a group of objects), where a display version of a feed tracked update can be generated dynamically and presented in a GUI as a feed item when a user requests a feed for the first record. In another implementation, a display version of a feed tracked update can be added when a record feed is stored and maintained for a record. As mentioned above, a feed may be maintained for only certain records. In one implementation, the feed of a record can be stored in the database associated with the record. For example, the feed can be stored as a field (e.g., as a child object) of the record. Such a field can store a pointer to the text to be displayed for the feed tracked update.

[0122] In some implementations, only the current feed tracked update (or other current feed item) may be kept or temporarily stored, e.g., in some temporary memory structure. For example, a feed tracked update for only a most recent change to any particular field is kept. In other implementations, many previous feed tracked updates may be kept in the feed. A time and/or date for each feed tracked update can be tracked. Herein, a feed of a record is also referred to as an entity feed, as a record is an instance of a particular entity object of the database.

[0123] In block 350, followers of the first record can be identified. A follower is a user following the first record, such as a subscriber to the feed of the first record. In one implementation, when a user requests a feed of a particular record, such an identification of block 350 can be omitted. In another implementation where a record feed is pushed to a user (e.g., as part of a news feed), then the user can be identified as a follower of the first record. Accordingly, this block can include the identification of records and other objects being followed by a particular user.

[0124] In one implementation, the database system can store a list of the followers for a particular record. In various implementations, the list can be stored with the first record or associated with the record using an identifier (e.g., a pointer) to retrieve the list. For example, the list can be stored in a field of the first record. In another implementation, a list of the records that a user is following is used. In one implementation, the database system can have a routine that runs for each user, where the routine polls the records in the list to determine if a new feed tracked update has been added to a feed of the record. In another implementation, the routine for the user can be running at least partially on a user device, which contacts the database to perform the polling.

[0125] In block 360, in one implementation, the feed tracked update can be stored in a table, as described in greater detail below. When the user opens a feed, an appropriate query is sent to one or more tables to retrieve updates to records, also described in greater detail below. In some implementations, the feed shows feed tracked updates in reverse chronological order. In one implementation, the feed tracked update is pushed to the feed of a user, e.g., by a routine that determines the followers for the record from a list associated with the record. In another implementation, the feed tracked update is pulled to a feed, e.g., by a user device. This pulling may occur when a user requests the feed, as occurs in block 370. Thus, these actions may occur in a different order. The creation of the feed for a pull may be a dynamic creation that identifies records being followed by the requesting user, generates the display version of relevant feed tracked updates from stored information (e.g., event and field change), and adds the feed tracked updates into the feed. A feed of feed tracked updates of records and other objects that a user is following is also generally referred to herein as a news feed, which can be a subset of a larger information feed in which other types of information updates appear, such as posts.

[0126] In yet another implementation, the feed tracked update could be sent as an email to the follower, instead of in a feed. In one implementation, email alerts for events can enable people to be emailed when certain events occur. In another implementation, emails can be sent when there are posts on a user profile and posts on entities to which the user subscribes. In one implementation, a user can turn on/off email alerts for all or some events. In an implementation, a user can specify what kind of feed tracked updates to receive about a record that the user is following. For example, a user can choose to only receive feed tracked updates about certain fields of a record that the user is following, and potentially about what kind of update was performed (e.g., a new value input into a specified field, or the creation of a new field).

[0127] In block 370, a follower can access his/her news feed to see the feed tracked update. In one implementation, the user has just one news feed for all of the records that the user is following. In one aspect, a user can access his/her own feed by selecting a particular tab or other object on a page of an interface to the database system. Once selected the feed can be provided as a list, e.g., with an identifier (e.g., a time) or including some or all of the text of the feed tracked update. In another implementation, the user can specify how the feed tracked updates are to be displayed and/or sent to the user. For example, a user can specify a font for the text, a location of where the feed can be selected and displayed, amount of text to be displayed, and other text or symbols to be displayed (e.g., importance flags).

[0128] FIG. 4 shows a block diagram of an example of components of a database system configuration 400 performing a method for tracking an update to a record according to some implementations. Database system configuration 400 can perform implementations of method 300, as well as implementations of other methods described herein.

[0129] A first user 405 sends a request 1 to update record 425 in database system 416. Although an update request is described, other events that are being tracked are equally applicable. In various implementations, the request 1 can be sent via a user interface (e.g., 30 of FIG. 1B) or an application program interface (e.g., API 32). An I/O port 420

can accommodate the signals of request 1 via any input interface, and send the signals to one or more processors 417. The processor 417 can analyze the request and determine operations to be performed. Herein, any reference to a processor 417 can refer to a specific processor or any set of processors in database system 416, which can be collectively referred to as processor 417.

[0130] Processor 417 can determine an identifier for record 425, and send commands with the new data 2 of the request to record database 412 to update record 425. In one implementation, record database 412 is where tenant storage space 112 of FIG. 1B is located. The request 1 and new data commands 2 can be encapsulated in a single write transaction sent to record database 412. In one implementation, multiple changes to records in the database can be made in a single write transaction.

[0131] Processor 417 can also analyze request 1 to determine whether a feed tracked update is to be created, which at this point may include determining whether the event (e.g., a change to a particular field) is to be tracked. This determination can be based on an interaction (i.e., an exchange of data) with record database 412 and/or other databases, or based on information stored locally (e.g., in cache or RAM) at processor 417. In one implementation, a list of record types that are being tracked can be stored. The list may be different for each tenant, e.g., as each tenant may configure the database system to its own specifications. Thus, if the record 425 is of a type not being tracked, then the determination of whether to create a feed tracked update can stop there.

[0132] The same list or a second list (which can be stored in a same location or a different location) can also include the fields and/or events that are tracked for the record types in the first list. This list can be searched to determine if the event is being tracked. A list may also contain information having the granularity of listing specific records that are to be tracked (e.g., if a tenant can specify the particular records to be tracked, as opposed to just type).

[0133] As an example, processor 417 may obtain an identifier associated with record 425 (e.g., obtained from request 1 or database 412), potentially along with a tenant identifier, and cross-reference the identifier with a list of records for which feed tracked updates are to be created. Specifically, the record identifier can be used to determine the record type and a list of tracked types can be searched for a match. The specific record may also be checked if such individual record tracking was enabled. The name of the field to be changed can also be used to search a list of tracking-enabled fields. Other criteria besides field and events can be used to determine whether a feed tracked update is created, e.g., type of change in the field. If a feed tracked update is to be generated, processor 417 can then generate the feed tracked update.

[0134] In some implementations, a feed tracked update is created dynamically when a feed (e.g., the entity feed of record 425) is requested. Thus, in one implementation, a feed tracked update can be created when a user requests the entity feed for record 425. In this implementation, the feed tracked update may be created (e.g., assembled), including re-created, each time the entity feed is to be displayed to any user. In one implementation, one or more event history tables can keep track of previous events so that the feed tracked update can be re-created.

[0135] In another implementation, a feed tracked update can be created at the time the event occurs, and the feed tracked update can be added to a list of feed items. The list of feed items may be specific to record 425, or may be an aggregate of feed items including feed items for many records. Such an aggregate list can include a record identifier so that the feed items for the entity feed of record 425 can be easily retrieved. For example, after the feed tracked update has been generated, processor 417 can add the new feed tracked update 3 to a feed of record 425. As mentioned above, in one implementation, the feed can be stored in a field (e.g., as a child object) of record 425. In another implementation, the feed can be stored in another location or in another database, but with a link (e.g., a connecting identifier) to record 425. The feed can be organized in various ways, e.g., as a linked list, an array, or other data structure.

[0136] A second user 430 can access the new feed tracked update 3 in various ways. In one implementation, second user 430 can send a request 4 for the record feed. For example, second user 430 can access a home page (detail page) of the record 425 (e.g., with a query or by browsing), and the feed can be obtained through a tab, button, or other activation object on the page. The feed can be displayed on the screen or downloaded.

[0137] In another implementation, processor 417 can add the new feed tracked update 5 to a feed (e.g., a news feed) of a user that is following record 425. In one implementation, processor 417 can determine each of the followers of record 425 by accessing a list of the users that have been registered as followers. This determination can be done for each new event (e.g., update 1). In another implementation, processor 417 can poll (e.g., with a query) the records that second user 430 is following to determine when new feed tracked updates (or other feed items) are available. Processor 417 can use a follower profile 435 of second user 430 that can contain a list of the records that the second user 430 is following. Such a list can be contained in other parts of the database as well. Second user 430 can then send a request 6 to his/her profile 435 to obtain a feed, which contains the new feed tracked update. The user's profile 435 can be stored in a profile database 414, which can be the same or different than database 412.

[0138] In some implementations, a user can define a news feed to include new feed tracked updates from various records, which may be limited to a maximum number. In one implementation, each user has one news feed. In another implementation, the follower profile 435 can include the specifications of each of the records to be followed (with the criteria for what feed tracked updates are to be provided and how they are displayed), as well as the feed.

[0139] Some implementations can provide various types of record (entity) feeds. Entity Feeds can exist for record types like account, opportunity, case, and contact. An entity feed can tell a user about the actions that people have taken on that particular record or on one its related records. The entity feed can include who made the action, which field was changed, and the old and new values. In one implementation, entity feeds can exist on all supported records as a list that is linked to the specific record. For example, a feed could be stored in a field that allows lists (e.g., linked lists) or as a child object.

[0140] IV. Tracking Actions of a User

[0141] In addition to knowing about events associated with a particular record, it can be helpful for a user to know what a particular user is doing. In particular, it might be nice to know what the user is doing without the user having to generate the feed tracked update (e.g., a user submitting a synopsis of what the user has done). Accordingly, implementations can automatically track actions of a user that trigger events, and feed tracked updates can be generated for certain events.

[0142] FIG. 5 shows a flowchart of an example of a method 500 for tracking actions of a user of a database system, performed in accordance with some implementations. Method 500 may be performed in addition to method 300. The operations of method 300, including order of blocks, can be performed in conjunction with method 500 and other methods described herein. Thus, a feed can be composed of changes to a record and actions of users.

[0143] In block 510, a database system (e.g., 16 of FIGS. 1A and 1B) identifies an action of a first user. In one implementation, the action triggers an event, and the event is identified. For example, the action of a user requesting an update to a record can be identified, where the event is receiving a request or is the resulting update of a record. The action may thus be defined by the resulting event. In another implementation, only certain types of actions (events) are identified. Which actions are identified can be set as a default or can be configurable by a tenant, or even configurable at a user level. In this way, processing effort can be reduced since only some actions are identified.

[0144] In block 520, it is determined whether the event qualifies for a feed tracked update. In one implementation, a predefined list of events (e.g., as mentioned herein) can be created so that only certain actions are identified. In one implementation, an administrator (or other user) of a tenant can specify the type of actions (events) for which a feed tracked update is to be generated. This block may also be performed for method 300.

[0145] In block 530, a feed tracked update is generated about the action. In an example where the action is an update of a record, the feed tracked update can be similar or the same as the feed tracked update created for the record. The description can be altered though to focus on the user as opposed to the record. For example, "John D. has closed a new opportunity for account XYZ" as opposed to "an opportunity has been closed for account XYZ."

[0146] In block 540, the feed tracked update is added to a profile feed of the first user when, e.g., the user clicks on a tab to open a page in a browser program displaying the feed. In one implementation, a feed for a particular user can be accessed on a page of the user's profile, in a similar manner as a record feed can be accessed on a detail page of the record. In another implementation, the first user may not have a profile feed and the feed tracked update may just be stored temporarily before proceeding. A profile feed of a user can be stored associated with the user's profile. This profile feed can be added to a news feed of another user.

[0147] In block 550, followers of the first user are identified. In one implementation, a user can specify which type of actions other users can follow. Similarly, in one implementation, a follower can select what actions by a user the follower wants to follow. In an implementation where different followers follow different types of actions, which users are followers of that user and the particular action can

be identified, e.g., using various lists that track what actions and criteria are being followed by a particular user. In various implementations, the followers of the first user can be identified in a similar manner as followers of a record, as described above for block 350.

[0148] In block 560, the feed tracked update is added to a news feed of each follower of the first user when, e.g., the follower clicks on a tab to open a page displaying the news feed. The feed tracked update can be added in a similar manner as the feed items for a record feed. The news feed can contain feed tracked updates both about users and records. In another implementation, a user can specify what kind of feed tracked updates to receive about a user that the user is following. For example, a user could specify feed tracked updates with particular keywords, of certain types of records, of records owned or created by certain users, particular fields, and other criteria as mentioned herein.

[0149] In block 570, a follower accesses the news feed and sees the feed tracked update. In one implementation, the user has just one news feed for all of the records that the user is following. In another implementation, a user can access his/her own feed (i.e. feed about his/her own actions) by selecting a particular tab or other object on a page of an interface to the database system. Thus, a feed can include feed tracked updates about what other users are doing in the database system. When a user becomes aware of a relevant action of another user, the user can contact the co-worker, thereby fostering teamwork.

[0150] V. Generation of a Feed Tracked Update

[0151] As described above, some implementations can generate text describing events (e.g., updates) that have occurred for a record and actions by a user that trigger an event. A database system can be configured to generate the feed tracked updates for various events in various ways.

[0152] A. Which Events to Generate a Feed Tracked Update

[0153] In a database system, there are various events that can be detected. However, the operator of the database system and/or a tenant may not want to detect every possible event as this could be costly with regards to performance. Accordingly, the operator and/or the tenant can configure the database system to only detect certain events. For example, an update of a record may be an event that is to be detected.

[0154] Out of the events that are detected, a tenant (including a specific user of the tenant) may not want a feed tracked update about each detected event. For example, all updates to a record may be identified at a first level. Then, based on specifications of an administrator and/or a specific user of a tenant, another level of inquiry can be made as to whether a feed tracked update is to be generated about the detected event. For example, the events that qualify for a feed tracked update can be restricted to changes for only certain fields of the record, and can differ depending on which user is receiving the feed. In one implementation, a database system can track whether an event qualifies for a feed tracked update for any user, and once the feed tracked update is generated, it can be determined who is to receive the feed tracked update.

[0155] Supported events (events for which a feed tracked update is generated) can include actions for standard fields, custom fields, and standard related lists. Regarding standard fields, for the entity feed and the profile feed, a standard field update can trigger a feed tracked update to be presented in that feed. In one implementation, which standard field can

create a feed tracked update can be set by an administrator to be the same for every user. In another implementation, a user can set which standard fields create a feed tracked update for that user's news feed. Custom fields can be treated the same or differently than standard fields.

[0156] The generation of a feed item can also depend on a relationship of an object to other objects (e.g., parent-child relationships). For example, if a child object is updated, a feed tracked update may be written to a feed of a parent of the child object. The level of relationship can be configured, e.g., only 1 level of separation (i.e. no grandparent-grand-child relationship). Also, in one implementation, a feed tracked update is generated only for objects above the objects being updated, i.e., a feed tracked update is not written for a child when the parent is updated.

[0157] In some implementations, for related lists of a record, a feed tracked update is written to its parent record (1 level only) when the related list item is added, and not when the list item is changed or deleted. For example: user A added a new opportunity XYZ for account ABC. In this manner, entity feeds can be controlled so as not to be cluttered with feed tracked updates about changes to their related items. Any changes to the related list item can be tracked on their own entity feed, if that related list item has a feed on it. In this implementation, if a user wants to see a feed of the related list item then the user can subscribe to it. Such a subscription might be when a user cares about a specific opportunity related to a specific account. A user can also browse to that object's entity feed. Other implementations can create a feed tracked update when a related entity is changed or deleted.

[0158] In one implementation, an administrator (of the system or of a specific tenant) can define which events of which related objects are to have feed tracked updates written about them in a parent record. In another implementation, a user can define which related object events to show. In one implementation, there are two types of related lists of related objects: first class lookup and second class lookup. Each of the records in the related lists can have a different rule for whether a feed tracked update is generated for a parent record. Each of these related lists can be composed as custom related lists. In various implementations, a custom related list can be composed of custom objects; the lists can contain a variety of records or items (e.g., not restricted to a particular type of record or item), and can be displayed in a customized manner.

[0159] In one implementation, a first class lookup contains records of a child record that can exist by itself. For example, the contacts on an account exist as a separate record and also as a child record of the account. In another implementation, a record in a first class lookup can have its own feed, which can be displayed on its detail page.

[0160] In one implementation, a second class lookup can have line items existing only in the context of their parent record (e.g., activities on an opportunity, contact roles on opportunity/contact). In one implementation, the line items are not objects themselves, and thus there is no detail page, and no place to put a feed. In another implementation, a change in a second class lookup can be reported on the feed of the parent.

[0161] Some implementations can also create feed tracked updates for dependent field changes. A dependent field change is a field that changes value when another field changes, and thus the field has a value that is dependent on

the value of the other field. For example, a dependent field might be a sum (or other formula) that totals values in other fields, and thus the dependent field would change when one of the fields being summed changes. Accordingly, in one implementation, a change in one field could create feed tracked updates for multiple fields. In other implementations, feed tracked updates are not created for dependent fields.

[0162] B. How the Feed Tracked Update is Generated

[0163] After it is determined that a feed tracked update is going to be generated, some implementations can also determine how the feed tracked update is generated. In one implementation, different methods can be used for different events, e.g., in a similar fashion as for the configurability of which events feed tracked updates are generated. A feed tracked update can also include a description of multiple events (e.g., john changed the account status and amount).

[0164] In one implementation, the feed tracked update is a grammatical sentence, thereby being easily understandable by a person. In another implementation, the feed tracked update provides detailed information about the update. In various examples, an old value and new value for a field may be included in the feed tracked update, an action for the update may be provided (e.g., submitted for approval), and the names of particular users that are responsible for replying or acting on the feed tracked update may be also provided. The feed tracked update can also have a level of importance based on settings chosen by the administrator, a particular user requesting an update, or by a following user who is to receive the feed tracked update, which fields is updated, a percentage of the change in a field, the type of event, or any combination of these factors.

[0165] The system may have a set of heuristics for creating a feed tracked update from the event (e.g., a request to update). For example, the subject may be the user, the record, or a field being added or changed. The verb can be based on the action requested by the user, which can be selected from a list of verbs (which may be provided as defaults or input by an administrator of a tenant). In one implementation, feed tracked updates can be generic containers with formatting restrictions,

[0166] As an example of a feed tracked update for a creation of a new record, "Mark Abramowitz created a new Opportunity for IBM-20,000 laptops with Amount as \$3.5M and Sam Palmisano as Decision Maker." This event can be posted to the profile feed for Mark Abramowitz and the entity feed for record of Opportunity for IBM-20,000 laptops. The pattern can be given by (AgentFullName) created a new (ObjectName)(RecordName) with [(FieldName) as (FieldValue) [, and]]*[[added/changed/removed] (RelatedListRecordName) [as/to/as] (RelatedListRecordValue) [./and]]*. Similar patterns can be formed for a changed field (standard or custom) and an added child record to a related list.

[0167] VI. Tracking Commentary from or about a User

[0168] Some implementations can also have a user submit text, instead of the database system generating a feed tracked update. As the text is submitted as part or all of a message by a user, the text can be about any topic. Thus, more information than just actions of a user and events of a record can be conveyed. In one implementation, the messages can be used to ask a question about a particular record, and users following the record can provide comments and responses.

[0169] FIG. 6 shows a flowchart of an example of a method 600 for creating a news feed from messages created by a user about a record or another user, performed in accordance with some implementations. In one implementation, method 600 can be combined with methods 300 and 500. In one aspect, a message can be associated with the first user when the first user creates the message (e.g., a post or comment about a record or another user). In another aspect, a message can be associated with the first user when the message is about the first user (e.g., posted by another user on the first user's profile feed).

[0170] In block 610, the database system receives a message (e.g., a post or status update) associated with a first user. The message (e.g., a post or status update) can contain text and/or multimedia content submitted by another user or by the first user. In one implementation, a post is for a section of the first user's profile page where any user can add a post, and where multiple posts can exist. Thus, a post can appear on the first user's profile page and can be viewed when the first user's profile is visited. For a message about a record, the post can appear on a detail page of a record. Note the message can appear in other feeds as well. In another implementation, a status update about the first user can only be added by the first user. In one implementation, a user can only have one status message.

[0171] In block 620, the message is added to a table, as described in greater detail below. When the feed is opened, a query filters one or more tables to identify the first user, identify other persons that the user is following, and retrieve the message. Messages and record updates are presented in a combined list as the feed. In this way, in one implementation, the message can be added to a profile feed of the first user, which is associated (e.g., as a related list) with the first user's profile. In one implementation, the posts are listed indefinitely. In another implementation, only the most recent posts (e.g., last 50) are kept in the profile feed. Such implementations can also be employed with feed tracked updates. In yet another implementation, the message can be added to a profile of the user adding the message.

[0172] In block 630, the database system identifies followers of the first user. In one implementation, the database system can identify the followers as described above for method 500. In various implementations, a follower can select to follow a feed about the actions of the first user, messages about the first user, or both (potentially in a same feed).

[0173] In block 640, the message is added to a news feed of each follower. In one implementation, the message is only added to a news feed of a particular follower if the message matches some criteria, e.g., the message includes a particular keyword or other criteria. In another implementation, a message can be deleted by the user who created the message. In one implementation, once deleted by the author, the message is deleted from all feeds to which the message had been added.

[0174] In block 650, the follower accesses a news feed and sees the message. For example, the follower can access a news feed on the follower's own profile page. As another example, the follower can have a news feed sent to his/her own desktop without having to first go to a home page.

[0175] In block 660, the database system receives a comment about the message. The database system can add the comment to a feed of the same first user, much as the original message was added. In one implementation, the comment

can also be added to a feed of a second user who added the comment. In one implementation, users can also reply to the comment. In another implementation, users can add comments to a feed tracked update, and further comments can be associated with the feed tracked update. In yet another implementation, making a comment or message is not an action to which a feed tracked update is created. Thus, the message may be the only feed item created from such an action.

[0176] In one implementation, if a feed tracked update or post is deleted, its corresponding comments are deleted as well. In another implementation, new comments on a feed tracked update or post do not update the feed tracked update timestamp. Also, the feed tracked update or post can continue to be shown in a feed (profile feed, record feed, or news feed) if it has had a comment within a specified timeframe (e.g., within the last week). Otherwise, the feed tracked update or post can be removed in an implementation.

[0177] In some implementations, all or most feed tracked updates can be commented on. In other implementations, feed tracked updates for certain records (e.g., cases or ideas) are not commentable. In various implementations, comments can be made for any one or more records of opportunities, accounts, contacts, leads, and custom objects.

[0178] In block 670, the comment is added to a news feed of each follower. In one implementation, a user can make the comment within the user's news feed. Such a comment can propagate to the appropriate profile feed or record feed, and then to the news feeds of the following users. Thus, feeds can include what people are saying, as well as what they are doing. In one aspect, feeds are a way to stay up-to-date (e.g., on users, opportunities, etc.) as well as an opportunity to reach out to co-workers/partners and engage them around common goals.

[0179] In some implementations, users can rate feed tracked updates or messages (including comments). A user can choose to prioritize a display of a feed so that higher rated feed items show up higher on a display. For example, in an implementation where comments are answers to a specific question, users can rate the different status posts so that a best answer can be identified. As another example, users are able to quickly identify feed items that are most important as those feed items can be displayed at a top of a list. The order of the feed items can be based on an importance level (which can be determined by the database system using various factors, some of which are mentioned herein) and based on a rating from users. In one implementation, the rating is on a scale that includes at least 3 values. In another implementation, the rating is based on a binary scale.

[0180] Besides a profile for a user, a group can also be created. In various implementations, the group can be created based on certain attributes that are common to the users, can be created by inviting users, and/or can be created by receiving requests to join from a user. In one implementation, a group feed can be created, with messages being added to the group feed when someone submits a message to the group as a whole through a suitable user interface. For example, a group page may have a group feed or a section within the feed for posts, and a user can submit a post through a publisher component in the user interface by clicking on a "Share" or similar button. In another implementation, a message can be added to a group feed when the message is submitted about any one of the members. Also,

a group feed can include feed tracked updates about actions of the group as a whole (e.g., when an administrator changes data in a group profile or a record owned by the group), or about actions of an individual member.

[0181] FIG. 7 shows an example of a group feed on a group page according to some implementations. As shown, a feed item **710** shows that a user has posted a document to the group object. The text “Bill Bauer has posted the document Competitive Insights” can be generated by the database system in a similar manner as feed tracked updates about a record being changed. A feed item **720** shows a post to the group, along with comments **730** from Ella Johnson, James Saxon, Mary Moore and Bill Bauer.

[0182] FIG. 8 shows an example of a record feed containing a feed tracked update, post, and comments according to some implementations. Feed item **810** shows a feed tracked update based on the event of submitting a discount for approval. Other feed items show posts, e.g., from Bill Bauer, that are made to the record and comments, e.g., from Erica Law and Jake Rapp, that are made on the posts.

[0183] VII. Infrastructure for a Feed

[0184] A. Tables Used to Create a Feed

[0185] FIG. 9A shows an example of a plurality of feed tracked update tables that may be used in tracking events and creating feeds according to some implementations. The tables of FIG. 9A may have entries added, or potentially removed, as part of tracking events in the database from which feed items are created or that correspond to feed items. In one implementation, each tenant has its own set of tables that are created based on criteria provided by the tenant.

[0186] An event history table **910** can provide a feed tracked update of events from which feed items are created. In one aspect, the events are for objects that are being tracked. Thus, table **910** can store and change feed tracked updates for feeds, and the changes can be persisted. In various implementations, event history table **910** can have columns of event ID **911**, object ID **912** (also called parent ID), and created by ID **913**. The event ID **911** can uniquely identify a particular event and can start at **1** (or other number or value).

[0187] Each new event can be added chronologically with a new event ID, which may be incremented in order. An object ID **912** can be used to track which record or user’s profile is being changed. For example, the object ID can correspond to the record whose field is being changed or the user whose feed is receiving a post. The created by ID **913** can track the user who is performing the action that results in the event, e.g., the user that is changing the field or that is posting a message to the profile of another user.

[0188] In some other implementations, event history table **910** can have one or more of the following variables with certain attributes: ORGANIZATION_ID being CHAR(15 BYTE), FEEDS_ENTITY_HIFEED_TRACKED_UPDATE_ID being CHAR(15 BYTE), PARENT_ID being CHAR(15 BYTE), CREATED_BY being CHAR(15 BYTE), CREATED_DATE being a variable of type DATE, DIVISION being a NUMBER, KEY_PREFIX being CHAR(3 BYTE), and DELETED being CHAR(1 BYTE). The parent ID can provide an ID of a parent object in case the change is promulgated to the parent. The key prefix can provide a key that is unique to a group of records, e.g., custom records (objects). The deleted variable can indicate that the feed items for the event are deleted, and thus the feed

items are not generated. In one implementation, the variables for each event entry or any entry in any of the tables may not be nullable. In another implementation, all entries in the event history table **910** are used to create feed items for only one object, as specified by the object ID **912**. For example, one feed tracked update cannot communicate updates on two records, such as updates of an account field and an opportunity field.

[0189] In one implementation, a name of an event can also be stored in table **910**. In one implementation, a tenant can specify events that they want tracked. In an implementation, event history table **910** can include the name of the field that changed (e.g., old and new values). In another implementation, the name of the field, and the values, are stored in a separate table. Other information about an event (e.g., text of comment, feed tracked update, post or status update) can be stored in event history table **910**, or in other tables, as is now described.

[0190] A field change table **920** can provide a feed tracked update of the changes to the fields. The columns of table **920** can include an event ID **921** (which correlates to the event ID **911**), an old value **922** for the field, and the new value **923** for the field. In one implementation, if an event changes more than one field value, then there can be an entry for each field changed. As shown, event ID **921** has two entries for event **E37**.

[0191] In some other implementations, field change table **920** can have one or more of the following variables with certain attributes: ORGANIZATION_ID being CHAR(15 BYTE), FEEDS_ENTITY_HIFEED_TRACKED_UPDATE_FIELDS_ID being CHAR(15 BYTE) and identifying each entry, FEEDS_ENTITY_HIFEED_TRACKED_UPDATE_ID being CHAR(15 BYTE), FIELD_KEY being VARCHAR2(120 BYTE), DATA_TYPE being CHAR(1 BYTE), OLDVAL_STRING VARCHAR2 being (765 BYTE), NEWVAL_STRING being VARCHAR2(765 BYTE), OLDVAL_FIRST_NAME being VARCHAR2(765 BYTE), NEWVAL_FIRST_NAME being VARCHAR2(765 BYTE), OLDVAL_LAST_NAME being VARCHAR2(765 BYTE), NEWVAL_LAST_NAME being VARCHAR2(765 BYTE), OLDVAL_NUMBER being NUMBER, NEWVAL_NUMBER being NUMBER, OLDVAL_DATE being DATE, NEWVAL_DATE being DATE, and DELETED being CHAR(1 BYTE). In one implementation, one or more of the variables for each entry in any of the tables may be nullable.

[0192] In one implementation, the data type variable (and/or other variables) is a non-API-insertable field. In another implementation, variable values can be derived from the record whose field is being changed. Certain values can be transferred into typed columns old/new value string, old/new value number or old/new value date depending upon the derived values. In another implementation, there can exist a data type for capturing add/deletes for child objects. The child ID can be tracked in the foreign-key column of the record. In yet another implementation, if the field name is pointing to a field in the parent entity, a field level security (FLS) can be used when a user attempts to a view a relevant feed item. Herein, security levels for objects and fields are also called access checks and determinations of authorization. In one aspect, the access can be for create, read, write, update, or delete of objects.

[0193] In one implementation, the field name (or key) can be either a field name of the entity or one of the values in a

separate list. For example, changes that do not involve the update of an existing field (e.g., a close or open) can have a field name specified in an enumerated list. This enumerated list can store “special” field name sentinel values for non-update actions that a tenant wants to track. In one aspect, the API just surfaces these values and the caller has to check the enumerated values to see if it is a special field name.

[0194] A comment table **930** can provide a feed tracked update of the comments made regarding an event, e.g., a comment on a post or a change of a field value. The columns of table **930** can include an event ID **921** (which correlates to the event ID **911**), the comment column **932** that stores the text of the comment, and the time/date **933** of the comment. In one implementation, there can be multiple comments for each event. As shown, event ID **921** has two entries for event **E37**.

[0195] In some other implementations, comment table **930** can have one or more of the following variables with certain attributes: ORGANIZATION_ID being CHAR(15 BYTE), FEEDS_COMMENTS_ID being CHAR(15 BYTE) and uniquely identifying each comment, PARENT_ID being CHAR(15 BYTE), CREATED_BY being CHAR(15 BYTE), CREATED_DATE being DATE, COMMENTS being VARCHAR2(420 BYTE), and DELETED being CHAR(1 BYTE).

[0196] A user subscription table **940** can provide a list of the objects being followed (subscribed to) by a user. In one implementation, each entry has a user ID **941** of the user doing the following and one object ID **942** corresponding to the object being followed. In one implementation, the object being followed can be a record or a user. As shown, the user with ID **U819** is following object IDs **0615** and **0489**. If user **U819** is following other objects, then additional entries may exist for user **U819**. Also as shown, user **U719** is also following object **O615**. The user subscription table **940** can be updated when a user adds or deletes an object that is being followed.

[0197] In some other implementations, user subscription table **940** can be composed of two tables (one for records being followed and one for users being followed). One table can have one or more of the following variables with certain attributes: ORGANIZATION_ID being CHAR(15 BYTE), ENTITY_SUBSCRIPTION_ID being CHAR(15 BYTE), PARENT_ID being CHAR(15 BYTE), CREATED_BY being CHAR(15 BYTE), CREATED_DATE being DATE, and DELETED being CHAR(1 BYTE). Another table can have one or more of the following variables with certain attributes: ORGANIZATION_ID being CHAR(15 BYTE), USER_SUBSCRIPTIONS_ID being CHAR(15 BYTE), USER_ID being CHAR(15 BYTE), CREATED_BY being CHAR(15 BYTE), and CREATED_DATE being DATE.

[0198] In one implementation, regarding a profile feed and a news feed, these are read-only views on the event history table **910** specialized for these feed types. Conceptually the news feed can be a semi-join between the user subscription table **940** and the event history table **910** on the object IDs **912** and **942** for the user. In one aspect, these entities can have polymorphic parents and can be subject to a number of restrictions detailed herein, e.g., to limit the cost of sharing checks.

[0199] In one implementation, entity feeds are modeled in the API as a feed associate entity (e.g., AccountFeed, CaseFeed, etc). A feed associate entity includes information composed of events (e.g., event IDs) for only one particular

record type. Such a list can limit the query (and sharing checks) to a specific record type. In one aspect, this structuring of the entity feeds can make the query run faster. For example, a request for a feed of a particular account can include the record type of account. In one implementation, an account feed table can then be searched, where the table has account record IDs and corresponding event IDs or pointers to particular event entries in event history table **910**. Since the account feed table only contains some of the records (not all), the query can run faster.

[0200] In one implementation, there may be objects with no events listed in the event history table **910**, even though the record is being tracked. In this case, the database service can return a result indicating that no feed items exist.

[0201] In another implementation, tables can also exist for audit tracking, e.g., to examine that operations of the system (e.g., access checks) are performing accurately. In one implementation, audit change-event history tables can be persisted (e.g., in bulk) synchronously in the same transaction as feed events are added to event history table **910**. In another implementation, entries to the two sets of table can be persisted in asynchronous manner (e.g., by forking a bulk update into a separate java thread). In one aspect, some updates to any of the tables can get lost if the instance of the table goes down while the update has not yet finished. This asynchronous manner can limit an impact performance on save operations. In some implementations, a field “persistence type” (tri state: AUDIT, FEEDS or BOTH) can be added to capture user preferences, as opposed to being hard coded.

[0202] B. Feed Item

[0203] A feed item can represent an individual field change of a record, creation and deletion of a record, or other events being tracked for a record or a user. In one implementation, all of the feed items in a single transaction (event) can be grouped together and have the same event ID. A single transaction relates to the operations that can be performed in a single communication with the database. In another implementation where a feed is an object of the database, a feed item can be a child of a profile feed, news feed, or entity feed. If a feed item is added to multiple feeds, the feed item can be replicated as a child of each feed to which the feed item is added.

[0204] In one implementation, a feed item is visible only when its parent feed is visible, which can be the same as needing read access on the feed’s parent (which can be by the type of record or by a specific record). The feed item’s field may be only visible when allowed under field-level security (FLS). Unfortunately, this can mean that the parent feed may be visible, but the child may not be because of FLS. Such access rules are described in more detail below. In one implementation, a feed item can be read-only. In this implementation, after being created, the feed item cannot be changed.

[0205] In multi-currency organizations, a feed item can have an extra currency code field. This field can give the currency code for the currency value in this field. In one aspect, the value is undefined when the data type is anything other than currency.

[0206] C. Feed Comment

[0207] In some implementations, a comment exists as an item that depends from feed tracked updates, posts, status updates, and other items that are independent of each other. Thus, a feed comment object can exist as a child object of

a feed item object. For example, comment table **930** can be considered a child table of event history table **910**. In one implementation, a feed comment can be a child of a profile feed, news feed, or entity feed that is separate from other feed items.

[0208] In various implementations, a feed comment can have various permissions for the following actions. For read permission, a feed comment can be visible if the parent feed is visible. For create permission, if a user has access to the feed (which can be tracked by the ID of the parent feed), the user can add a comment. For delete, only a user with modify all data permission or a user who added the comment can delete the comment. Also delete permission can involve access on the parent feed. An update of a comment can be restricted, and thus not be allowed.

[0209] In one implementation, regarding a query restriction, a feed comment cannot be queried directly, but can be queried only via the parent feed. An example is “select id, parentid, (select . . . from feedcomment) from entityfeed”. In another implementation, a feed comment can be directly queried, e.g., by querying comment table **930**. A query could include the text of a comment or any other column of the table.

[0210] In another implementation, regarding soft delete behavior, a feed comment table does not have a soft delete column. A soft delete allows an undelete action. In one implementation, a record can have a soft delete. Thus, when the record is deleted, the feed (and its children) can be soft deleted. Therefore, in one aspect, a feed comment cannot be retrieved via the “query” verb (which would retrieve only the comment), but can be retrieved via “queryAll” verb though. An example is queryAll(“select id, (select id, commentbody from feedcomments) from accountfeed where parentid=‘001x000xxx3MkADAA0’”); // where ‘001x000xxx3MkADAA0’ has been soft deleted. When a hard delete (a physical delete) happens, the comment can be hard deleted from the database.

[0211] In one implementation, regarding an implicit delete, feeds with comments are not deleted by a reaper (a routine that performs deletion). In another implementation, a user cannot delete a feed. In yet another implementation, upon lead convert (e.g., to an opportunity or contact), the feed items of the lead can be hard deleted. This implementation can be configured to perform such a deletion for any change in record type. In various implementations, only the comments are hard deleted upon a lead convert, other convert, or when the object is deleted (as mentioned above).

[0212] In one implementation, viewing a feed pulls up the most recent messages or feed tracked updates (e.g., 25) and searches the most recent (e.g., 4) comments for each feed item. The comments can be identified via the comment table **930**. In one implementation, a user can request to see more comments, e.g., by selecting a see more link.

[0213] In some implementations, user feeds and/or entity feeds have a last comment date field. In various implementations, the last comment date field is stored as a field of a record or a user profile. For feeds with no comments, this can be the same as the created date. Whenever a new comment is created, the associated feed’s last comment date can be updated with the created date of the comment. The last comment date is unchanged if a feed comment is deleted. A use case is to allow people to order their queries to see the feeds, which have been most recently commented on.

[0214] D. Creating Custom Feeds by Customizing the Event History Table

[0215] In some implementations, a tenant (e.g., through an administrator) or a specific user of a tenant can specify the types of events for which feed items are created. A user can add more events or remove events from a list of events that get added to the event history table **910**. In one implementation, a trigger can be added as a piece of code, rule, or item on a list for adding a custom event to the event history table **910**. These custom events can provide customers the ability to create their own custom feeds and custom feed items to augment or replace implicitly generated feeds via event history table **910**. Implicitly generated feed data can be created when feed-tracking is enabled for certain entities/field-names. In one implementation, in order to override implicit feeds, feed tracking can be turned off and then triggers can be defined by the user to add events to the event history table **910**. In other implementations, users are not allowed to override the default list of events that are added to table **910**, and thus cannot define their own triggers for having events tracked.

[0216] For example, upon lead convert or case close, a default action to be taken by the system may be to add multiple events to event history table **910**. If a customer (e.g., a tenant or a specific user) does not want each of these events to show up as feed items, the customer can turn off tracking for the entities and generate custom feeds by defining customized triggers (e.g., by using an API) upon the events. As another example, although data is not changed, a customer may still want to track an action on a record (e.g., status changes if not already being tracked, views by certain people, retrieval of data, etc.).

[0217] In one implementation, if a user does not want a feed item to be generated upon every change on a given field, but only if the change exceeds a certain threshold or range, then such custom feeds can be conditionally generated with the customized triggers. In one implementation, the default tracking for the record or user may be turned off for this customization so that the events are only conditionally tracked. In another implementation, a trigger can be defined that deletes events that are not desired, so that default tracking can still be turned on for a particular object type. Such conditional tracking can be used for other events as well.

[0218] In some implementations, defining triggers to track certain events can be done as follows. A user can define an object type to track. This object type can be added to a list of objects that can be tracked for a particular tenant. The tenant can remove object types from this list as well. Custom objects and standard objects can be on the list, which may, for example, be stored in cache or RAM of a server or in the database. Generally only one such list exists for a tenant, and users do not have individual lists for themselves, although in some implementations, they may particularly when the number of users in a tenant is small.

[0219] In one implementation, a tenant can select which records of an object type are to be tracked. In another implementation, once an object type is added to the tracking list of object types, then all records of that type are tracked. The tenant can then specify the particulars of how the tracking is to be performed. For example, the tenant can specify triggers as described above, fields to be tracked, or any of the customizations mentioned herein.

[0220] In some implementations, when a feed is defined as an object in the database (e.g., as a child object of entity records that can be tracked), a particular instance of the feed object (e.g., for a particular record) can be create-able and delete-able. In one implementation, if a user has access to a record then the user can customize the feed for the record. In one implementation, a record may be locked to prevent customization of its feed.

[0221] One method of creating a custom feed for users of a database system according to implementations is now described. Any of the following blocks can be performed wholly or partially with the database system, and in particular by one or more processor of the database system.

[0222] In block A, one or more criteria specifying which events are to be tracked for possible inclusion into a feed to be displayed are received from a tenant. In block B, data indicative of an event is received. In block C, the event is analyzed to determine if the criteria are satisfied. In block D, if the criteria are satisfied, at least a portion of the data is added to a table (e.g., one or more of the tables in FIG. 9A) that tracks events for inclusion into at least one feed for a user of the tenant. The feed in which feed items of an event may ultimately be displayed can be a news feed, record feed, or a profile feed.

[0223] E. Creating Custom Feeds with Filtering

[0224] After feed items have been generated, they can be filtered so that only certain feed items are displayed, which may be tailored to a specific tenant and/or user. In one implementation, a user can specify changes to a field that meet certain criteria for the feed item to show up in a feed displayed to the user, e.g., a news feed or even an entity feed displayed directly to the user. In one implementation, the criteria can be combined with other factors (e.g., number of feed items in the feed) to determine which feed items to display. For instance, if a small number of feed items exist (e.g., below a threshold), then all of the feed items may be displayed.

[0225] In one implementation, a user can specify the criteria via a query on the feed items in his/her new feed, and thus a feed may only return objects of a certain type, certain types of events, feed tracked updates about certain fields, and other criteria mentioned herein. Messages can also be filtered according to some criteria, which may be specified in a query. Such an added query can be added onto a standard query that is used to create the news feed for a user. A first user could specify the users and records that the first user is following in this manner, as well as identify the specific feed items that the first user wants to follow. The query could be created through a graphical interface or added by a user directly in a query language. Other criteria could include receiving only posts directed to a particular user or record, as opposed to other feed items.

[0226] In one implementation, the filters can be run by defining code triggers, which run when an event, specific or otherwise, occurs. The trigger could then run to perform the filtering at the time the event occurs or when a user (who has certain defined triggers, that is configured for a particular user) requests a display of the feed. A trigger could search for certain terms (e.g., vulgar language) and then remove such terms or not create the feed item. A trigger can also be used to send the feed item to a particular person (e.g., an administrator) who does not normally receive the feed item were it not for the feed item containing the flagged terms.

[0227] F. Access Checks

[0228] In one implementation, a user can access a feed of a record if the user can access the record. The security rules for determining whether a user has access to a record can be performed in a variety of ways, some of which are described in commonly assigned U.S. Pat. No. 8,095,531, titled METHODS AND SYSTEMS FOR CONTROLLING ACCESS TO CUSTOM OBJECTS IN A DATABASE, by Weissman et al., issued on Jan. 10, 2012, and hereby incorporated by reference in its entirety and for all purposes. For example, a security level table can specify whether a user can see a particular type of record and/or particular records. In one implementation, a hierarchy of positions within a tenant is used. For example, a manager can inherit the access levels of employees that the manager supervises. Field level security (FLS) can also be used to determine whether a particular feed tracked update about an update to a field can be seen by the user. The field change table 920 can be used to identify a field name or field ID, and then whether the user has read access to that field can be determined from an FLS table. For example, if a user could not see a field of a social security number, the feed of the user provided to the user would not include any feed items related to the social security number field.

[0229] In one implementation, a user can edit a feed of a record if the user has access to the record, e.g., deleting or editing a feed item. In another implementation, a user (besides an administrator) cannot edit a feed item, except for performing an action from which a feed item can be created. In one example, a user is first has to have access to a particular record and field for a feed item to be created based on an action of the user. In this case, an administrator can be considered to be a user with MODIFY-ALL-DATA security level. In yet another implementation, a user who created the record can edit the feed.

[0230] G. Posts

[0231] In one implementation, the text of posts are stored in a child table (post table 950), which can be cross-referenced with event history table 910. Post table 950 can include event ID 951 (to cross-reference with event ID 911), post text 952 to store the text of the post, and time/date 953. An entry in post table 950 can be considered a feed post object. Posts for a record can also be subject to access checks. In one implementation, if a user can view a record then all of the posts can be seen, i.e. there is not an additional level of security check as there is for FLS. In another implementation, an additional security check could be done, e.g., by checking on whether certain keywords (or phrases) exist in the post. For instance, a post may not be not provided to specified users if a certain keyword exists, or only provided to specified users if a keyword exists. In another implementation, a table can exist for status updates.

[0232] VIII. Subscribing to Users and Records to Follow

[0233] As described above, a user can follow users, groups, and records. Implementations can provide mechanisms for a user to manage which users, groups, and records that the user is currently following. In one implementation, a user can be limited to the number of users and records (collectively or separately) that the user can follow. For example, a user may be restricted to only following 10 users and 15 records, or as another example, 25 total. Alternatively, the user may be permitted to follow more or less users.

[0234] In one implementation, a user can go to a page of a record and then select to follow that object (e.g., with a

button marked “follow” or “join”). In another implementation, a user can search for a record and have the matching records show up in a list. The search can include criteria of records that the user might want to follow. Such criteria can include the owner, the creation date, last comment date, and numerical values of particular fields (e.g., an opportunity with a value of more than 10,000).

[0235] A follow button (or other activation object) can then reside next to each record in the resulting list, and the follow button can be selected to start following the record. Similarly, a user can go to a profile page of a user and select to follow the user, or a search for users can provide a list, where one or more users can be selected for following from the list. The selections of subscribing and unsubscribing can add and delete rows in table 920.

[0236] In some implementations, a subscription center acts as a centralized place in a database application (e.g., application platform 18) to manage which records a user subscribes to, and which field updates the user wants to see in feed tracked updates. The subscription center can use a subscription table to keep track of the subscriptions of various users. In one implementation, the subscription center shows a list of all the items (users and records) a user is subscribed to. In another implementation, a user can unsubscribe to subscribed objects from the subscription center.

[0237] A. Automatic Subscription

[0238] In one implementation, an automatic subscription feature can ensure that a user is receiving certain feeds. In this manner, a user does not have to actively select certain objects to follow. Also, a tenant can ensure that a user is following objects that the user needs to be following.

[0239] In various implementations for automatically following users, a default for small organizations can be to follow everyone. For big organizations, the default can be to follow a manager and peers. If a user is a manager, the default can be to follow the manager’s supervisor, peers, and people that the manager supervises (subordinates). In other implementations for automatically following records, records that the user owns may be automatically followed and/or records recently viewed (or changed) may be automatically followed.

[0240] In one example, a new record is created. The owner (not necessarily the user who created the entity) is subscribed to the entity. If ownership is changed, the new owner may automatically be subscribed to follow the entity. Also, after a lead convert, the user doing the lead convert may be automatically subscribed to the new account, opportunity, or contact resulting from the lead convert. In one implementation, the auto subscription is controlled by user preference. That is a user or tenant can have the auto subscribe feature enabled or not. In one aspect, the default is to have the auto-subscribe turned on.

[0241] FIG. 9B shows a flowchart of an example of a method 900 for automatically subscribing a user to an object in a database system, performed in accordance with some implementations. Any of the following blocks can be performed wholly or partially with the database system, and in particular by one or more processor of the database system.

[0242] In block 901, one or more properties of an object stored in the database system are received. The properties can be received from administrators of the database system, or from users of the database system (which may be an administrator of a customer organization). The properties can be records or users, and can include any of the fields of

the object that are stored in the database system. Examples of properties of a record include: an owner of the record, a user that converted the record from one record type to another record type, whether the first user has viewed the record, and a time the first user viewed the record. Examples of properties of a user include: which organization (tenant) the user is associated with, the second user’s position in the same organization, and which other users the user had emailed or worked with on projects.

[0243] In block 902, the database system receives one or more criteria about which users are to automatically follow the object. The criteria can be received from administrators of the database system, or from one or more users of the database system. The users may be an administrator of a customer organization, which can set tenant-wide criteria or criteria for specific users (who may also set the criteria themselves). Examples of the criteria can include: an owner or creator of a record is to follow the record, subordinates of an owner or creator of a record are to follow the record, a user is to follow records recently viewed (potentially after a specific number of views), records that a user has changed values (potentially with a date requirement), and records created by others in a same business group as the user. Examples of the criteria can also include: a user is to follow his/her manager, the user’s peers, other users in the same business group as the user, and other users that the user has emailed or worked with on a project. The criteria can be specific to a user or group of users (e.g., users of a tenant).

[0244] In block 903, the database system determines whether the one or more properties of the object satisfy the one or more criteria for a first user. In one implementation, this determination can occur by first obtaining the criteria and then determining objects that satisfy the criteria. The determination can occur periodically, at time of creation of an object, or at other times. If different users have different criteria, then the criteria for a particular user or group could be searched at the same time. Since users of different tenants normally cannot view objects of another tenant, certain criteria do not have to be checked. In another implementation, this determination can occur by looking at certain properties and then identifying any criteria that are met. In yet another implementation, the criteria and properties can be used to find users that satisfy the criteria.

[0245] In block 904, if the criteria are satisfied, the object is associated with the first user. The association can be in a list that stores information as to what objects are being followed by the first user. User subscription table 940 is an example of such a list. In one implementation, the one or more criteria are satisfied if one property satisfies at least one criterion. Thus, if the criteria are that a user follows his/her manager and the object is the user’s manager, then the first user will follow the object.

[0246] In one implementation, a user can also be automatically unsubscribed, e.g., if a certain action happens. The action could be a change in the user’s position within the organization, e.g., a demotion or becoming a contractor. As another example, if a case gets closed, then users following the case may be automatically unsubscribed.

[0247] B. Feed and Subscription API

[0248] In one implementation, a feed and subscription center API can enable tenants to provide mechanisms for tracking and creating feed items, e.g., as described above for creating custom feeds by allowing users to add custom events for tracking. For example, after some initial feed

items are created (e.g., by administrators of the database system), outside groups (e.g., tenants or software providers selling software to the tenants) can ‘enable objects’ for feeds through a standard API. The groups can then integrate into the subscription center and the feed tracked update feeds on their own. In one implementation, the feed and subscription center API can use a graphical user interface implemented for the default feed tracking. In one implementation, API examples include subscribing to an entity by creating a new entity subscription object for a particular user ID, or for all users of a tenant (e.g., user subscription table 940). In one implementation, obtaining all subscriptions for a given user can be performed by using a query, such as “select . . . from EntitySubscription where userid=‘ . . . ’”.

[0249] Some implementations have restriction on non-admin users, e.g., those without view all data permissions (VAD). One restriction can be a limit clause on entity subscription queries (e.g., queries on user subscription table 940), e.g., where the limit of the number of operations is less than 100. In one implementation, users are not required to specify an order-by, but if an order-by is specified they can only order on fields on the entity subscription entity. In one implementation, filters on entity subscription can likewise only specify fields on the entity subscription entity. In one aspect, the object ID being followed can be sorted or filtered, but not the object name.

[0250] In one implementation, one or more restrictions can also be placed on the identification of feed items in a feed that a user can access. For example, if a low-level user (i.e. user can access few objects) is attempting to see a profile feed of a high level user, a maximum number of checks (e.g., 500) for access rights may be allowed. Such a restriction can minimize a cost of a feed request. In some implementations, there are restriction on the type of queries (e.g., fields for filtering) allowed to construct on feeds (e.g., on tables in FIG. 9A).

[0251] C. Sharing

[0252] As mentioned above, users may be restricted from seeing records from other tenants, as well as certain records from the tenant to which the user belongs (e.g., the user’s employer). Sharing rules can refer to the access rules that restrict a user from seeing records that the user is not authorized to see or access. Additionally, in one implementation, a user may be restricted to only seeing certain fields of a record, field-level security (FLS).

[0253] In an implementation, access rule checks are done upon subscription. For example, a user is not allowed to subscribe to a record or type of record that the user cannot access. In one aspect, this can minimize (but not necessarily eliminate) cases where a user subscribes to entities they cannot access. Such cases can slow down news feed queries, when an access check is performed (which can end up removing much of the feed items). Thus, a minimization of access checks can speed up operation. In another implementation, when feed items are created dynamically, access rule checks may be done dynamically at the time of subsequent access, and not upon subscription or in addition to at time of subscription.

[0254] An example case where access checks are still performed is when a first user follows a second user, but the second user performs some actions on records or is following records that the first user is not allowed to see. The first user may be allowed to follow the second user, and thus the subscription is valid even though the first user may not be

able to see all of the feed items. Before a feed tracked update is provided to a news feed of the first user, a security check may be performed to validate whether the first user has access rights to the feed item. If not, the feed item is not displayed to the first user. In one implementation, users can be blocked from feed items that contain certain terms, symbols, account numbers, etc. In one implementation, any user can follow another user. In another implementation, users may be restricted as to which users, objects, and/or records he/she can follow.

[0255] Regarding viewing privileges of a feed, in one implementation, a user can see all of his own subscriptions (even if he’s lost read access to a record). For example, a user can become a contractor, and then the user may lose access to some records. But, the user may still see that he/she is following the object. This can help if there is a limit to the number of objects that can be followed. To unsubscribe a user may need to know what they are following so they can unsubscribe and subscribe to objects the user can see. In another implementation, for access to other people’s subscriptions, a user can be required to need read-access on the record-id to see the subscription. In some implementations, users with authorization to modify all data can create/delete any subscription. In other implementations, a user can create/delete subscriptions only for that user, and not anyone else.

[0256] D. Configuration of which Field to Follow

[0257] There can be various feed settings for which feed items get added to profile and record feeds, and which get added to news feeds. In one implementation, for profile feeds and entity feeds, feed tracked updates can be written for all standard and custom fields on the supported objects. In one implementation, feed settings can be set to limit how many and which fields of a record are tracked for determining whether a feed tracked update is to be generated. For example, a user or administrator can choose specific fields to track and/or certain ones not to track. In another implementation, there is a separate limit for the number of trackable fields (e.g., 20) for a record. Thus, only certain changes may be tracked in an entity feed tracked update and show up in the feed. In yet another implementation, default fields may be chosen for tracking, where the defaults can be exposed in the subscriptions center.

[0258] IX. Adding Items to a Feed

[0259] As described above, a feed includes feed items, which include feed tracked updates and messages, as defined herein. Various feeds can be generated. For example, a feed can be generated about a record or about a user. Then, users can view these feeds. A user can separately view a feed of a record or user, e.g., by going to a home page for the user or the record. As described above, a user can also follow another user or record and receive the feed items of those feeds through a separate feed application (e.g., in a page or window), which is termed “chatter” in certain examples. The feed application can provide each of the feeds that a user is following and, in some examples, can combine various feeds in a single information feed.

[0260] A feed generator can refer to any software program running on a processor or a dedicated processor (or combination thereof) that can generate feed items (e.g., feed tracked updates or messages) and combine them into a feed. In one implementation, the feed generator can generate a feed item by receiving a feed tracked update or message, identifying what feeds the item should be added to, and

adding the feed. Adding the feed can include adding additional information (metadata) to the feed tracked update or message (e.g., adding a document, sender of message, a determined importance, etc.). The feed generator can also check to make sure that no one sees feed tracked updates for data that they don't have access to see (e.g., according to sharing rules). A feed generator can run at various times to pre-compute feeds or to compute them dynamically, or combinations thereof.

[0261] In one implementation, the feed generator can de-dupe events (i.e. prevent duplicates) that may come in from numerous records (and users). For example, since a feed tracked update can be published to multiple feeds (e.g., John Choe changed the Starbucks Account Status) and a person can be subscribed to both the Starbucks account and John Choe, implementations can filter out duplicates before adding or displaying the items in a news feed. Thus, the Feed Generator can collapse events with multiple records and users for a single transaction into a single feed tracked update and ensure the right number of feed tracked updates for the particular feed. In some implementations, an action by a user does not create a feed item for that user (e.g., for a profile feed of that user), and it is only the feed of the object being acted upon (e.g., updated) for which a feed item is created. Thus, there should not be duplicates. For example, if someone updates the status of a record, the feed item is only for the record and not the user.

[0262] In one implementation, processor 417 in FIG. 4 can identify an event that meets criteria for a feed tracked update, and then generate the feed tracked update. Processor 417 can also identify a message. For example, an application interface can have certain mechanisms for submitting a message (e.g., "submit" buttons on a profile page, detail page of a record, "comment" button on post), and use of these mechanisms can be used to identify a message to be added to a table used to create a feed or added directly to a list of feed items ready for display.

[0263] A. Adding Items to a Pre-Computed Feed

[0264] In some implementations, a feed of feed items is created before a user requests the feed. Such an implementation can run fast, but have high overall costs for storage. In one implementation, once a profile feed or a record feed has been created, a feed item (messages and feed tracked updates) can be added to the feed. The feed can exist in the database system in a variety of ways, such as a related list. The feed can include mechanisms to remove items as well as add them.

[0265] As described above, a news feed can be an aggregated feed of all the record feeds and profile feeds to which a user has subscribed. The news feed can be provided on the home page of the subscribing user. Therefore, a news feed can be created by and exist for a particular user. For example, a user can subscribe to receive entity feeds of certain records that are of interest to the user, and to receive profile feeds of people that are of interest (e.g., people on a same team, that work for the user, are a boss of the user, etc.). A news feed can tell a user about all the actions across all the records (and people) whom have explicitly (or implicitly) been subscribed to via the subscriptions center (described above).

[0266] In one implementation, only one instance of each feed tracked update is shown on a user's news feed, even if the feed tracked update is published in multiple entities to which the user is subscribed. In one aspect, there may be

delays in publishing news articles. For example, the delay may be due to queued up messages for asynchronous entity feed tracked update persistence. Different feeds may have different delays (e.g., delay for new feeds, but none of profile and entity feeds). In another implementation, certain feed tracked updates regarding a subscribed profile feed or an entity feed are not shown because the user is not allowed access, e.g., due to sharing rules (which restrict which users can see which data). Also, in one implementation, data of the record that has been updated (which includes creation) can be provided in the feed (e.g., a file or updated value of a feed can be added as a flash rendition).

[0267] Examples are provided below as how it can be determined which feed items to add to which news feeds. In one implementation, the addition of items to a news feed is driven by the following user. For example, the user's profile can be checked to determine objects the user is following, and the database may be queried to determine updates to these objects. In another implementation, the users and records being followed drive the addition of items to a news feed. Implementations can also combine these and other aspects. In one implementation, a database system can be follower-driven if the number of subscriptions (users and records the user is following) is small. For example, since the number subscriptions are small, then changes to a small number of objects need to be checked for the follower.

[0268] Regarding implementations that are follower-driven, one implementation can have a routine run for a particular user. The routine knows the users and records that the user is following. The routine can poll the database system for new feed tracked updates and messages about the users and records that are being followed. In one implementation, the polling can be implemented as queries. In one implementation, the routine can run at least partially (even wholly) on a user device.

[0269] Regarding implementations where a news feed is driven by the record (or user) being followed, processor 417 can identify followers of the record after a feed item is added to the record feed. Processor 417 can retrieve a list of the followers from the database system. The list can be associated with the record, and can be stored as a related list or other object that is a field or child of the record.

[0270] In one implementation, profile and record feeds can be updated immediately with a new feed item after an action is taken or an event occurs. A news feed can also be updated immediately. In another implementation, a news feed can be updated in batch jobs, which can run at periodic times.

[0271] B. Dynamically Generating Feeds

[0272] In some implementations, a feed generator can generate the feed items dynamically when a user requests to see a particular feed, e.g., a profile feed, entity feed, or the user's news feed. In one implementation, the most recent feed items (e.g., top 50) are generated first. In one aspect, the other feed items can be generated as a background process, e.g., not synchronously with the request to view the feed. However, since the background process is likely to complete before a user gets to the next 50 feed items, the feed generation may appear synchronous. In another aspect, the most recent feed items may or may not include comments, e.g., that are tied to feed tracked updates or posts.

[0273] In one implementation, the feed generator can query the appropriate subset of tables shown in FIG. 9A and/or other tables as necessary, to generate the feed items for display. For example, the feed generator can query the

event history table **910** for the updates that occurred for a particular record. The ID of the particular record can be matched against the ID of the record. In one implementation, changes to a whole set of records can be stored in one table. The feed generator can also query for status updates, posts, and comments, each of which can be stored in different parts of a record or in separate tables, as shown in FIG. **9A**. What gets recorded in the entity event history table (as well as what is displayed) can be controlled by a feed settings page in setup, which can be configurable by an administrator and can be the same for the entire organization, as is described above for custom feeds.

[0274] In one implementation, there can be two feed generators. For example, one generator can generate the record and profile feeds and another generator can generate news feeds. For the former, the feed generator can query identifiers of the record or the user profile. For the latter, the news feed generator can query the subscribed profile feeds and record feeds, e.g., user subscription table **940**. In one implementation, the feed generator looks at a person's subscription center to decide which feeds to query for and return a list of feed items for the user. The list can be de-duped, e.g., by looking at the event number and values for the respective table, such as field name or ID, comment ID, or other information.

[0275] C. Adding Information to Feed Tracked Update Tables

[0276] FIG. **10** shows a flowchart of an example of a method **1000** for saving information to feed tracking tables, performed in accordance with some implementations. In one implementation, some of the blocks may be performed regardless of whether a specific event or part of an event (e.g., only one field of an update is being tracked) is being tracked. In various implementations, a processor or set of processors (hardwired or programmed) can perform method **1000** and any other method described herein.

[0277] In block **1010**, data indicative of an event is received. The data may have a particular identifier that specifies the event. For example, there may be a particular identifier for a field update. In another implementation, the transaction may be investigated for keywords identifying the event (e.g., terms in a query indicating a close, change field, or create operations).

[0278] In block **1020**, it is determined whether the event is being tracked for inclusion into feed tracked update tables. The determination of what is being tracked can be based on a tenant's configuration as described above. In one aspect, the event has an actor (person performing an event), and an object of the event (e.g., record or user profile being changed).

[0279] In block **1030**, the event is written to an event history table (e.g., table **910**). In one implementation, this feed tracking operation can be performed in the same transaction that performs a save operation for updating a record. In another implementation, a transaction includes at least two roundtrip database operations, with one roundtrip being the database save (write), and the second database operation being the saving of the update in the feed tracked update table. In one implementation, the event history table is chronological. In another implementation, if user A posts on user B's profile, then user A is under the "created by" **913** and user B is under the object ID **912**.

[0280] In block **1040**, a field change table (e.g., field change table **920**) can be updated with an entry having the

event identifier and fields that were changed in the update. In one implementation, the field change table is a child table of the event history table. This table can include information about each of the fields that are changed. For example, for an event that changes the name and balance for an account record, an entry can have the event identifier, the old and new name, and the old and new balance. Alternatively, each field change can be in a different row with the same event identifier. The field name or ID can also be included to determine which field the values are associated.

[0281] In block **1050**, when the event is a post, a post table (e.g., post table **950**) can be updated with an entry having the event identifier and text of the post. In one implementation, the field change table is a child table of the event history table. In another implementation, the text can be identified in the transaction (e.g., a query command), stripped out, and put into the entry at the appropriate column. The various tables described herein can be combined or separated in various ways. For example, the post table and the field change table may be part of the same table or distinct tables, or may include overlapping portions of data.

[0282] In block **1060**, a comment is received for an event and the comment is added to a comment table (e.g., comment table **930**). The comment could be for a post or an update of a record, from which a feed tracked update can be generated for display. In one implementation, the text can be identified in the transaction (e.g., a query command), stripped out, and put into the entry at the appropriate column.

[0283] D. Reading Information from Feed Tracked Update Tables

[0284] FIG. **11** shows a flowchart of an example of a method **1100** for reading a feed item as part of generating a feed for display, performed in accordance with some implementations. In one implementation, the feed item may be read as part of creating a feed for a record.

[0285] In block **1110**, a query is received for an events history table (e.g., event history table **910**) for events related to a particular record. In one implementation, the query includes an identifier of the record for which the feed is being requested. In various implementations, the query may be initiated from a detail page of the record, a home page of a user requesting the record feed, or from a listing of different records (e.g., obtained from a search or from browsing).

[0286] In block **1120**, the user's security level can be checked to determine if the user can view the record feed. Typically, a user can view a record feed, if the user can access the record. This security check can be performed in various ways. In one implementation, a first table is checked to see if the user has a classification (e.g., a security level that allows him to view records of the given type). In another implementation, a second table is checked to see if the user is allowed to see the specific record. The first table can be checked before the second table, and both tables can be different sections of a same table. If the user has requested the feed from the detail page of the record, one implementation can skip the security level check for the record since the check was already done when the user requested to view the detail page.

[0287] In one implementation, a security check is determined upon each request to view the record feed. Thus, whether or not a feed item is displayed to a user is determined based on access rights, e.g., when the user requests to

see a feed of a record or a news feed of all the objects the user is following. In this manner, if a user's security changes, a feed automatically adapts to the user's security level when it is changed. In another implementation, a feed can be computed before being requested and a subsequent security check can be made to determine whether the person still has access right to view the feed items. The security (access) check may be at the field level, as well as at the record level.

[0288] In block **1130**, if the user can access the record, a field level security table can be checked to determine whether the user can see particular fields. In one implementation, only those fields are displayed to the user. Alternatively, a subset of those the user has access to is displayed. The field level security check may optionally be performed at the same time and even using the same operation as the record level check. In addition, the record type check may also be performed at this time. If the user can only see certain fields, then any feed items related to those fields (e.g., as determined from field change table **920**) can be removed from the feed being displayed.

[0289] In block **1140**, the feed items that the user has access to are displayed. In one implementation, a predetermined number (e.g., 20) of feed items are displayed at a time. The method can display the first 20 feed items that are found to be readable, and then determine others while the user is viewing the first 20. In another implementation, the other feed items are not determined until the user requests to see them, e.g., by activating a see more link.

[0290] FIG. **12** shows a flowchart of an example of a method **1200** for reading a feed item of a profile feed for display, performed in accordance with some implementations. In one implementation, the query includes an identifier of the user profile feed that is being requested. Certain blocks may be optional, as is also true for other methods described herein. For example, security checks may not be performed.

[0291] In block **1210**, a query is directed to an event history table (e.g., event history table **910**) for events having a first user as the actor of the event (e.g., creation of an account) or on which the event occurred (e.g., a post to the user's profile). In various implementations, the query may be initiated by a second user from the user's profile page, a home page of a user requesting the profile feed (e.g., from a list of users being followed), or from a listing of different users (e.g., obtained from a search or from browsing). Various mechanisms for determining aspects of events and obtaining information from tables can be the same across any of the methods described herein.

[0292] In block **1220**, a security check may also be performed on whether the second user can see the first user's profile. In one implementation any user can see the profile of another user of the same tenant, and block **1220** is optional.

[0293] In block **1230**, a security (access) check can be performed for the feed tracked updates based on record types, records, and/or fields, as well security checks for messages. In one implementation, only the feed tracked updates related to records that the person has updated are the ones that need security check as the feed items about the user are readable by any user of the same tenant. Users of other tenants are not navigable, and thus security can be enforced at a tenant level. In another implementation, messages can

be checked for keywords or links to a record or field that the second user does not have access.

[0294] As users can have different security classifications, it is important that a user with a low-level security cannot see changes to records that have been performed by a user with high-level security. In one implementation, each feed item can be checked and then the viewable results displayed, but this can be inefficient. For example, such a security check may take a long time, and the second user would like to get some results sooner rather than later. The following blocks illustrate one implementation of how security might be checked for a first user that has a lot of feed items, but the second user cannot see most of them. This implementation can be used for all situations, but can be effective in the above situation.

[0295] In block **1231**, a predetermined number of entries are retrieved from the event history table (e.g., starting from the most recent, which may be determined from the event identifier). The retrieved entries may just be ones that match the user ID of the query. In one implementation, entries are checked to find the entries that are associated with the user and with a record (i.e. not just posts to the user account). In another implementation, those entries associated with the user are allowed to be viewed, e.g., because the second user can see the profile of the first user as determined in block **1220**.

[0296] In block **1232**, the record identifiers are organized by type and the type is checked on whether the second user can see the record types. Other checks such as whether a record was manually shared (e.g., by the owner) can also be performed. In one implementation, the queries for the different types can be done in parallel.

[0297] In block **1233**, if a user can see the record type, then a check can be performed on the specific record. In one implementation, if a user can see a record type, then the user can see all of the records of that type, and so this block can be skipped. In another implementation, the sharing model can account for whether a user below the second user (e.g., the second user is a manager) can see the record. In such an implementation, the second user may see such a record. In one implementation, if a user cannot see a specific record, then comments on that record are also not viewable.

[0298] In block **1234**, field level sharing rules can be used to determine whether the second user can see information about an update or value of certain fields. In one implementation, messages can be analyzed to determine if reference to a particular field name is made. If so, then field level security can be applied to the messages.

[0299] In block **1280**, blocks **1231-1234** are repeated until a stopping criterion is met. In one implementation, the stopping criteria may be when a maximum number (e.g., 100) of entries that are viewable have been identified. In another implementation, the stopping criteria can be that a maximum number (e.g., 500) of entries from the entity feed tracked update table have been analyzed, regardless of whether the entries are viewable or not.

[0300] In one implementation, a news feed can be generated as a combination of the profile feeds and the entity feeds, e.g., as described above. In one implementation, a list of records and user profiles for the queries in blocks **1110** and **1210** can be obtained from user subscription table **940**. In one implementation, there is a maximum number of objects that can be followed.

[0301] In various implementations, the entity feed table can be queried for any one or more of the following matching variables as part of determining items for a feed: CreatedDate, CreatedById, CreatedBy.FirstName, CreatedBy.LastName, ParentId, and Parent.Name. The child tables can also be queried for any one or more of the following matching variables as part of determining items for a feed: DataType, FieldName, OldValue, and NewValue. A query can also specify how the resulting feed items can be sorted for display, e.g., by event number, date, importance, etc. The query can also include a number of items to be returned, which can be enforced at the server.

[0302] The two examples provided above can be done periodically to create the feeds ahead of time or done dynamically at the time the display of a feed is requested. Such a dynamic calculation can be computationally intensive for a news feed, particularly if many users and records are being followed, although there can be a low demand for storage. Accordingly, one implementation performs some calculations ahead of time and stores the results in order to create a news feed.

[0303] E. Partial Pre-Computing of Items for a Feed

[0304] FIG. 13 shows a flowchart of an example of a method 1300 of storing event information for efficient generation of feed items to display in a feed, performed in accordance with some implementations. In various implementations, method 1300 can be performed each time an event is written to the event history table, or periodically based on some other criteria (e.g., every minute, after five updates have been made, etc.).

[0305] In block 1310, data indicative of an event is received. The data may be the same and identified in the same way as described for block 1010. The event may be written to an event history table (e.g., table 910).

[0306] In block 1320, the object(s) associated with the event are identified. In various implementations, the object may be identified by according to various criteria, such as the record being changed, the user changing the record, a user posting a message, and a user whose profile the message is being posted to.

[0307] In block 1330, the users following the event are determined. In one implementation, one or more objects that are associated with the event are used to determine the users following the event. In one implementation, a subscription table (e.g., table 940) can be used to find the identified objects. The entries of the identified objects can contain an identifier (e.g., user ID 941) of each of the users following the object

[0308] In block 1340, the event and the source of the event, e.g., a record (for a record update) or a posting user (for a user-generated post) are written to a news feed table along with an event identifier. In one implementation, such information is added as a separate entry into the news feed table along with the event ID. In another implementation, each of the events for a user is added as a new column for the row of the user. In yet another implementation, more columns (e.g., columns from the other tables) can be added.

[0309] News feed table 960 shows an example of such a table with user ID 961 and event ID or pointer 962. The table can be organized in any manner. One difference from event history table 910 is that one event can have multiple entries (one for each subscriber) in the news feed table 960. In one implementation, all of the entries for a same user are grouped together, e.g., as shown. The user U819 is shown as

following events E37 and E90, and thus any of the individual feed items resulting from those events. In another implementation, any new entries are added at the end of the table. Thus, all of the followers for a new event can be added as a group. In such an implementation, the event IDs would generally be grouped together in the table. Of course, the table can be sorted in any suitable manner.

[0310] In an implementation, if the number of users is small, then the feed items in one or more of the tables may be written as part of the same write transaction. In one implementation, the determination of small depends on the number of updates performed for the event (e.g., a maximum number of update operations may be allowed), and if more operations are performed, then the addition of the feed items is performed. In one aspect, the number of operations can be counted by the number of rows to be updated, including the rows of the record (which depends on the update event), and the rows of the feed tracked update tables, which can depend on the number of followers. In another implementation, if the number of users is large, the rest of the feed items can be created by batch. In one implementation, the feed items are written as part of a different transaction, i.e., by batch job.

[0311] In one implementation, security checks can be performed before an entry is added to the news feed table 960. In this manner, security checks can be performed during batch jobs and may not have to be performed at the time of requesting a news feed. In one implementation, the event can be analyzed and if access is not allowed to a feed item of the event, then an entry is not added. In one aspect, multiple feed items for a same user may not result from a same event (e.g., by how an event is defined in table 910), and thus there is no concern about a user missing a feed item that he/she should be able to view.

[0312] In block 1350, a request for a news feed is received from a user. In one implementation, the request is obtained when a user navigates to the user's home page. In another implementation, the user selects a table, link, or other page item that causes the request to be sent.

[0313] In block 1360, the news feed table and other tables are accessed to provide displayable feed items of the news feed. The news feed can then be displayed. In one implementation, the news feed table can then be joined with the event history table to determine the feed items. For example, the news feed table 960 can be searched for entries with a particular user ID. These entries can be used to identify event entries in event history table 910, and the proper information from any child tables can be retrieved. The feed items (e.g., feed tracked updates and messages) can then be generated for display.

[0314] In one implementation, the most recent feed items (e.g., 100 most recent) are determined first. The other feed items may then be determined in a batch process. Thus, the feed item that a user is most likely to view can come up first, and the user may not recognize that the other feed items are being done in batch. In one implementation, the most recent feed items can be gauged by the event identifiers. In another implementation, the feed items with a highest importance level can be displayed first. The highest importance being determined by one or more criteria, such as, who posted the feed item, how recently, how related to other feed items, etc.

[0315] In one implementation where the user subscription table 940 is used to dynamically create a news feed, the query would search the subscription table, and then use the

object IDs to search the event history table (one search for each object the user is following). Thus, the query for the news feed can be proportional to the number of objects that one was subscribing to. The news feed table allows the intermediate block of determining the object IDs to be done at an earlier stage so that the relevant events are already known. Thus, the determination of the feed is no longer proportional to the number of object being followed.

[0316] In some implementations, a news feed table can include a pointer (as opposed to an event identifier) to the event history table for each event that is being followed by the user. In this manner, the event entries can immediately be retrieved without having to perform a search on the event history table. Security checks can be made at this time, and the text for the feed tracked updates can be generated.

[0317] X. Display of a Feed

[0318] Feeds include messages and feed tracked updates and can show up in many places in an application interface with the database system. In one implementation, feeds can be scoped to the context of the page on which they are being displayed. For example, how a feed tracked update is presented can vary depending on which page it is being displayed (e.g., in news feeds, on a detail page of a record, and even based on how the user ended up at a particular page). In another implementation, only a finite number of feed items are displayed (e.g., 50). In one implementation, there can be a limit specifically on the number of feed tracked updates or messages displayed. Alternatively, the limit can be applied to particular types of feed tracked updates or messages. For example, only the most recent changes (e.g., 5 most recent) for a field may be displayed. Also, the number of fields for which changes are displayed can also be limited. Such limits can also be placed on profile feeds and news feeds. In one implementation, feed items may also be subject to certain filtering criteria before being displayed, e.g., as described below.

[0319] A. Sharing Rules for Feeds

[0320] As mentioned above, a user may not be allowed to see all of the records in the database, and not even all of the records of the organization to which the user belongs. A user can also be restricted from viewing certain fields of a record that the user is otherwise authorized to view. Accordingly, certain implementations use access rules (also called sharing rules and field-level security FLS) to ensure that a user does not view a feed tracked update or message that the user is not authorized to see. A feed of a record can be subject to the same access rules as the parent record.

[0321] In one implementation, access rules can be used to prevent subscription to a record that the user cannot see. In one implementation, a user can see a record, but only some of the fields. In such instances, only items about fields that the user can access may be displayed. In another implementation, sharing rules and FLS are applied before a feed item is being added to a feed. In another implementation, sharing rules and FLS are applied after a feed item has been added and when the feed is being displayed. When a restriction of display is mentioned, the enforcement of access rules may occur at any stage before display.

[0322] In some implementations, the access rules can be enforced when a query is provided to a record or a user's profile to obtain feed items for a news feed of a user. The access rules can be checked and cross-referenced with the feed items that are in the feed. Then, the query can only return feed items for which the user has access.

[0323] In other implementations, the access rules can be enforced when a user selects a specific profile feed or record feed. For example, when a user arrives on a home page (or selects a tab to see the record feed), the database system can check to see which feed items the user can see. In such an implementation, each feed item can be associated with metadata that identifies which field the feed item is about. Thus, in one implementation, a feed tracked update is not visible unless the associated record and/or field are visible to the user.

[0324] In one example, when a user accesses a feed of a record, an access check can be performed to identify whether the user can access the object type of the record. In one implementation, users are assigned a profile type, and the profile type is cross-referenced (e.g., by checking a table) to determine whether the profile type of the user can see the object type of the record.

[0325] In some implementations, access to specific records can be checked, e.g., after it has been determined that the user can access the record type. Rules can be used to determine the records viewable by a user. Such rules can determine the viewable records as a combination of those viewable by profile type, viewable due to a profile hierarchy (e.g., a boss can view records of profile types lower in the hierarchy), and viewable by manual sharing (e.g., as may be done by an owner of a record). In one implementation, the records viewable by a user can be determined beforehand and stored in a table. In one implementation, the table can be cross-referenced by user (or profile type of a user) to provide a list of the records that the user can see, and the list can be searched to determine if the record at issue is among the list. In another implementation, the table can be cross-referenced by record to determine a list of the profile types that can access the record, and the list can be searched to find out if the requesting user is in the list. In another implementation, the records viewable by a user can be determined dynamically at the time of the access check, e.g., by applying rules to data (such as user profile and hierarchy information) obtained from querying one or more tables.

[0326] In other implementations, checks can be made as to whether a user has access to certain fields of a record, e.g., after it has been determined that the user can access the record. In one aspect, the access check on fields can be performed on results already obtained from the database, to filter out fields that the user cannot see. In one implementation, the fields associated with retrieved feed items are determined, and these fields are cross-referenced with an access table that contains the fields accessible by the user (e.g., using the profile type of the user). Such an access table could also be a negative access table by specifying fields that the user cannot see, as can other access tables mentioned herein. In one implementation, the field level access table is stored in cache at a server.

[0327] In one implementation, a user can see the same fields across all records of a certain type (e.g., as long as the user can see the record). In one implementation, there is a field level access table for each object type. The access table can be cross-referenced by user (e.g., via profile type) or field. For example, a field can be identified along with the profile types that can see the field, and it can be determined whether the user's profile type is listed. In another example, the user can be found and the fields to which the user has access can be obtained. In another implementation, the accessible fields could be specified for each record.

[0328] Regarding profile feeds and news feeds, a first user may perform an action on a record, and a feed tracked update may be generated and added to the first user's profile feed. A second user who is allowed to follow the first user may not have access rights to the record. Thus, the feed tracked update can be excluded from a news feed of the second user, or when the second user views the first user's profile feed directly. In one implementation, if a user is already on the detail page, then another access check (at least at the record level) may optionally not be performed since a check was already done in order to view the detail page.

[0329] In some implementations, for profile feeds and news feeds, the feed items can be organized by object type. IT can then be determined whether the requesting user can access to those object types. Other access checks can be done independently or in conjunction with these access checks, as is described above.

[0330] B. API Implementation

[0331] Various implementations can implement the access rules in various ways. In one implementation, all recent feed items (or more generally events) are retrieved from a feed that is ready for display (e.g., after a feed generator performs formatting) or a table. Then, bulk sharing checks can be applied on the retrieved items. The viewable feed items of the most recent set can then be displayed.

[0332] In another implementation regarding a profile feed, for non-VAD (view all data) users, i.e. users who can see everything, certain functions can be overridden. In one implementation, a FROM clause in a query can be overridden to be a pipelined function, e.g., with different parts of the query being operated on at the same time, but with different operations of a pipeline. This pipeline function can be given a row limit and the maximum number of sharing checks to run. It can loop, selecting the next batch of rows, run sharing checks against them in bulk, and pipe back any IDs which are accessible. In one aspect, in nearly all cases, the user feed can contain accessible IDs so the sharing checks can pass on the first loop. However, it is possible the sharing may have changed such that this user's access is greatly reduced. In one worst case, implementations can run sharing checks on up to the maximum number of sharing check rows (e.g., a default 500) and then terminate the function with the IDs which passed so far, possibly zero. Such an example includes a low level person viewing profile feed of CEO.

[0333] In some implementations, if the user has a small number of subscriptions (e.g., <25), then implementations can first run sharing checks on those IDs and then drive the main query from those accessible IDs, as opposed to a semi-join against the subscription and running sharing checks on the resulting rows. In other implementations, FLS is enforced by building up a TABLE CAST of the accessible field IDs from the cached values. A main query can then join against this table to filter only accessible fields.

[0334] XI. Filtering and Searching Feeds

[0335] It can be possible that a user subscribes to many users and records, which can cause a user's news feed to be very long and include many feed items. In such instances, it can be difficult for the user to read every feed item, and thus some important or interesting feed items may not be read. In some implementations, filters may be used to determine which feed items are added to a feed or displayed in the feed, even though a user may be authorized to see more than what is displayed. Section VII.E also provides a description of filtering based on criteria.

[0336] In one implementation, an "interestingness" filter can function as a module for controlling/recommending which feed tracked updates make it to the news feed when the number of items that a user subscribes to is large. In one such implementation, a user can specify a filter, which is applied to a user's news feed or to record and profile feeds that the user requests. Different filters can be used for each. For example, processing can be done on the news feed to figure out which feed tracked updates are the most relevant to the user. One implementation can use an importance weight and level/ranking, as described herein. Other implementations can include a user specifying keywords for a message and specifying which records or users are most important.

[0337] In one implementation, a filter can be used that only allows certain feed items to be added to a feed and/or to be displayed as part of a feed. A filter can be used such that the removal or non-addition of certain feed items automatically occur for any new feed items after the filter criteria are entered. The filter criteria can also be added retroactively. The criteria of such a filter can be applied via a query mechanism as part of adding a feed item to a table or displaying a feed, as described in sections above. In various implementations, a user can directly write a query or create the query through a graphical user interface.

[0338] FIG. 14 shows a flowchart of an example of a method 1400 for creating a custom feed for users of a database system using filtering criteria, performed in accordance with some implementations. Any of the following blocks can be performed wholly or partially with the database system, and in particular by one or more processor of the database system.

[0339] In block 1410, one or more criteria specifying which feed items are to be displayed to a first user are received from a tenant. In one implementation, the criteria specify which items to add to the custom feed. For example, the criteria could specify to only include feed items for certain fields of a record, messages including certain keywords, and other criteria mentioned herein. In another implementation, the criteria specify which items to remove from the custom feed. For example, the criteria could specify not to include feed items about certain fields or including certain keywords.

[0340] In block 1420, the database system identifies feed items of one or more selected objects that match the criteria. The feed items can be stored in the database, e.g., in one or more of the tables of FIG. 9A. In one implementation, the one or more selected objects are the objects that the first user is following. In another implementation, the one or more selected objects is a single record whose record feed the first user is requesting.

[0341] In block 1430, the feed items that match the criteria are displayed to the first user in the custom feed. The generation of text for a feed tracked update can occur after the identification of the feed items (e.g., data for a field change) and before the display of the final version of the feed item.

[0342] In one implementation, the criteria are received before a feed item is created. In another implementation, the criteria are received from the first user. In one aspect, the criteria may only be used for determining feeds to display to the first user. In yet another implementation, the criteria are received from a first tenant and apply to all of the users of the first tenant. Also, in an implementation where a plurality

of criteria are specified, the criteria may be satisfied for a feed item if one criterion is satisfied.

[0343] Some implementations can provide mechanisms to search for feed items of interest. For example, the feed items can be searched by keyword, e.g., as entered by a user. As another example, a tab (or other selection device) can show feed items about or from a particular user. In one implementation, only messages (or even just comments) from a particular user can be selected.

[0344] In another implementation, a user can enter search criteria so that the feed items currently displayed are searched and a new list of matching feed items is displayed. A search box can be used to enter keywords. Picklists, menus, or other mechanisms can be used to select search criteria. In yet another implementation, feed comments are text-indexed and searchable. Feed comments accessibility and visibility can apply on the search operation too.

[0345] In one implementation, when a user performs a search of feeds, there can be an implicit filter of the user (e.g., by user ID). This can restrict the search to only the news feed of the user, and thus to only record feeds and profile feeds that the user is subscribed. In another implementation, searches can also be done across feeds of users and records that are not being subscribed.

[0346] Besides searching for feed items that match criteria, one also could search for a particular feed item. However, in one implementation, a user cannot directly query a feed item or feed comment. In such an implementation, a user can query to obtain a particular profile or record feed, and then navigate to the feed item (e.g., as child of the parent feed). In another implementation, the relationship from a feed to its parent entity (e.g., a record or user profile) is uni-directional. That is a user can navigate from the feed to the parent but not vice versa.

[0347] In one implementation, a user can directly query the child tables, e.g., comment table 930. Thus, a user could search for comments only that user has made, or comments that contain certain words. In another implementation, a user can search for a profile feed of only one user. In yet another implementation, a user can search for profile feeds of multiple users (e.g., by specifying multiple user names or IDs), which can be combined into a single feed.

[0348] XII. Maintaining Records for Follower's Feeds

[0349] If every feed item is stored and maintained on a follower's feed or even in the profile and/or record feeds, the amount of data to be stored could be massive, enough to cause storage issues in the system. In one implementation, the N (e.g., 50) most recent feed items for each feed are kept. However, there can be a need to keep certain older feed items. Thus, implementations can remove certain feed items, while keeping others. In other implementations, old feed tracked updates may be archived in a data store separate from where recent feed items are stored.

[0350] In some implementations, feeds are purged by a routine (also called a reaper) that can remove items deemed not worthy to keep (e.g., old items). Any underlying data structures from which feed items are created can also be purged. In one implementation, the reaper can remove certain items when new items are added (e.g., after every 5th item added). As another example, feed items may be deleted synchronously during the save operation itself. However, this may slow down each save operation. In one implementation, however, this may be better than incurring a larger cost when the items are removed at longer intervals. In

another implementation, the reaper can run periodically as a batch process. Such routines can ensure that a table size does not become too large. In one aspect, a reaper routine can keep the event history table relatively small so the sharing checks are not extremely expensive.

[0351] In various implementations, the reaper can maintain a minimum number (e.g., 50 or 100) of feed items per record, maintain a minimum number of records per user (e.g., per user ID), and not deleting feed items (or entire records), which have comments against it. Such implementations can ensure that the detail page and profile page have sufficient data to display in a feed. Note that the sharing checks for feed queries can cut down the number of records further for users with less access. Thus, the number of records finally displayed for specific users can be significantly less than a minimum number for a specific profile or record feed. In one implementation, a reaper deletes data that is older than a specified time (e.g., 6 months or a year).

[0352] In one implementation, the reaper can perform the deletion of feed items (purging) as a batch up deletion. This can avoid deletion of large number of records that may lead to locking issues. In another implementation, the reaper can be run often so that the table does not become difficult to manage (e.g., size-wise). In this way the reaper can work on a limited set of records. In one implementation, the reaper may have logic that deletes certain items (e.g., by an identification) from tables (e.g., those in FIG. 9A), or sections of the tables.

[0353] XIII. Informing Users of Relevant Social Network Data

[0354] FIG. 15 shows a flowchart of an example of a method 1500 for informing a user of group data of an online social network when the group data is relevant to the user, performed in accordance with some implementations. For example, one or more keywords can be used to determine whether group data is relevant to a user and, thus, whether the user should be informed of the group data. Keywords can be in the form of individual words, alphanumeric characters, numbers, symbols, and combinations thereof. In some applications, keywords represent topics of interest to users and/or groups of the online social network. For example, as described in greater detail below, a user can be provided with a list of available keywords, from which the user can select according to the user's interests. A user can select keywords at a stage of registration for a user profile in the social network and, in some implementations, later update the selected keywords according to changes in the user's interests. Groups can similarly select and customize keywords during setup of a group and at various stages thereafter. In this way, the particular keywords associated with a user or a group can change over time according to changes in the user's or group's interests.

[0355] In FIG. 15, at block 1504, a computing device such as any of the servers described above is configured to identify one or more keywords as being associated with both a particular group and a particular user of the online social network. That is, when the same one or more keywords has been associated with one or more groups and one or more users, any such users and groups can be identified as being related to one another by virtue of the common keyword. At block 1504, in some instances, the identification of one or more keywords as being associated with a group can occur immediately in response to a user identifying such keyword or keywords in association with the user's profile. For

instance, during an initial registration process to obtain a profile, a user can be prompted to enter one or more keywords representing topics of interest. As soon as such keywords are selected or otherwise identified, the one or more computing devices performing method **1500** can check a list of keywords associated with various groups of the online social network to determine which group or groups may be of interest to the user. Any groups having the same keyword or keywords identified by the user can be presented in a list of names of such groups as part of a user interface during registration or upon completion of registration, for example, in a user interface when the user logs in to the online social network, or by other communication channels such as email or text messaging. By the same token, during group setup, any users having a keyword or keywords identified in association with the group can be flagged and presented to the group leader so the group leader can invite such users to join the group, for example, by posting to the various user's profile feeds or by other communication channels.

[0356] In FIG. **15**, at block **1508**, when one or more keywords have been identified as associated with both a group and a user, a communication can be generated with information identifying the group data. For example, an email, text message in a cellular network, information update posted to a feed accessible by the particular user, automated phone call, tweet, or other communication can be generated. In some instances, the communication includes the group data of interest, while in other instances, the communication describes or otherwise identifies the group data of interest. The group data can be any data associated with the particular group or groups of block **1504**. Examples of group data include posts, comments, record updates, and other various information updates submitted to a group feed, such as the information feed of the XYZ competitive group as shown in FIG. **7**. Other various group data can include records stored in association with the group, new members or existing members of the group, status updates, hyperlinks, groups identified as having similar or the same keywords, and any other data associated with a group.

[0357] In FIG. **15**, at block **1512**, when any of the various types of communications described above are generated at block **1508**, the communication can be transmitted to a display device accessible by the user. For example, when method **1500** is performed by one or more servers, the communication generated at block **1508** can be delivered over an appropriate communications network or combination of networks to a user system **12** operated by the user identified at block **1504** to be displayed on a display of the user system. Thus, when the communication is in the form of a text message, the communication can be sent over a wireless network to a smartphone or tablet operated by the user. In other instances, when the communication is in the form of an information update, the communication can be posted to the user's profile feed. When the communication is in the form of an email, the email can be sent over the Internet to an email address associated with the user, and so forth. In some other implementations, method **1500** is performed at a computing device operated by the user. In this case, the communication can be generated at block **1508** and provided at block **1512** at the same computing device at which block **1504** is performed.

[0358] FIG. **16** shows a flowchart of an example of a method **1600** for informing a user of group data of an online

social network when the group data is relevant to the user, performed in accordance with some implementations. In method **1600**, one or more keywords have been selected and associated with a user's profile of the online social network. Similarly, one or more keywords have been associated with a group of the online social network. Examples of the selection, updating, and customization of keywords to be associated with users' profiles and groups are described in greater detail below. In some instances, one or more of the same keywords may be associated with both the user's profile and the group while, in other instances, there is no overlap among keywords associated with the respective user's profile and group.

[0359] In FIG. **16**, at block **1602**, a computing device such as a server or user system is configured to retrieve any keywords associated with a user's profile. For instance, such keywords can be stored on a suitable storage medium and linked with or otherwise flagged as being associated with a particular user's profile. In some instances, such keywords are stored as part of the data of a user's profile. At block **1604**, any keywords associated with a group of the online social network are also retrieved from one or more storage mediums. Again, such keywords can be stored as part of the group data or, in other instances, stored on a suitable storage medium and linked with, flagged, or otherwise identified as being associated with the group. Blocks **1602** and **1604** can be performed to retrieve keywords associated with more than one user's profile and more than one group of the online social network.

[0360] Any number of groups can exist in a social network environment. For instance, an organization can have various groups, and such an organization can be one of many tenants in a multi-tenant database environment, where the various organizations have respective groups. The storage of group data of an organization is generally described above, and various database implementations are possible to manage and link keywords with respective users and groups. In one example, a database table or set of tables is configured to store user profile data. Another table or set of tables is configured to store group data, and a third database table is configured to store a list of keywords. The database table or tables storing user profiles can have pointers to particular entries in the keyword database table, and the group database tables can be similarly configured with separate pointers to selected keywords in the keyword database table. Other various implementations for storing and linking such data are possible, as will be understood by those skilled in the art.

[0361] In FIG. **16**, at block **1606**, when keywords of a particular group and user's profile have been retrieved, any keywords associated with both the user and the group are identified. When any keywords match, at block **1608**, a first communication can be generated as generally described above at block **1508** of method **1500**. In this example, the first communication includes content indicating both the group data and the fact that one or more keywords were identified as being associated with the user and the group. In some instances, the content of the first communication identifies the particular keyword or keywords. At block **1612**, the first communication generated at block **1608** is provided to a display device accessible by the user, as generally described above at block **1512** of method **1500**.

[0362] In the example of FIG. **16**, at block **1616**, a second communication can be received from the user. This second

communication can take various forms similar to the first communication of blocks **1608** and **1612**. For example, when the first communication is an email, the second communication can be a reply email. By the same token, when the first communication is a post to the user's profile feed, the second communication can be a comment on such a post. In this example, the second communication includes information requesting that the user receive information updates associated with the group. For example, the second communication can request that the user be admitted as a member of the group or request that the user begin following the group, so the user is able to be notified of any future information updates submitted to the group feed.

[0363] In FIG. **16**, at block **1620**, the one or more computing devices performing method **1600** is configured to determine whether the user's request of block **1616** should be granted. That is, in the preceding examples, the one or more computing devices can automatically determine that the user is to be admitted as a new member of the group or automatically determine that the user can subscribe to the group, thus including the user as a recipient of any future email notifications of various information updates submitted to the group feed. In some instances, following block **1616**, it is determined that the user is not allowed to subscribe to the group or otherwise receive information updates associated with the group, for any of various reasons.

[0364] In FIG. **16**, at block **1624**, when it is determined that the user will be permitted to receive information updates, any such information updates associated with the group are provided for access by the user. For example, any group-related updates can be reported to the user in the form of an email daily digest or automatically included in the user's profile feed.

[0365] In some examples, the selection and attachment of keywords to a user or group can be determined according to a number of instances that such keywords are included in an information feed associated with a group. For instance, returning to FIG. **7**, the name "Bill" occurs at least three times, that is, as the author of post **710**, as the author of one of the comments **730**, and as the subject of Ella Johnson's comment. In such cases, the name "Bill" can be identified as occurring more than or equal to a specified number representing a threshold of relevance, such as the number '3' in this example. In such cases, an additional communication can be generated by the one or more computing devices performing any of the various methods described herein to notify a user or group leader that such a keyword may be relevant and of interest to attach or otherwise associate with the user's profile or group. In some instances, the user or group leader can be provided with an option of selecting such keywords and authorizing the association of the keywords. Any such communications identifying keywords of interest can be communicated to various computing devices operated by a group leader or user. Such communications can take the various forms mentioned above with respect to the first and second communications of method **1600**.

[0366] FIG. **17** shows a flowchart of an example of a method **1700** for informing a first user of data of an online social network when the data is relevant to the first user, performed in accordance with some implementations. Various social network data in addition to the group data described above with reference to FIGS. **15** and **16** can be pushed to users, such as various information updates communicated to or from various users, groups, records, and

other entities, and identifications of relevant users, groups, records, and related content to users having profiles deemed relevant to the social network data, and vice versa.

[0367] In the example of FIG. **17**, at block **1704**, social network data associated with an entity of the online social network is received at one or more computing devices performing method **1700**. In this example, the first user is not following the entity with which the social network data is associated at the time block **1704** occurs. For example, the first user may not be aware of the data, for example, when the data is associated with a record or a user's profile, which the user has not subscribed to or friended. The social network data can have any of various types of content such as information updates, records, users' profiles, data describing such items, and combinations thereof. The social network data can be associated with any combination of various entities. For example, the social network data can be associated with one or more groups, one or more user profiles, and/or one or more records stored in a multi-tenant database environment. Such entities can be identified by headers, content, or links attached to the social network data, depending on the type of social network data and context in which the social network data appears.

[0368] FIG. **18** shows an example of a group page **2000** in the form of a graphical user interface (GUI) configured to provide group data and be accessible by various members of the group, according to some implementations. The GUI shows another state of the XYZ Competitive Group page of FIG. **7**, where a keywords region **2004** is presented as part of page **2000** and displays particular keywords associated with the XYZ Competitive Group. In this example, keywords currently associated with the XYZ Competitive Group include "XYZ", and "Cirrus Computers". As mentioned above, the keywords associated with a group can change over time, and an update button **2008** allows a group leader, system administrator, or authorized member of the group to modify the list of keywords in region **2004**. For example, clicking on update button **2008** with a user input device such as a mouse can cause a pane to open, allowing a user to enter additional keywords, delete existing keywords, and/or select from a list of authorized keywords, depending on the particular implementation. Such changes in the keyword associations can then be stored and listed in keywords region **2004** of page **2000**.

[0369] In addition, in FIG. **18**, information updates presented as feed items in group feed **2012** can have an associated "Tag" selection **2016** or similar selection allowing a user to select individual updates in the feed with which to associate keywords. In this example, in which page **2000** is presented to Parker Harris, Parker is provided with a "Tag" selection **2016**. Clicking on the tag selection **2016** causes a keywords pop-up window **2020** to be generated. The window **2020** includes a list of keyword selections **2024**, any of which Parker can click on to select and link with Parker's post **720**. In addition, keywords window **2020** includes a custom data entry field **2028**, in which Parker can enter custom keywords to attach to post **720**, such as the words "Bill" and "Saxon". In this example, only the author of a post, such as Parker Harris, is provided with the capability of tagging his post with appropriate keywords. In other examples, other users can tag any of the various posts in an information feed with appropriate keywords. Thus, individual information updates in a given feed can be identified as relevant to any of various users and other

entities of the online social network and communicated to any such users having matching keywords, as generally described herein. By the same token, in some instances, the group keywords in region **2004** can be customized and edited by authorized group members or, in some instances, any of various users of the online social network.

[0370] Returning to FIG. 17, at block **1706**, the one or more computing devices performing method **1700** are configured to retrieve one or more keywords from a suitable storage medium such as a database. At block **1706**, the retrieved keywords can be associated with one or more users' profiles, a group, an information update, a record, etc. For instance, the retrieved keywords can be mentioned in a user's job description or a user's biographical information in the user's profile.

[0371] In FIG. 17, at block **1708**, the one or more computing devices performing method **1700** are configured to identify the social network data of interest as being related to one or more of the keywords retrieved at block **1706**. For example, returning to FIG. 18, when Parker Harris has submitted his post **720** to group feed **2012** and tagged post **720** with keywords using pop-up window **2020**, Parker's post **720** can be identified as being related to the one or more keywords of block **1706** when such keywords match Parker's selections in pop-up window **2020**. When one or a combination of such keywords associated with Parker's post **720** match the keywords retrieved at block **1706**, method **1700** proceeds to block **1712** at which the content of Parker's post can be sent to a display device in a presentation accessible by another user having a profile associated with the keywords retrieved at block **1706**, by way of example. The social network data can be communicated in various forms as described above with reference to blocks **1508** and **1512** of method **1500** and blocks **1608** and **1612** of method **1600**. Thus, by way of example, Parker Harris' post **720** can be included in another user's profile feed.

[0372] FIG. 19 shows a flowchart of an example of a method **1900** for informing a user of group data of an online social network when the group data is relevant to the user, performed in accordance with some implementations. As generally described above, in some implementations, a user can identify and associate keywords with the user's profile to identify the user as interested in various topics or categories represented by the keywords. When groups or other entities in the online social network have data associated with the same or similar keywords, the interested user can be notified of such data without having to subscribe to the entity. Thus, for example, information updates in any other user's profile feed or a group feed that is possibly related to the user's keywords can be communicated to the user.

[0373] In some implementations, the identification, selection, association, updating, and deletion of keywords for a group or other entity of possible interest to a user is carried out in an asynchronous manner with respect to the identification, selection, association, updating, and deletion of keywords associated with the user's profile. That is, regardless of the status of keyword association with a group, at any given time, a user can initiate a keyword selection process or, after being registered, update the keywords attached to that user's profile. Similarly, groups can be created with associated keywords, and those keywords can be updated regardless of the status of various users in the online social network.

[0374] In FIG. 19, at block **1904**, a computing device is configured to associate one or more group keywords with a particular group of the online social network. Returning to FIG. 18, after a group is created and any keywords have been associated with the group, a master list of keywords can be maintained. As mentioned above, keywords region **2004** lists part or all of the associated keywords. Such keywords can be updated over time, for instance, by deleting certain keywords that are no longer pertinent to the group's purpose or by adding keywords that later become interesting based on conversations happening in the group feed **2012** or any changes to the purpose or description of the group. In FIG. 18, the update button **2008** can be selected to modify the keywords associated with the XYZ Competitive Group, by way of example.

[0375] FIG. 20 shows an example of a group setup window **2100** for generating and defining information about a particular group, according to some implementations. The group setup window **2100** can be presented in a user interface to a user when the user initiates creation of the group in the online social network, or it can be presented after the group has been established. FIG. 20 represents one of many possible examples of an interface for obtaining and defining group information when the user wishes to change any of the group attributes.

[0376] In the example of FIG. 20, the group setup window **2100** includes a number of fields, any one of which can be used to identify or trigger the selection and identification of keywords to be associated with the group. In this example, group setup window **2100** includes a name field **2104**, in which a user can enter any combination of letters, symbols, other characters, and words to define the name of the group. In addition, a description field **2108** allows a user accessing the group setup window **2100** to enter words and phrases describing the purpose of the group and any other relevant information about the group to facilitate an understanding of the group by other users. In this example, the XYZ Competitive Group identified in field **2104** is described as "a place to share information that will allow us to compete against XYZ more effectively", as show in field **2108**. Members of the group can be identified in field **2112** and, in this example, include Parker Harris and Joe Olsen, among other users.

[0377] In the example of FIG. 21, system-generated keywords of possible interest to the user accessing the group setup window **2100** are generated and listed in regions **2116** and **2120**. In this example, the combination of characters, "XYZ", has generated keywords of possible relevance to "XYZ". The identification and listing of keywords such as "PC", "tablets", and "cloud computing", can be based on possible keywords in various sources and can include various schemes for determining which keywords to present in region **2116**. Similarly, the keywords listed in region **2120** are identified and selected based on the occurrence of the word "Compete" in field **2108**. Some of the various techniques for identifying and selecting keywords to be presented in regions **2116** or **2120** are described in greater detail below.

[0378] In FIG. 20, a custom keyword region **2124** includes a data entry field **2128** in which a user can type in any characters or combinations of characters, words, phrases, symbols, etc. to associate with the XYZ Competitive Group. In some implementations, a prerequisite to a group being established is that a user at least select or enter in field **2128**

keywords of interest to the user. In this way, any groups existing in the online social network can have at least one keyword and thus be identified using the techniques herein.

[0379] Similar setup windows as that shown in FIG. 21 can be used to define other entities than groups. For instance, records, users, and other various objects can be named, described, and otherwise defined with data to establish one or more keywords associated with the entity.

[0380] In some implementations, the keywords relevant to a group may change as the group changes over time. For instance, the initial description and purpose of a group entered by the group leader in group setup window 2100 can evolve as one or more conversations develop in group feed 2012 of group page 2000. Some implementations provide for the automatic selection and updating of additional keywords to track any changes to the group or group conversations. For example, some implementations provide for identifying words, phrases, symbols or other characters in the information updates of a given feed, such as group feed 2012, and identifying any such data that occurs with a designated frequency or having a count of occurrences meeting or exceeding a designated threshold. For example, when any word appears in the group feed more than the designated threshold, a pop-up window can be generated and presented in the user interface with group page 2000. Such a pop-up window can prompt the user to select any such data as a keyword to be associated with the group. Thus, in the example of FIG. 18, when the acronym “SVUC” appears more than three times, such a pop-up window can be generated to suggest that SVUC be added as a keyword associated with the XYZ Competitive Group. If the user confirms that such a keyword is to be added, the list of keywords in region 2004 can be updated accordingly.

[0381] In some other implementations, a pop-up window or other similar prompt for user input can be omitted, and the one or more computing devices performing method 1900 can automatically update the list of keywords in region 2004 based on various criteria, such as the occurrence of the keyword more than a designated number of times in the feed items of group feed 2012. The user to be prompted whether to add a keyword to an entity such as a group is often the group leader or another user having specifically defined permissions to access and update attributes of the entity such as a group. In some implementations, a system administrator is the only person provided with the capability of adding or removing keywords from the list associated with a particular entity.

[0382] The updating of keywords associated with an entity such as a group can be initiated manually or scheduled at various times. For example, an update procedure can be scheduled to be performed by a server once every two weeks at a scheduled time to cause the list of keywords associated with an entity such as a group to automatically be updated. For instance, code can be executed to automatically scan all of the posts, comments, record updates, and other various data in the group feed 2012 of the example in FIG. 18 and identify any words, phrases, symbols, or other characters that occur more than a designated number of times among other criteria. Another example of a criterion is the number of other users who have replied to an information update having one or more words, phrases, symbols, or other characters of possible interest to be added to the keyword list. In addition, such a procedure can automatically scan the existing list of keywords associated with a group to deter-

mine whether a possible new keyword is already included. In this case, any additional operations to prompt a user for selection or to otherwise screen the possible keyword can be skipped.

[0383] In FIG. 19, at block 1906, a first communication is received by the one or more computing devices performing method 1900. This first communication can take various forms, such as email, text messages, posts to an information feed, tweets, data entered into a user interface, etc. In this example, the first communication includes one or more user keywords, that is, keywords selected or identified by a user or in association with a user’s profile. At block 1906, one or more keywords can be received to be associated with a user’s profile at various stages of an online social network. For example, keywords can be received and attached to a user’s profile during an initial registration process for the user to obtain a login and password to access the online social network. In other instances, keywords can be received after a user’s profile has been established to update the keywords associated with that profile.

[0384] FIG. 21 shows an example of a new user setup window 2200 as part of a registration process for a new user to obtain a user’s profile in an online social network, according to some implementations. In this example, a user can enter a desired login or ID in data entry field 2204 and can select a password in field 2208. In this example, a job description field 2212 allows a user to type words, phrases, symbols and other characters pertaining to the user’s job in an organization or group within an organization. In addition, a “bio” field 2216 allows the user to enter biographical information.

[0385] In this example, similar to the functionality of group setup window 2100, interest regions 2220 and 2224 are generated to include suggested keywords of possible interest to the new user based on the identification of any data entered in the various fields 2204-2216. In this example, the mention of the words “patent” and “legal” in job description field 2212 has triggered an automatic system-generated retrieval of suggested keywords related to those respective words. Thus, in this example, the suggested keywords of “enforcement”, “prosecution”, “defense”, and “licensing” have been proposed as possible keywords relevant to the word “patent”. Keywords possibly relevant to the term “legal” are similarly retrieved and set forth in region 2224. A custom keyword region 2228 has a data entry field 2232 allowing a user to enter any keywords representing topics of interest to the user. In this way, a user can customize the user’s experience in an online social network to leverage the methods disclosed herein to obtain social network data of interest to the user.

[0386] In some implementations, it is desirable to have a defined list of keywords for possible selection by groups, users, and other entities to be associated therewith. For instance, a designated set of keywords can simplify the management and processing of information to align users with groups and other entities. Thus, in FIGS. 20 and 21, in some implementations, keywords listed in the respective regions 2116 and 2120, as well as region 2220 and 2224, can be pulled from the same master list of possible keywords. For example, the possible keywords presented to a user in new user setup window 2200 can be retrieved from a master list of keywords already associated with one or more groups of the online social network. Thus, when groups have identified themselves as maintaining and discussing data

related to certain topics, users can access data related to any of those topics in the context of a social network.

[0387] In FIG. 21, in some implementations, the identification, selection, and generation of keywords to be presented to a user for possible association with the user's profile can be performed similar to the techniques described above for identifying keywords to be associated with groups. In some implementations, the keywords suggested to a user for association with the user's profile can be based on other data and actions associated with the particular user. For instance, all of the posts and comments submitted by a particular author can be filtered to identify frequent words or phrases used by the particular user that can be presented to the user as a possible keyword of interest. Any keywords selected by a user in an interface such as new user setup window 2200 can be linked with that user's profile and used to identify groups and other entities, as well any social network data having the same keywords. Such information can then be pushed to the user as described herein.

[0388] In some instances, an administrator or other user than the user setting up his or her profile is granted sufficient access permissions to add or update the keywords associated with the user's profile. In some instances, new user setup window 2200 is generated and displayed in a user interface to an administrator, who sets up the user's profile on that user's behalf. After the user's profile is established, the particular user can access and modify the various information of his or her profile to add or delete keywords as desired, for example, using a similar window 2200 as shown in FIG. 21. This can be helpful in situations where the user's interests change over time as the user becomes more involved with certain projects or groups within an organization and, by the same token, less interested in other entities. Thus, at any given time during the user's membership in the online social network, the user can freely update keywords associated with that user's profile to attempt to access information and connect with entities of most relevance to that user.

[0389] In other implementations, keywords are automatically identified and associated with groups or users' profiles by one or more computing devices. For example, keywords can be suggested to a user based on the keywords associated with the profiles of other users in that particular user's role, team, department, manager, and any entities, which the particular is following. For instance, a new employee could be notified of groups that fellow team members frequently post to or otherwise interact with. As the new user becomes involved and participates in conversations in various groups and regarding various entities, the keywords associated with that user's profile can change accordingly over time. As mentioned above, various criteria can be used to identify keywords or limit the number of keywords suggested to a particular user or automatically associated with that user's profile. Criteria for selecting keywords for association with a user's profile include a frequency of visits by the particular user to a designated group or other entity in the online social network, that is, whether the frequency meets or exceeds a designated threshold, as well as criteria such as how often or the total number of times a particular keyword has been used in that user's posts and comments submitted to feeds of the online social network.

[0390] In some implementations, the identification of social network data to provide to a user or the presentation of keywords selections for a user to attach to the user's

profile can include the retrieval of keywords associated with one or more groups of the online social network. In some instances, the keywords ultimately used or presented to the user are a subset the group keywords. For instance, in some cases, the only keywords identified are common to two or a designated number of groups. In some cases, keywords of possible interest to a user are scored according to the number of times that those keywords appear in a particular group feed or appear in association with various entities in the online social network. Thus, a relevance measure can be based on such a score, with the keywords being presented to the user in a ranked list, that is, with the keywords having the highest score at the top of the list. Returning to FIG. 21, by way of example, the keywords in region 2220 can be listed in such a manner.

[0391] In other examples, the act of joining or following a group, or the act of following any other social network entity may cause one or more keywords associated with the entity to be automatically associated with the user's profile. In some cases, when keywords are scored according to a relevance measure, keywords having a score higher than a designated threshold can be automatically associated with the user's profile.

[0392] In some examples, the score of a keyword in terms of a relevance measure is determined according to the number of times the keyword is used in a conversation. For instance, a conversation could have fifty posts and twenty comments to selected posts. A word or phrase mentioned once in the entire conversation would have a lower score than a word or phrase mentioned twenty times throughout the conversation. In some cases, any keywords mentioned fewer times than a designated threshold are automatically disregarded and not considered as suggestions or automatic additions to a user's profile. Another factor or criterion to determine the relevance, identification, or association of a particular keyword can be the age of the group with which the keyword is associated. For instance, keywords associated with groups more than a year old can be deemed more relevant than keywords associated with a group created one week ago. When several criteria are applied to determine a relevance measure for relevance, identification, or association of keywords with respect to a user or group, each criterion can have an associated score which is summed with the score of other criteria to determine a total score to be compared with the scores of other criteria and thus rank the keywords according to total relevance.

[0393] In some instances, an automated process periodically suggests new keywords for association with a group or with any user following the group based on the content of the conversation in the group feed. The same is true for entities other than groups in the online social network. For example, a group could initially be dedicated to the discussion of user interfaces and, over time, have a conversation in the group feed that shifts to a discussion of knitting or white water rafting. By performing an automated process to update the keywords at designated intervals, the contents of the various posts and comments can be filtered to identify the more recent and repeated subjects of knitting and white water rafting as suggested keywords to associate with the particular entity. For instance, when the word "knitting" appears more than ten times, a pop-up window can be generated and caused to appear in the user interface of the group leader and any users accessing the group page. The pop-up window can prompt the group leader or user to

decide whether to associate the new keyword with the group, the user's profile, or other entities.

[0394] Returning to FIG. 19, at block 1908, the one or more keywords received at block 1906, for instance, from one or more computing devices configured to automatically identify and select keywords or from one or more users choosing the keywords, are associated with the user's profile.

[0395] In FIG. 19, at block 1912, the one or more computing devices performing method 1900 are configured to determine whether the one or more group keywords of block 1904 include the one or more user keywords of blocks 1906 and 1908. When there is overlap among any group keywords and any user keywords, at block 1916, a second communication is provided including an identification of the group data to a display device for display to the user. That is, for example, any posts or comments in a group feed, as well as any other information updates appearing in the group feed and any documents or files uploaded to the group can be communicated to the user. In some cases, a description of the group data is provided to the user, for example, in terms of a post to the user's profile feed, an email to the user's email address, a text message to the user's cell phone number, as well as other various communication options described above. In some cases, the communication of block 1916 includes the content of the group data itself, for instance, with the content of a relevant post embedded in an email or the document of interest attached to an email sent to a user.

[0396] In some examples, the second communication of block 1916 can be in the form of an email, for instance, sent to the user as a daily digest with a summary of the names of groups having conversations in which the user's keywords appear. By the same token, the digest can identify any groups in which the conversations no longer include any keywords of interest to the user. Thus, on a daily or other periodic basis, a user can choose to follow new groups of interest or unsubscribe from groups no longer of interest to the user. By the same token, one or more computing devices can be configured to send a communication in the form of an email or other format to a group leader including identifications of any users' profiles having keywords matching those associated with the group. In this way, a group leader could identify and invite selected users to join or follow the group. In other cases, a group leader could be notified of any users mentioning group keywords in the user's posts and other messages submitted in the online social network that include the keywords.

[0397] The specific details of the specific aspects of implementations disclosed herein may be combined in any suitable manner without departing from the spirit and scope of the disclosed implementations. However, other implementations may be directed to specific implementations relating to each individual aspect, or specific combinations of these individual aspects.

[0398] While the disclosed examples are often described herein with reference to an implementation in which an on-demand database service environment is implemented in a system having an application server providing a front end for an on-demand database service capable of supporting multiple tenants, the present implementations are not limited to multi-tenant databases nor deployment on application servers. Implementations may be practiced using other data-

base architectures, i.e., ORACLE®, DB2® by IBM and the like without departing from the scope of the implementations claimed.

[0399] It should be understood that some of the disclosed implementations can be embodied in the form of control logic using hardware and/or using computer software in a modular or integrated manner. Other ways and/or methods are possible using hardware and a combination of hardware and software.

[0400] Any of the software components or functions described in this application may be implemented as software code to be executed by a processor using any suitable computer language such as, for example, Java, C++ or Perl using, for example, conventional or object-oriented techniques. The software code may be stored as a series of instructions or commands on a computer-readable medium for storage and/or transmission, suitable media include random access memory (RAM), a read only memory (ROM), a magnetic medium such as a hard-drive or a floppy disk, or an optical medium such as a compact disk (CD) or DVD (digital versatile disk), flash memory, and the like. The computer-readable medium may be any combination of such storage or transmission devices. Computer-readable media encoded with the software/program code may be packaged with a compatible device or provided separately from other devices (e.g., via Internet download). Any such computer-readable medium may reside on or within a single computing device or an entire computer system, and may be among other computer-readable media within a system or network. A computer system, or other computing device, may include a monitor, printer, or other suitable display for providing any of the results mentioned herein to a user.

[0401] While various implementations have been described herein, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of the present application should not be limited by any of the implementations described herein, but should be defined only in accordance with the following and later-submitted claims and their equivalents.

1-20. (canceled)

21. A system for supporting dynamic membership of teams of users in an organization, the system comprising:

a database system implemented using a server system comprising one or more hardware processors, the database system configurable to cause:

sending, to a client device, rule builder data capable of being processed to cause display of a rule builder user interface at the client device, the rule builder user interface comprising one or more selectable options for configuring a rule;

obtaining, from the client device, rule configuration data generated using the rule builder user interface;

configuring the rule according to the rule configuration data;

processing the rule, the processing of the rule comprising:

obtaining a specific criterion or criteria stored in a database,

applying the specific criterion or criteria to a plurality of user attributes stored in a database, the user attributes identifying characteristics of a plurality of users in an organization,

determining that one or more of the user attributes satisfies the criterion or criteria, and

- identifying one or more of the users as having the one or more user attributes;
- defining membership of a team of users in the organization according to the rule, the defining of the membership of the team of users comprising adding or removing the one or more users to or from the team of users; and
- broadcasting a communication to the team of users.
- 22.** The system of claim **21**, the database system further configurable to cause:
- detecting one or more changes in the plurality of user attributes, the detecting of the one or more changes causing the processing of the rule.
- 23.** The system of claim **22**, wherein the one or more changes in the plurality of user attributes indicates that a user has been added to or removed from a group of users.
- 24.** The system of claim **21**, the database system further configurable to cause:
- validating one or more properties, operators, and/or values of the rule.
- 25.** The system of claim **21**, wherein the criterion or criteria specifies that the one or more user attributes starts with a designated string, character, word, phrase or symbol.
- 26.** The system of claim **21**, wherein the communication is broadcast on a channel of the team of users, the channel dedicated to a specific topic, department, or project of the organization, the channel comprising conversations, meetings and/or files.
- 27.** The system of claim **21**, wherein the communication is broadcast via one or more of email, text message, phone call, tweet or post to a feed.
- 28.** A computer program product comprising computer-readable program code capable of being executed by one or more processors when retrieved from a non-transitory computer-readable medium, the program code comprising instructions configurable to cause:
- sending, to a client device, rule builder data capable of being processed to cause display of a rule builder user interface at the client device, the rule builder user interface comprising one or more selectable options for configuring a rule;
- obtaining, from the client device, rule configuration data generated using the rule builder user interface;
- configuring the rule according to the rule configuration data;
- processing the rule, the processing of the rule comprising:
- obtaining a specific criterion or criteria stored in a database,
- applying the specific criterion or criteria to a plurality of user attributes stored in a database, the user attributes identifying characteristics of a plurality of users in an organization,
- determining that one or more of the user attributes satisfies the criterion or criteria, and
- identifying one or more of the users as having the one or more user attributes;
- defining membership of a team of users in the organization according to the rule, the defining of the membership of the team of users comprising adding or removing the one or more users to or from the team of users; and
- broadcasting a communication to the team of users.
- 29.** The computer program product of claim **28**, the instructions further configurable to cause:
- detecting one or more changes in the plurality of user attributes, the detecting of the one or more changes causing the processing of the rule.
- 30.** The computer program product of claim **29**, wherein the one or more changes in the plurality of user attributes indicates that a user has been added to or removed from a group of users.
- 31.** The computer program product of claim **28**, the instructions further configurable to cause:
- validating one or more properties, operators, and/or values of the rule.
- 32.** The computer program product of claim **28**, wherein the criterion or criteria specifies that the one or more user attributes starts with a designated string, character, word, phrase or symbol.
- 33.** The computer program product of claim **28**, wherein the communication is broadcast on a channel of the team of users, the channel dedicated to a specific topic, department, or project of the organization, the channel comprising conversations, meetings and/or files.
- 34.** The computer program product of claim **28**, wherein the communication is broadcast via one or more of email, text message, phone call, tweet or post to a feed.
- 35.** A method comprising:
- sending, to a client device, rule builder data capable of being processed to cause display of a rule builder user interface at the client device, the rule builder user interface comprising one or more selectable options for configuring a rule;
- obtaining, from the client device, rule configuration data generated using the rule builder user interface;
- configuring the rule according to the rule configuration data;
- processing the rule, the processing of the rule comprising:
- obtaining a specific criterion or criteria stored in a database,
- applying the specific criterion or criteria to a plurality of user attributes stored in a database, the user attributes identifying characteristics of a plurality of users in an organization,
- determining that one or more of the user attributes satisfies the criterion or criteria, and
- identifying one or more of the users as having the one or more user attributes;
- defining membership of a team of users in the organization according to the rule, the defining of the membership of the team of users comprising adding or removing the one or more users to or from the team of users; and
- broadcasting a communication to the team of users.
- 36.** The method of claim **35**, further comprising:
- detecting one or more changes in the plurality of user attributes, the detecting of the one or more changes causing the processing of the rule.
- 37.** The method of claim **36**, wherein the one or more changes in the plurality of user attributes indicates that a user has been added to or removed from a group of users.
- 38.** The method of claim **35**, further comprising:
- validating one or more properties, operators, and/or values of the rule.
- 39.** The method of claim **35**, wherein the criterion or criteria specifies that the one or more user attributes starts with a designated string, character, word, phrase or symbol.

40. The method of claim 35, wherein the communication is broadcast on a channel of the team of users, the channel dedicated to a specific topic, department, or project of the organization, the channel comprising conversations, meetings and/or files.

* * * * *