

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
27 March 2008 (27.03.2008)

PCT

(10) International Publication Number  
**WO 2008/036464 A2**

(51) International Patent Classification:

**G06F 15/173** (2006.01)

(21) International Application Number:

PCT/US2007/074475

(22) International Filing Date: 26 July 2007 (26.07.2007)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

60/820,481	26 July 2006 (26.07.2006)	US
60/820,485	27 July 2006 (27.07.2006)	US
60/835,570	4 August 2006 (04.08.2006)	US
60/822,551	16 August 2006 (16.08.2006)	US
60/823,767	29 August 2006 (29.08.2006)	US
60/823,780	29 August 2006 (29.08.2006)	US
60/862,004	18 October 2006 (18.10.2006)	US
60/862,600	23 October 2006 (23.10.2006)	US
11/557,271	7 November 2006 (07.11.2006)	US
60/866,864	22 November 2006 (22.11.2006)	US
60/872,118	1 December 2006 (01.12.2006)	US
60/868,548	5 December 2006 (05.12.2006)	US
60/884,667	12 January 2007 (12.01.2007)	US
60/887,316	30 January 2007 (30.01.2007)	US

60/890,813 20 February 2007 (20.02.2007) US

60/914,107 26 April 2007 (26.04.2007) US

60/914,228 26 April 2007 (26.04.2007) US

11/750,301 17 May 2007 (17.05.2007) US

60/950,726 19 July 2007 (19.07.2007) US

(71) Applicant (for all designated States except US): **NEWS-LIKE MEDIA GROUP, INC.** [US/US]; 78 Winter Street, Lincoln, Massachusetts 01773 (US).

(72) Inventors; and

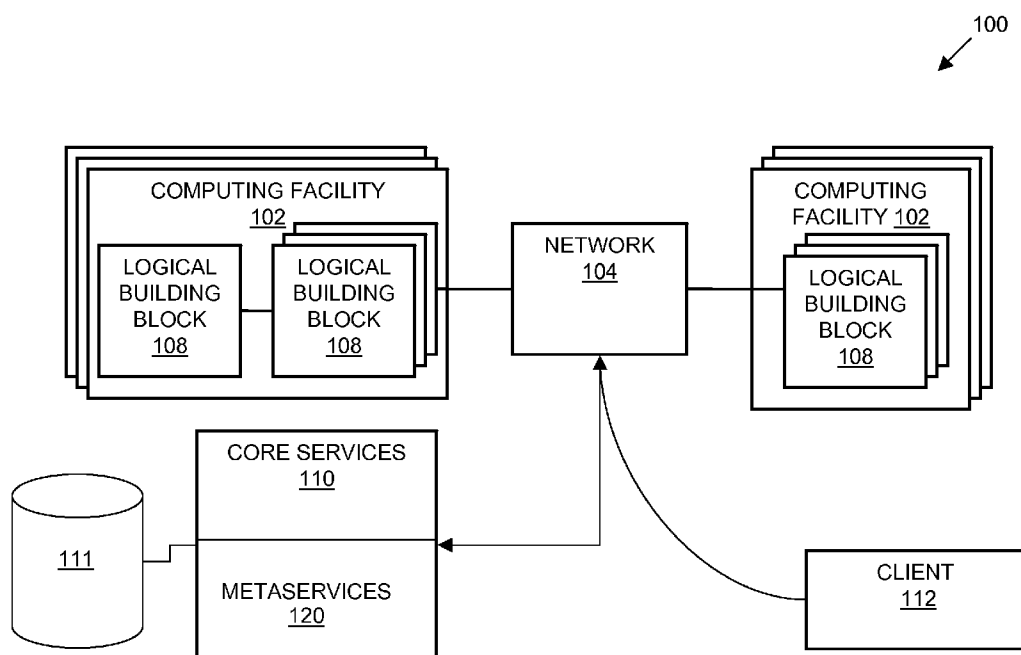
(75) Inventors/Applicants (for US only): **MOORE, James** [US/US]; 78 Winter Street, Lincoln, Massachusetts 01773 (US). **WOOD, Charles** [US/US]; 7101 Foxtree Cove, Austiin, Texas 78750 (US).

(74) Agent: **NORTRUP, John**; Strategic Patents, c/o Intelle-vate, P.O. Box 52050, Minneapolis, Minnesota 55402 (US).

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

[Continued on next page]

(54) Title: SYNDICATION-BASED APPLICATION CONNECTORS



(57) Abstract: Application connectors use a syndicated format for exchanging data between platforms such as calendaring systems and other local applications and/or web-based applications. In addition, connectors may be employed to enrich data, synchronize data, or otherwise manipulate data forwarded to a platform. In another aspect, advertising is improved by permitting real-time deployment of advertisements in response to detected surges in web-based activity, and by permitting insertion of advertisements during data exchange using the application connectors above.



PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY,  
TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA,  
ZM, ZW.

FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL,  
PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM,  
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**(84) Designated States** (*unless otherwise indicated, for every  
kind of regional protection available*): ARIPO (BW, GH,  
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,  
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),  
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

**Published:**

- *without international search report and to be republished  
upon receipt of that report*
- *with information concerning one or more priority claims  
considered void*

## Syndication-Based Application Connectors

## RELATED APPLICATIONS

[0001] This application claims priority to the following commonly owned U.S. Patent Applications, each of which is hereby incorporated by reference in its entirety: U.S. Patent App. No. 11/557,271 filed on November 7, 2006; U.S. Patent App. No. 11/380,923 filed on November 23, 2006; U.S. Patent App. No. 11/750,301 filed on May 17, 2007; U.S. Provisional App. No. 60/820,481 filed on July 26, 2006; U.S. Provisional App. No. 60/820,485 filed on July 27, 2006; U.S. Provisional App. No. 60/835,570 filed on August 4, 2006; U.S. Provisional App. No. 60/822,551 filed on August 16, 2006; U.S. Provisional App. No. 60/823,767 filed on August 29, 2006; U.S. Provisional App. No. 60/823,780 filed on August 29, 2006; U.S. Provisional App. No. 60/862,004 filed on October 18, 2006; U.S. Provisional App. No. 60/862,600 filed on October 23, 2006; U.S. Provisional App. No. 60/866,864 filed on November 22, 2006; U.S. Provisional App. No. 60/868,548 filed on December 5, 2006; U.S. Provisional App. No. 60/872,118 filed on December 1, 2006; U.S. Provisional App. No. 60/884,667 filed on January 12, 2007; U.S. Provisional App. No. 60/887,316 filed on January 30, 2007; U.S. Provisional App. No. 60/890,813 filed on February 20, 2007; U.S. Provisional App. No. 60/914,107 filed on April 26, 2007; U.S. Provisional App. No. 60/914,228 filed on April 26, 2007; and U.S. Provisional App. No. 60/950,726 filed on July 19, 2007.

## BACKGROUND

[0002] The radical proliferation of syndicated content on the World Wide Web, followed more recently by the emergence of so-called Web 2.0 services, evidences a strong desire within the Internet user community for a self-defined environment. In one sense, the use of syndicated content and services such as mashups is highly personalized. Each use can define a very specific, individual view of a universe of syndicated content, and offer customized services built upon third-party programming interfaces. In another sense, the syndication environment is highly collaborative and participatory, with communities of interest arising quickly and assembling around topics or points of view, and sometimes dissipating just as quickly. In either case, content and services change rapidly according to user activity.

**[0003]** Numerous approaches to dynamically deploying services across a network have been devised such as Microsoft's .NET technology, or web services using, e.g., the Web Services Description Language ("WSDL"). While these technology infrastructures support discovery and use of services across a network and can accommodate an evolving set of network-accessible services, would-be users are constrained by the corresponding conceptual and syntactic frameworks. A current trend in Internet services appears to be sharply diverging from this canned approach. Instead, various web-accessible programming interfaces are being published for general use, and these interfaces are discovered and combined on an ad hoc basis by end users. As new programming interfaces and services appear, additional uses and combinations are recognized and deployed in "mashups" that can in turn be republished as web sites or new programming interfaces.

**[0004]** There remains a need for improved tools and techniques to interconnect the variety of web-based resources, web content, and local applications that are currently available. There also remains a need for improved advertising techniques responsive to current web usage patterns.

## SUMMARY

**[0005]** A calendar synchronization system uses syndicated data to transfer calendar items among calendars in disparate calendaring systems.

**[0006]** A method disclosed herein includes accessing a first calendar, the first calendar supported by a first calendaring system and the first calendar containing one or more calendar items; converting the one or more calendar items into a syndicated format to provide a calendar feed; processing the calendar feed; and providing the calendar feed to a second calendar supported by a second calendaring system.

**[0007]** Providing the calendar feed to the second calendaring system may include passing the calendar feed to a syndication-to-calendar application connector.

**[0008]** Processing the calendar feed may include adding a calendar item. Processing the calendar feed may include filtering the calendar feed. Processing the calendar feed may include aggregating at least two calendars into a single feed. Processing the calendar feed may include converting additional calendar items from the second calendar system into a second

calendar feed and resolving one or more conflicts between the calendar feed and the second calendar feed.

**[0009]** The first calendaring system may include a system selected from the group consisting of Microsoft Outlook, Google calendar, iCalendar, and a Salesforce calendar.

**[0010]** The second calendaring system may include a system selected from the group consisting of Microsoft Outlook, Google calendar, iCalendar, and a Salesforce calendar.

**[0011]** The first calendaring system may be different from the second calendaring system.

**[0012]** A computer program product disclosed herein includes computer executable code that when executing on one or more computer devices performs the steps of: accessing a first calendar, the first calendar supported by a first calendaring system and the first calendar containing one or more calendar items; converting the one or more calendar items into a syndicated format to provide a calendar feed; processing the calendar feed; and providing the calendar feed to a second calendar supported by a second calendaring system.

**[0013]** A password management process handles passwords at a remote service that operates as an intermediary between a user and a web service.

**[0014]** A method disclosed herein includes providing a remote service including an RSS interface to a secure enterprise system; receiving login credentials for accessing the secure enterprise system, the login credentials including a username and a password, the password being encrypted for secure communication between a client and the remote service; validating a subscription to the remote service with the username; decrypting the password; encrypting the password for secure communication between the remote service and the secure enterprise system; logging in to the secure enterprise system; and providing the client with the RSS interface to the secure enterprise system.

**[0015]** The secure enterprise system may include a customer relationship management system. The method may further include providing sales leads to the client through the RSS interface.

**[0016]** The secure enterprise system may include a calendar. The method may further include providing calendar entries to the client through the RSS interface.

**[0017]** The secure enterprise system may include a database. The secure enterprise system may include a web application.

**[0018]** The method may further include locally caching a session identifier at the remote service in order to maintain a session between the remote service and the secure enterprise system.

**[0019]** The remote service may include a fee-based service for converting data from a customer relationship management system to a syndicated data feed. The syndicated data feed may include one or more calendar items.

**[0020]** A computer program product disclosed herein includes computer executable code that when executing on one or more computer devices performs the steps of: providing a remote service including an RSS interface to a secure enterprise system; receiving login credentials for accessing the secure enterprise system, the login credentials including a username and a password, the password being encrypted for secure communication between a client and the remote service; validating a subscription to the remote service with the username; decrypting the password; encrypting the password for secure communication between the remote service and the secure enterprise system; logging in to the secure enterprise system; and providing the client with the RSS interface to the secure enterprise system.

**[0021]** The secure enterprise system includes a customer relationship management system. The computer program product may further include computer executable code that performs the step of providing sales leads to the client through the RSS interface.

**[0022]** The secure enterprise system may include a calendar. The computer program product may further include computer executable code that performs the step of providing calendar entries to the client through the RSS interface.

**[0023]** The secure enterprise system may include a database. The secure enterprise system may include a web application.

**[0024]** A method disclosed herein includes receiving login credentials from a user at a proxy for logging in to a network service, the login credentials including a username and a password; determining whether the login credentials have been previously and unsuccessfully applied to log in to the network service; transmitting the login credentials to the network service when the login credentials have not been previously and unsuccessfully applied to log in to the network service; and transmitting an error message to the user when the login credentials have been previously and unsuccessfully applied to log in to the network service.

**[0025]** Determining whether the login credentials may have been previously and unsuccessfully applied includes maintaining a database of usernames and corresponding hashes of passwords for each unsuccessful login attempt. Maintaining the database may include removing a username and corresponding hash of a password when the username is submitted with a password having a hash that does not match the hash in the database.

**[0026]** Syndicated feeds are generated from enterprise data using an interface that dynamically adapts user requests to an enterprise system and dynamically conforms the responsively provided enterprise data to a format suitable for presentation in a syndicated data feed.

**[0027]** A method disclosed herein includes receiving a request from a client to access an enterprise system; accessing the enterprise system; querying the enterprise system to obtain enterprise data; converting the enterprise data into a syndicated feed; and providing the syndicated feed to the client.

**[0028]** The method may further include identifying an object identifier in the enterprise data; performing a supplemental request to the enterprise system for additional data relating to the object identifier; and inserting the additional data into the syndicated feed. The supplemental request may be based upon a type of object specified by the identifier. The type may include one or more of a name, a customer number, a date, and an address.

**[0029]** The request from the client may include at least one relative term, the method may further include converting the relative term into a specific term for presentation to the enterprise system. The relative term may specify a range of times. The range of times may be one or more of this year, this month, last month, today, and next three days.

**[0030]** Converting the enterprise data into a syndicated feed may include using a syntax for formatting during presentation of the feed. The syntax may specify one or more of a local variable, a date, a dollar amount, and a percentage.

**[0031]** Converting the enterprise data into a syndicated feed may include converting at least one error into a human-readable message. The human-readable message may describe one or more of an inability to reach the enterprise system, a failure of login credentials, and a lapsed subscription to a service. The human-readable message may be communicated through a medium including one or more of an electronic mail, an instant message, and a cellular phone text message.

**[0032]** Converting the enterprise data into a syndicated feed may include mapping the enterprise data to at least one RSS extension. Converting the enterprise data into a syndicated feed may include encoding the enterprise query into one or more custom namespace elements of an RSS output, the custom namespace elements providing a machine-readable form to facilitate transfer of the enterprise data to an application. The application may be one or more of a local application and a second enterprise system.

**[0033]** Converting the enterprise data into a syndicated feed may include one or more of filtering, sorting, categorizing and formatting the enterprise data.

**[0034]** A computer program product disclosed herein includes computer executable code that when executing on one or more computer devices performs the steps of receiving a request from a client to access an enterprise system; accessing the enterprise system; querying the enterprise system to obtain enterprise data; converting the enterprise data into a syndicated feed; and providing the syndicated feed to the client.

**[0035]** The computer program product may further include computer executable code that performs the steps of identifying an object identifier in the enterprise data; performing a supplemental request to the enterprise system for additional data relating to the object identifier; and inserting the additional data into the syndicated feed. The supplemental request may be based upon a type of object specified by the identifier. The type may include one or more of a name, a customer number, a date, and an address.

**[0036]** A background data synchronization process employs human-readable communication channels to contact a user and request resolution of synchronization conflicts as they arise.

**[0037]** A method disclosed herein includes performing a background data synchronization process across at least two data sets; identifying a conflict among the at least two data sets requiring manual resolution; creating a human-readable message requesting resolution of the conflict; communicating the human-readable message over a human-readable communication channel; receiving a response to the human readable-message resolving the conflict over the human-readable communication channel; and applying the response in the background data synchronization process.

**[0038]** The human-readable communication channel may carry cellular phone text messages. The human-readable communication channel may carry instant messages. The



human-readable communication channel may carry electronic mail. The human-readable communication channel may carry telephonic communications.

**[0039]** The background data synchronization process may synchronize data for one or more of a calendar, a contacts database, and a mailing list.

**[0040]** Performing the background data synchronization process may include synchronizing data for two or more web applications. Performing the background data synchronization process may include synchronizing data between a web application and a client-based application.

**[0041]** A method disclosed herein includes performing a background data synchronization process across at least two data sets; identifying a conflict among the at least two data sets requiring manual resolution; creating a human-readable message requesting resolution of the conflict; communicating the human-readable message over a first human-readable communication channel; receiving a response to the human readable-message resolving the conflict over a second human-readable communication channel, the second human-readable communication channel being different from the first human-readable communication channel; and applying the response in the background data synchronization process.

**[0042]** At least one of the first human-readable communication channel and the second human-readable communication channel may carry cellular phone text messages. At least one of the first human-readable communication channel and the second human-readable communication channel may carry instant messages. At least one of the first human-readable communication channel and the second human-readable communication channel may carry electronic mail.

**[0043]** A computer program product disclosed herein includes computer executable code that when executing on one or more computer devices performs the steps of performing a background data synchronization process across at least two data sets; identifying a conflict among the at least two data sets requiring manual resolution; creating a human-readable message requesting resolution of the conflict; communicating the human-readable message over a human-readable communication channel; receiving a response to the human readable-message resolving the conflict over the human-readable communication channel; and applying the response in the background data synchronization process.

**[0044]** The human-readable communication channel may carry cellular phone text messages. The human-readable communication channel may carry instant messages. The human-readable communication channel may carry electronic mail. The human-readable communication channel may carry voice traffic.

**[0045]** The background data synchronization process may synchronize data for one or more of a calendar, a contacts database, and a mailing list.

**[0046]** Performing a background data synchronization process may include synchronizing data for two or more web applications. Performing a background data synchronization process may include synchronizing data between a web application and a client-based application.

**[0047]** A calendar-based advertising system inserts advertisements into a calendar.

**[0048]** A method disclosed herein includes accessing a calendar through a programming interface and inserting an advertisement into the calendar as a calendar item.

**[0049]** The calendar item may have a time associated with an occurrence of an advertised event. The advertised event may be one or more of a sale, a television program, a show, a presentation, and a web presentation.

**[0050]** Accessing the calendar may include obtaining access in exchange for use of a free web-based calendar synchronization service.

**[0051]** The advertisement may include one or more of a video presentation, an audio presentation, a Flash media presentation, and a hyperlink to a web page.

**[0052]** The calendar item may have a time that is the current time.

**[0053]** The method may further include generating a reminder for the advertisement from a calendar system that maintains the calendar.

**[0054]** The calendar may be maintained by a web-based calendar system.

**[0055]** Accessing the calendar may include subscribing the calendar to an RSS feed of advertisements. The advertisement may be selected from a plurality of advertisements according to one or more user preferences.

**[0056]** A computer program product disclosed herein includes computer executable code that when executing on one or more computer devices performs the steps of accessing a calendar through a programming interface and inserting an advertisement into the calendar as a calendar item.

**[0057]** The calendar item may have a time associated with an occurrence of an advertised event. The advertised event may be one or more of a sale, a television program, a show, a presentation, and a web presentation.

**[0058]** Accessing the calendar may include obtaining access in exchange for use of a free web-based calendar synchronization service.

**[0059]** The advertisement may include one or more of a video presentation, an audio presentation, a Flash media presentation, and a hyperlink to a web page.

**[0060]** The calendar item may have a time that is the current time.

**[0061]** The computer program product may further include computer executable code that performs the step of generating a reminder for the advertisement from a calendar system that maintains the calendar.

**[0062]** The calendar may be maintained by a web-based calendar system.

**[0063]** Accessing the calendar may include subscribing the calendar to an RSS feed of advertisements. The advertisement may be selected from a plurality of advertisements according to one or more user preferences.

**[0064]** A syndication platform provides a supplemental feed of most-viewed items in a syndicated data feed.

**[0065]** A method disclosed herein includes publishing a first feed, the first feed including a plurality of items; monitoring remote accesses to the plurality of items; analyzing the remote accesses to identify a most-viewed subset of the plurality of items; and publishing a second feed, the second feed consisting of the most-viewed subset of the plurality of items.

**[0066]** The most-viewed subset may be determined for a time period of one hour. The most-viewed subset may be determined for a time period of one day. The most-viewed subset may be determined for a time period of one week. The most-viewed subset may be determined for a time period of one month. The most-viewed subset may consist of ten items.

**[0067]** Monitoring remote accesses to the plurality of items may include measuring a number of views. Monitoring remote accesses to the plurality of items may include measuring a number of citations.

**[0068]** Analyzing the remote accesses may include filtering the plurality of items to identify one or more of the plurality of items containing multimedia attachments. The second feed may consist of items having at least one multimedia attachment.

[0069] A computer program product disclosed herein includes computer executable code that when executing on one or more computer devices performs the steps of publishing a first feed, the first feed including a plurality of items; monitoring remote accesses to the plurality of items; analyzing the remote accesses to identify a most-viewed subset of the plurality of items; and publishing a second feed, the second feed consisting of the most-viewed subset of the plurality of items.

[0070] The most-viewed subset may be determined for a time period of one hour. The most-viewed subset may be determined for a time period of one day. The most-viewed subset may be determined for a time period of one week. The most-viewed subset may be determined for a time period of one month. The most-viewed subset may consist of ten items.

[0071] Monitoring remote accesses to the plurality of items may include measuring a number of views. Monitoring remote accesses to the plurality of items may include measuring a number of citations.

[0072] Analyzing the remote accesses may include filtering the plurality of items to identify one or more of the plurality of items containing multimedia attachments. The second feed may consist of items having at least one multimedia attachment.

[0073] Surge-based advertising permits deployment of advertisements in response to substantially real-time changes in user activity.

[0074] A system disclosed herein includes a monitoring system that monitors traffic patterns at a plurality of web logs in substantially real time and provides an analysis of the traffic patterns including an identification of a surge in activity; a presentation system for displaying the surge in human-readable form; and an advertisement delivery system for deploying an advertisement in response to the surge.

[0075] The surge may include a rapid increase in activity around one or more of a current event, a book, an individual, and a political issue.

[0076] The traffic patterns may include one or more of traffic, subscriptions, identity, navigation, rate of change in traffic, demographics, and change in demographics.

[0077] The analysis of traffic patterns may include an analysis of semantic patterns that include one or more of identity of referrals, source popularity, content, persistence of a topic, trajectory of a topic, relationship of topic to location of readers, and relationship of topic to web location.

[0078] Deploying the advertisement may include modifying a placement of the advertisement. Modifying the placement of the advertisement may include directing the advertisement at a locus of the surge. Deploying the advertisement may include modifying content of the advertisement.

[0079] A method disclosed herein includes monitoring traffic patterns and semantic patterns on the web; detecting a surge in web activity in substantially real time based upon the traffic patterns and the semantic patterns; and deploying an advertisement in response to the surge.

[0080] The surge may include a rapid increase in activity around one or more of a current event, a book, an individual, and a political issue.

[0081] The traffic patterns may include one or more of traffic, subscriptions, identity, navigation, rate of change in traffic, demographics, and change in demographics.

[0082] The semantic patterns may include one or more of identity of referrals, source popularity, content, persistence of a topic, trajectory of a topic, relationship of topic to location of readers, and relationship of topic to web location.

[0083] Deploying the advertisement may include modifying a placement of the advertisement. Modifying the placement of the advertisement may include directing the advertisement at a locus of the surge. Deploying the advertisement may include modifying content of the advertisement.

[0084] A computer program product disclosed herein includes computer executable code that when executing on one or more computer devices performs the steps of monitoring traffic patterns and semantic patterns on the web; detecting a surge in web activity in substantially real time based upon the traffic patterns and the semantic patterns; and deploying an advertisement in response to the surge.

[0085] The surge may include a rapid increase in activity around one or more of a current event, a book, an individual, and a political issue.

[0086] The traffic patterns may include one or more of traffic, subscriptions, identity, navigation, rate of change in traffic, demographics, and change in demographics.

[0087] The semantic patterns may include one or more of identity of referrals, source popularity, content, persistence of a topic, trajectory of a topic, relationship of topic to location of readers, and relationship of topic to web location.

[0088] Deploying the advertisement may include modifying placement of the advertisement. Modifying placement of the advertisement may include directing the advertisement at a locus of the surge.

[0089] These and other systems, methods, objects, features, and advantages of the present invention will be apparent to those skilled in the art from the following detailed description of the preferred embodiment and the drawings. All documents mentioned herein are hereby incorporated in their entirety by reference.

#### BRIEF DESCRIPTION OF THE FIGURES

[0090] The foregoing and other objects and advantages of the invention will be appreciated more fully from the following further description thereof, with reference to the accompanying drawings, wherein:

- [0091] Fig. 1 shows computing facilities communicating through a network facility.
- [0092] Fig. 2 shows communication among software modules of separate computing facilities.
- [0093] Fig. 3 illustrates aspects of software modules.
- [0094] Fig. 4 depicts an operational kernel interface.
- [0095] Figs. 5-10 depict aspects of various processing interfaces that may be exposed through programming interfaces.
- [0096] Fig. 11 depicts a complex processing operation.
- [0097] Fig. 12 depicts a high level programming interface.
- [0098] Fig. 13 shows an architecture for integrating and manipulating distributed data and services.
- [0099] Fig. 14 shows a spreadsheet paradigm that may be used as a user interface for the navigation, management, and visualization tools described herein.
- [00100] Fig. 15 shows a graphical user interface for managing services.
- [00101] Fig. 16 shows a method for calendar synchronization.
- [00102] Fig. 17 shows a method for managing passwords at a remote service.
- [00103] Fig. 18 shows a usage-based post prioritization process.
- [00104] Fig. 19 shows a surge-based advertising process.

## DETAILED DESCRIPTION

**[00105]** A platform for processing in a wide scale network may be supported through a collection of logical software modules exposed to external users through an interface such as an HTTP get/post interface. The server supporting these services may access other services and provide services to other such services, in an arrangement where services can also act, for example, as a client of a remote (or local) service, each communicating through get, put, post, and delete methods. This allows the logical software modules to be arranged in user-defined or machine-defined configurations, with the output of one module being provided as the input to another, whose output is the input to yet another, and so on. In addition, this allows services to access external services as a client, permitting implementation of any services that can be defined using the core services described herein, either alone, or in combination with exposed services available on the network. Thus, it will be appreciated that while the system described below with reference to Figs. 1-12 is one useful approach to deploying individual services in a manner that accommodates use and combination with other services, the techniques described herein may more generally be applied to the creation and management of composite services, and all such variations that would be clear to one of ordinary skill in the art are intended to fall within the scope of this disclosure.

**[00106]** Referring to Fig. 1, a system 100 may include a plurality of computing facilities 102 that are operatively coupled via a network 104. Within the computing facilities 102 are one or more logical building blocks 108, which themselves may be operatively coupled. In embodiments, this coupling may be between at least two logical building blocks 108 within the same computing facility 102. Alternatively or additionally, this coupling may be between at least two logical building blocks 108 within different computing facilities.

**[00107]** The network 104 may be an IP-based data network providing data communications between at least two computing facilities 102. This network may include the Internet, a WAN, a MAN, a LAN, a personal-area network, or any other IP-based data network, including any IP-based network component, in any arrangement or configuration. The network 104 may also, or instead, employ non-IP communications such as Asynchronous Transmission Mode communications or any other suitable communications protocol(s).

**[00108]** The computing facility 102 may be a microprocessor-based computer. This computer may include a rack-mount server, a workstation, a tower computer, a laptop computer,

a personal computer, a palmtop computer, a portable computer, a cellular phone, a set top box, a router or other network element, a portable communications device such as a Blackberry, an embedded computer, a computer associated with a sensor (such as may be used in a distributed sensor network), and so forth.

**[00109]** The logical building block 108 may be implemented as a software program. This program may be associated with one more processes and/or one or more threads of execution. The building block 108 may be composed of a number of software components, which are described in great detail hereinafter. It will be understood that, while a microprocessor is one common embodiment of a device executing software such as the logical building block 108, the computing facility 102 may also, or instead, include an ASIC, FPGA, PLA, PLD, or any other special-purpose hardware device that is fabricated and/or programmed to execute the building block 108. Throughout this disclosure, it should be appreciated that terms such as “software”, “program”, and “execution”, are to be interpreted broadly as any implementation of a logical process, unless a different meaning is explicitly provided or otherwise clear from the context.

**[00110]** Core services 110, which may be for example any of the services described below, along with related methods and interfaces, may be available through the network 104. The core services 110 may provide any functionality suitable for supporting, combining, and publishing new services formed from the services of the computing facilities 102, which may be ad hoc services, and any services selected from the core services 110. It will be understood that the computing facilities 102 as described herein may generally provide any ad hoc services along with self-defined programming interfaces. As will be discussed in greater detail below, the core services 110 may include services for discovery, indexing, documentation, and assessment of such services in order to facilitate end use by clients 112. The core services 110 may also include any number of services that support creation of new composite services by providing, e.g., security, conditional access, transaction processing, data conversion, and any other functions that might commonly be employed to build sophisticated services from ad hoc functional building blocks available on the Internet. In one aspect, the core services 110 may operate generally as a server, or a single point of contact on a network for various services.

**[00111]** Note that in Fig. 1, metaservices 120 are depicted separately from core services 110. This reflects the general notion that metaservices 120, i.e., services for managing



services, are distinct from other services that may be employed to support a robust metaservices infrastructure, i.e., core services 110. It will be appreciated that this distinction may blur with respect to certain services. For example, where a filter is employed in passing RSS data from one service to another service, this may be viewed as a metaservice or a core service. As such, these terms should be understood to be correspondingly flexible in the following discussion, unless a specific meaning is indicated or otherwise clear from the context. It will also be understood that, while Fig. 1 depicts the core services 110 and metaservices 120 at a single network location, these services may be separate, and/or may be distributed across a network at two or more redundant and/or unique locations.

**[00112]** In one aspect, the core services 110 may be viewed as a coherent integration solution stack including a number of discrete layers. Each layer may provide a well-defined interface to two adjacent layers, as in a conventional protocol stack. In this manner, each functional area may be developed independently by numerous parties each of whom may improve, customize, optimize, or otherwise adapt the layer to specific or generalized usage. Alternatively, each layer may operate as a stand-alone collection of services that may be invoked independently of other layers. Numerous other configurations are possible, and will be clear to one of ordinary skill in the art. All such arrangements are intended to fall within the scope of this disclosure. The relevant features may be decomposed in a variety of manners. One example is set out in U.S. App. No. 11/223,826, the entire contents of which are incorporated herein by reference. As another example, an integration stack may include the following services generally intended to support integration of other remote services into composite services or platforms.

**[00113]** Layer seven: One layer of the stack may contemplate various modes of human communication and interaction, and enable sharing and usage among communities and users in various combinations. This may include, for example, communities, swarms, cross-functional teams, collaborations, dialogues.

**[00114]** Layer six: One layer of the stack may relate specifically to media outputs of various forms, such as interactive media, communication, audio, visual, audio-visual, presentation, and other media, as well as information services and the like.

**[00115]** Layer five: One layer of the stack may address metaservices, such as the discovery, integration, modification, and adaptation of services, along with searching and

publication thereof. This may include integration of web services, superservices, scripts, metatools, superservice libraries, automated testing of end-to-end integrations of services such as those described herein, and any other services and/or content, and the like. As noted above, metaservices 120 may optionally be deployed as separate and discrete from core services 110, in which case this layer of an integration stack may be omitted, or may simply point to or interface with a separate metaservices 120 component.

**[00116]** Layer four: One layer of the stack may address certification of operability and interoperability with reference to one or more standards, such as objective, publicly available standards for operability of the layer three web superservices. Generally, this may address performance matters such as usability, relevance of performance achieved, stability, reliability, scalability, openness and extensibility, software compatibility, hardware compatibility, end-to-end compatibility, and so forth. This may also, or instead, address standards compatibility with relevant standards such as XML, HTML, RSS, OPML, WSDL, and so forth.

**[00117]** Layer three: One layer of the stack may address decomposition and reuse of services such as web superservices. This may include development of utilities to compose, publish, secure, authenticate, gather, archive, search, filter, analyze, display, email, or otherwise manipulate services. Alternatively, some or all of this low-level service/superservice manipulation may be incorporated into the Layer five metaservices described above.

**[00118]** Layer two: One layer of the stack may embrace participation by a worldwide community of users, activist, developers, entrepreneurs, or otherwise contemplate inclusion of various disparate users and sources of services. This may advantageously provide a common, shareable platform for developing superservices and metaservices. It will be noted that this layer is distinguished from layer seven, which relates to the manner in which discrete services or composite services are presented to end users, while layer two relates to the manner in which developers and others participate in creation of new services.

**[00119]** Layer one: One layer of the stack may provide low-level physical connectivity for the variety of simple, stable, ubiquitous standards (URL, SOAP, RSS, OPML, XML, HTML, etc.). This layer ensures that inputs to and outputs from other layers can communicate with external resources and users.

**[00120]** It will be appreciated that integration of services may be accomplished in a number of different ways, and may include different allocations of components in the integration stack. In one embodiment, significant advantages may be realized from a standardized, end-to-end model to interconnect communities of users with low-level physical protocols and services deployed thereon. In general, this conceptual architecture provides a platform for customizing and integrating the functionality of arbitrary combinations of ad hoc services deployed as remote, third-party programming interfaces. Using the platform described herein, application programming interfaces such as those available from Google, Google Maps, MSN Search, eBay, Amazon, Yahoo, and myriad lesser-known providers of network-accessible programming interfaces, can be integrated into a new, composite service which may be used privately or released as a new programming interface or as a self-contained Application Interaction Interface (“AII”) – a web application adapted for direct human use through a browser or other client.

**[00121]** A database 111 may support the core services 111 both by storing procedures and code for the core services 110, and by providing a data repository or database for users of the core services 110. In addition, the core services 111 may provide a data store for external services, such as ad hoc services running on the computing facilities 102. As will be appreciated from the description below, this may advantageously expand the functionality of ad hoc services by providing a buffer for inputs to or outputs from these services when sequencing of a chain of operations from different ad hoc service locations. More generally, those of ordinary skill in the art will appreciate many advantageous uses of persistent memory. Further, the core services 110 may provide differential levels of database services. For certain users, such as authenticated users, the database 111 may be available for general usage in connection with core services 110 or otherwise. For other users, the database 111 may not be available. In this latter environment, the core services 110 may provide a service that permits a user to incorporate the user’s local storage, such as storage on the client 112, as a database or short term memory store. While the database 111 is depicted as a conventional database 111 behind the core services 110 and/or metaservices 120, it will be understood that other techniques may be employed to provide an actual or effective database in connection with composite services and metaservices.

**[00122]** In one aspect, RSS or a similar syndication technology may be employed for data persistence between stages of a composite service. Thus, a metaservice 120 or other program coordinating execution of a composite service may direct a first service to output an RSS feed. The URL of the RSS output may then be used as an input to a second service, and so forth. As a significant advantage, this approach provides a simple, convenient, ubiquitous, and readily accessible resource as a buffer for composite service processing. Using techniques described, for example, in U.S. App. No. 11/223,826, entitled "Enhanced Syndication" these RSS streams may, in turn, be secured to provide for conditional access based on user identity (which may be encoded by the metaservice or composite service that is using the RSS buffer). Access to these process-oriented RSS feeds may be permissions based, or otherwise restricted. In one aspect, intermediate or final RSS feeds may be useful in multiple ways, and it may be desirable to have intermediate data streams available for general, public use. In another aspect, intermediate or final RSS feeds may be highly proprietary, and it may be desirable to have some level of security associated with content therein.

**[00123]** As an additional component, it may be useful to monitor pools of data associated with processing of composite services. That is, a large amount of data may be generated and distributed among numerous RSS sources on a network. An audit tool may be provided for reviewing and analyzing levels of security on such sources. This may include an analysis of the content and vulnerability of such sources, with respect to either each source as a network resource or the underlying data, or both. While this tool may provide for useful security audits of an RSS-based data store for composite service processing, it will be understood that a tool for security audits of RSS data may have significant value independent of the composite services discussed herein. That is, an enterprise, publisher, or other entity may periodically audit RSS sources for vulnerabilities, with respect to, for example, whether data is secured in the manner intended. Where security flaws are identified, the audit tool may employ remedial measures such as securing the source of RSS data, e.g., by requiring suitable encryption on RSS output, or by securing or quarantining the offending RSS feeds. The audit tool may also, or instead, evaluate security risks based upon the data sources available to an RSS feed. In such cases, suitable responses may include filtering output from the feed to remove any secured source data, along with hardening the source itself against filter circumvention. Thus, there is disclosed herein a general tool for evaluating security exposures

associated with syndicated data, and more generally, any pools of unstructured or structured data. The tool may provide a security profile characterizing data exposure. The tool may also or instead, actively secure sources according to a security policy, or make recommendations concerning exposure and risk mitigation.

**[00124]** The client 112 may be any device communicating with the network 104. In general, the client 112 may access various combinations of the core services 110 and the ad hoc services from the computing facilities 102 to provide a composite service, described in greater detail below. The composite service may in turn be published as a new ad hoc service through a user-defined programming interface, either through the core services 110 and related infrastructure, or on a user-selected server available through the network 104.

**[00125]** Thus, in general there is disclosed herein a technique for supporting use and combination of ad hoc remote services through one or more core services and/or that are available on a network.

**[00126]** Referring now to Fig. 2, a system 200 including two computing facilities 102 may support communication among services. Each of the computing facilities 102 in this depiction contains one logical building block 108. Within each of the logical building blocks 108 are a number of software components. These may include, for example, a link protocol layer 202; a network protocol layer 204; a transport protocol layer 208; an application protocol layer 201; an operational kernel 212 with a kernel interface 214; a plurality of aspects 216, each of which has an aspect interface 218; and a plurality of service software modules 220 that are operatively coupled to the operational kernel 212 and the aspects 216 via the interfaces 218. The link protocol layer 202 may be implemented to provide a logical coupling to the network 104, such as via Ethernet, WiFi, ATM, and so forth; the network protocol 204 layer may implement IP (v4 and/or v6); the transport protocol layer 208 may implement TCP; and the application protocol layer 210 may implement HTTP. The operational kernel 212 implements methods of the kernel interface 214, which are described in detail hereinafter. The aspects 216 of the operational kernel implement methods of the aspect interfaces 218, which are also described in detail hereinafter. The server software modules may implement arbitrary services 220 by utilizing any native functionality of the computing facility 102 combined with the implementations provided by the operational kernel 212 and its aspects 216. The native functionality of the computing facility 102 may be any of the functions or features of a

computer and a resident operating system on the computer. It will be appreciated that the resident operating system may be any operating system, including a proprietary operating system (such as Windows XP), an open-source operating system (such as OpenBSD), a real-time operating system, an embedded operating system, and so forth. Likewise, it will be appreciated that the functions and features of the computing facility 102 may vary from implementation to implementation and that the present invention is not limited to any particular type of computing facility 102 or operating system. Moreover, it will be appreciated that there is not an intrinsic need for an operating system, and in some embodiments the logical building block 108 may run directly on hardware of the computing facility 102.

**[00127]** As depicted, the boundaries between the software modules are logical boundaries. According to software engineering practices, these software modules may be implemented as individual software modules. However, the software modules may also be implemented in a more monolithic fashion, with logical boundaries omitted or loosely defined in the implementing source code. For example and without limitation, a network protocol stack of several layers may be implemented in a single, monolithic tract of source code. It should be appreciated that various levels of integration or modularity may be reflected in a particular implementation of the software modules. All such implementations are within the scope of the present invention. In an embodiment, the operational kernel 212, its aspects 216, and the service software modules 220 are written in a modular fashion, with the aspects 216 coupled to the operational kernel 212 via a well-defined interface (depicted simply as a boundary between the aspect 216 and the operational kernel 212) and with the service software modules 220 accessing the services provided by the operational kernel 212 and its methods solely via the methods of their interfaces 214.

**[00128]** The methods of the interfaces 214, 218 that are implemented by the operational kernel 212 and its aspects 216 provide an abstraction of the underlying software modules and computing facility 102. Some of these services may be implemented and provided by the operational kernel 212 itself, some may be implemented and provided by the aspects 216 of the operational kernel, and others may be implemented and provided by the service software modules 220. As a general guideline, certain core services 110 may be provided by the operational kernel where those services that are commonly used or required, while services that are application-specific may be implemented by the service software modules 220. It will be

appreciated that which services should be implemented in which modules may vary, or may change over time.

**[00129]** Generally, as referred to here, a service provides a useful, concrete, and tangible result, such as by executing a logical process of a logical building block 108. This logical process can include an implementation of an interface 214 and/or 218, an implementation of a service software module 220, an implementation of an operational kernel 212, an implementation of software provided to the logical building block 108, an implementation of a software module of the logical building block 108, or the implementation of any other software associated with the logical building block 108. Certain services, such as superservices, web services, composite services, and metaservices are discussed in greater detail below. In general, services provided through non-standard application programming interfaces from remote network resources – interfaces such as Google Maps – are referred to herein as ad hoc or unstructured services, and are also intended to fall within the scope of services as that term is used herein.

**[00130]** The system described herein may employ message-passing to communicate an object representation 222 among logical building blocks 108. When building blocks 108 exist in different computing facilities 102, the network 104 provides the communication of the object representation 222 between the logical building blocks 108. In this case, the object representation 222 is transmitted and received by the link protocol layers 202 of the logical building blocks 108. The communication of the object representation 222 may be performed in a one-to-one fashion, with a single building block 108 communicating the object representation 222 to another single building block 108. In alternate embodiments, the communication may be performed in a one-to-many or many-to-many fashion. In these alternate embodiments, the communication may utilize a multicast or broadcast technique or protocol.

**[00131]** The object that is represented by the object representation 222 can be any data including an element of a file-stream; a file-stream; a pool of file-streams; a relation or configuration; a tag; a service; an external reference such as a URI or URL, a description, specification, or outline of any of the foregoing; or any other data. For convenience, all of these things that are herein referred to as the object. The data may include an RSS feed; an OPML file or object; an XML file; an HTML file; an HTTP stream; an RTP stream; an image file; an audio file; an audiovisual file; a text file; a Word document; a JavaScript file; and so forth.

**[00132]** Thus there is described herein a generalized technique for sharing instructions and data among ad hoc services in a networked computing environment. As noted above, where a composite service employs a number of services in sequence (or in parallel), an RSS-based buffer or other database 111 may be employed to cache interim and/or final results.

**[00133]** Fig. 3 shows aspects of an operational kernel. The aspects may be arranged around a canonical organization of core or atomic functions desirable for facilitating generalized use of loosely structured or ad hoc network services – the core services 110 described in reference to Fig. 1. The aspects of the operational kernel may without limitation include an application aspect 602, a data aspect 604, a syndication aspect 608, and other or hybrid aspect 610 (referred to hereinafter as the other aspect), a semantic aspect 612, and an infrastructure aspect 614. Other aspects may be useful in various processing contexts, and may be included in the operational kernel or as services associated therewith, such as transactions (i.e., events involving exchange of funds), security, encryption, and authentication. Numerous arrangements and hierarchies are possible for these core services. Figure 3 suggests one canonical arrangement of services by way of example and not by way of limitation. Figure 3 provides high-level descriptors for data processing, semantic processing, syndication process, infrastructure processing, and so forth. Each of these aspects is discussed in greater detail below.

**[00134]** Fig. 4 shows a high level depiction of an HTTP-based programming interface for accessing services. Referring to Fig. 4, the operational kernel interface 212 for accessing aspects of the operational kernel such as those described above may include a Get method 702 and a Post method 704. The Get method 702 provides a way of getting an object from a URL or providing the object at the URL. The Post method 704 provides a way of posting an object to a service at a URL or accepting an object posted to a service at a URL. The operational kernel interface may also include a Put method and a Delete method. The Put provides a way of putting an object to a URL or accepting an object that is put to a URL. The Delete method provides a way of deleting an object at a URL, including requesting that the object be deleted or accepting the request and deleting the object in response to the request. These methods may collectively provide a bi-directional HTTP-based programming interface where the Get 702, Put, Post 704, and Delete methods are implemented according to HTTP, while the operational kernel 212 operates as an HTTP client and an HTTP server. In other embodiments, the Get 702,



Put, Post 704, and Delete methods may be implemented via SMTP or any other protocol that allows bi-directional communication.

**[00135]** The system may present a variety of services or functions to external users through a programming interface accessed using the methods of the operation kernel. A number of such functions and services that might be usefully provided in a processing environment are described below. In general, these services may provide a functional platform for integrating disparate services. This can accommodate ad hoc combinations of unstructured services, each of which may be available as a programming interface on a network, by providing a set of core services to augment functionality. Thus, for example, ad hoc combinations of services can further incorporate security measures such as conditional access or authentication with reference to a trusted third party, or incorporate semantic processing, search, data processing, and so forth.

**[00136]** Referring now to Fig. 5, the application-aspect interface 802 may provide a variety of methods related to interaction with client-side applications. This may, for example, include a program method 804, a media viewer method 808, a social network method 810, a user interface method 812, and an e-commerce method 814. The program method 804 may provide accessing to an application executed or interpreted by a logical building block 108. The media viewer method 808 may provide a service that processes a media object, such as by preparing for display or converting between media formats. The social network method 810 may provide a service that relates to social networking such as attention brokering, popularity measurement, rankings, and so forth, and may support social networking among participants by supporting groups, buddy lists, affiliations, preferences, and so forth. The user interface method 812 may control presentation of content to an end user according to, e.g., display hardware, screen layout preferences, substantive preferences, and so forth. The e-commerce 814 method may provide services specific to a market, industry, or user group including without limitation formatting of data or performing transactions according to industry customs or standards. This may also include forms of electronic payment, credit evaluations, and e-commerce techniques such as shopping carts, catalogs, and so forth.

**[00137]** Referring now to Fig. 6, the client-aspect interface 902 may include a variety of methods specific to client interactions with a service. For example, this may include a format-display method 904, a transaction method 908, an identity method 910, and a conditional

access method 812. The format-display method 904 provides a way of formatting or specifying a format of an object for rendering on a target display.

**[00138]** The transaction method 908 may provide a way of conducting a transaction. It will be appreciated that a wide array of transactions and payments may be usefully employed with the systems described herein. Transactions may include, for example, receiving and/or executing financial transactions using a variety of payment infrastructures including ACH, credit card, wire transfer, PayPal and similar commercial payment services, and so forth. As another example, and not by way of limitation, transactions may include financial transactions related to use of the core services 110, metaservices 120, and other, third party services as described generally herein. For example, the core services 110 may support pay-per-use or subscription models for internal services and remote services. Where remote services are employed, the system may track usage and provide periodic reporting. The system may further support automated or manual payment for such services through the core services 110 transaction method 908.

**[00139]** More generally, the transaction method 908 may support tracking of usage charges for complex composite services. That is, a user may create and publish a composite service through the system that employs other ad hoc services, one or more of which require payment (e.g., a subscription, a database access charge, a time charge, a processor time charge, or the like). At the same time, the composite service publisher may specify fees for the composite service, which may be fixed or variable, and may depend on third party usage costs. The transaction method 908 may bill charges to, or collect charges from, a user of the composite service, and may further manage payment among the publisher and any of the ad hoc services. When coupled with security features provided by other core services 110 described herein, this may support, for example, an enterprise computer platform that outsources certain services such as payroll processing or access to digital libraries on a pay-as-you-go or per-user basis. More generally, this platform supports integration of disparate, commercial services for individual or enterprise use, which may also be seamlessly combined with any related non-commercial ad hoc services.

**[00140]** In another aspect, the transaction method 908 may cooperate with e.g., methods of the infrastructure-aspect interface 1302 or the data-aspect interface 1002 to manage payment for enhanced service. Thus, for example, a publisher or user of a composite service

that includes commercial, third-party, ad hoc services may pay for guarantees or service levels related to QoS, bandwidth, processing throughput, and the like. Similarly, a user (or publisher) of a composite service may coordinate cost-effective usage of services, such as by scheduling use of certain commercial services at lower-cost, off-peak times. In one embodiment, the composite service may simply be a scheduler for scheduling work to a commercial service provider in a cost-effective manner. In various embodiments, a composite service may provide a single login access point for combined authentication, service usage, and payment.

**[00141]** The identity method 910 may provide a way of accessing, establishing, verifying, evaluating or otherwise processing an identity or identity attribute. The conditional access method 921 may provide a way of specifying or enforcing a conditional access rule, or otherwise controlling access to data on a conditional basis. As illustrated by some of these examples, one or more aspects may reside in multiple interfaces, or reasonably be incorporated into different interfaces. For example, the identity and conditional access methods may be associated with a security interface or infrastructure interface. All such variations are intended to fall within the scope of this disclosure.

**[00142]** Referring now to Fig. 7, the data-aspect or search-aspect interface 1002 may include methods generally related to data manipulation. This may include a database method 1004, a data quality method 1008, a data transformation method 1010, a search/filter/cluster method 1012, a search engine method 1014, a spider method 1016, and an information relationship/hierarchy/categorization method 1018. The database method 1004 may provide a way of accessing or providing a database or database management system. The data quality method 1008 may provide a way of accessing or providing a quality metric associated with the object. The data transformation method 1010 may provide a way of accessing or providing a service that transforms an object from one representation to another. The search/filter/cluster method 1012 may provide a way of accessing or providing a service that searches, filters, and/or clusters objects. The search engine method 1014 may provide a way of accessing or providing a search engine. The spider method 1016 may provide a way of accessing or providing a spider. The information relationship/hierarchy/categorization method 1018 may provide a way of accessing or providing service that relates to an information relationship, hierarchy, or categorization. Uses of these and other atomic services are described in greater detail below.

**[00143]** Referring now to Fig. 8, the semantic-aspect interface 1102 may include methods that generally support semantic processing. This may include, for example, a metadata creation method 1104, an interpretation method 1108, a metadata enrichment method 1110, a knowledge structure creation method 1112, and a dictionary or thesaurus method 1114. The metadata creation method 1110 may provide a way of accessing or providing a service that creates metadata. The interpretation 1108 method may provide a way of accessing or providing a service that interprets or translates an object using natural language processing. The metadata enrichment method 1110 may provide a way of accessing or providing a service that enriches metadata. The knowledge structure creation method 1112 may provide a way of accessing or providing a service that creates a knowledge structure. The dictionary or thesaurus method 1114 may provide a way of accessing or providing a dictionary or thesaurus. In general, these methods may support semantic processing and manipulation of data. Other supporting functions may be similarly provided, such as an identification method that can provide a globally unique identifier upon request, and/or add such a globally unique identifier to metadata for a media object or other content. In one embodiment, the core services 110 may include this service. In another embodiment, the identifier service may be provided by a trusted third party that can certify identification numbers and usage thereof.

**[00144]** Referring now to Fig. 9, the syndication-aspect interface 1202 may support syndication functions. This may include a publish method 1204, a subscribe method 1208, an aggregate method 1210, a republish method 1212, and a manage syndication information method 1214. The publish method 1204 may provide a way of accessing or providing a service that publishes an object. The subscribe method 1208 may provide a way of accessing or providing a service that subscribes to an object. The aggregate method 1210 may provide a way of accessing or providing a service that aggregates objects. The republish method 1212 may provide a way of accessing or providing a service that republishes objects. The manage syndication information method 1214 may provide a way of accessing or providing a service that manages syndication information. In general, these methods may support syndication functions. In general, these methods may support fundamental syndication functions that enable subscription-based distribution and re-distribution of data.

**[00145]** Referring now to Fig. 10, the infrastructure-aspect interface 1302 may provide methods related to network or service management infrastructure. In general, these

methods may support an infrastructure for using ad hoc services by providing fundamental network and data functions. This may include a security method 1304, a logging method 1308, an authentication method 1310, a communications method 1312, a traffic management method 1314, and a pinging method 1316. The security method 1304 may provide a way of accessing or providing a service that may provide network or data security. The logging method 1308 may provide a way of accessing or providing a logging service. The authentication method 1310 may provide a way of accessing or providing an authentication service.

**[00146]** The communications method 1312 may provide a way of accessing or providing a communications service. This may include, for example, access to low-level functions such as network and physical layer protocols. This may also, or instead, include various protocols for conventional communications types such as e-mail (e.g., SMTP, POP, Microsoft Exchange Server), collaborative platforms (e.g., Lotus Notes), VoIP, instant messaging, video conferencing, text messaging, telecommunications, and so forth. In an alternative embodiment, the communications method 1312 may support network communications protocols while, for example, the social network method 810 of the application-aspect interface supports higher-level communications protocols.

**[00147]** The traffic management method 1314 may provide a way of accessing or providing a traffic management service. In one aspect, this method may provide reporting on current or historical traffic and usage, which may be provided by corresponding services, or may be independently monitored and reported within the core services 110, or both. These metrics may be reported on a per user basis, on a per service basis, or in any other combination useful to a recipient. It will be understood that, as with many of the other methods described herein, the method may be adapted to receive highly parameterized requests for data, such as traffic request for a specific service as used by a specific group of users over a specific time period, or the method may provide very simple, low-level functions, with other core services 110 or metaservices 120 providing functionality to extract desired reports from raw data extracted by the method. When used in combination with other core services 110 or other services, this method may be configured to generate and forward periodic reports. In another aspect, this method may provide tools for proactively managing usage of services. This may include, for example, scheduling and prioritization of usage, and reports on status of currently executing composite services.

**[00148]** The ping method 1316 may provide a way of accessing or providing a ping service. The ping method 1316 may provide conventional ping services such as testing the availability and network characteristics (e.g., round trip time) or network hosts or other remote services. The ping method 1316 may also, or instead, provide syndication-related ping services such as notifying a ping server of content updates (such as for RSS authoring) or checking for the availability of new content (such as for RSS aggregators or readers).

**[00149]** A validation method 1318 may support evaluation and validation of remote services. This may generate user-specified or automated test calls to remote services to ensure proper functioning, such as by reference to a published specification of a corresponding programming interface. More generally, this method may support a host of metrics for remote, ad hoc services including reliability, mean time between failure, performance, bandwidth, latency, quality of service, availability, and the like. Related services may include audits for security, reliability, and so forth. This method may also be used in combination with the traffic management method 1310 described above to more efficiently schedule processes, or to optimize system usage based upon variations in current and anticipated usage of various services underlying a composite service.

**[00150]** Fig. 11 shows a composite service 1100 formed from other services. Referring to Fig. 11, a system 100 may include a plurality of sources 1404, a spider 1402, an aggregator 1602, a group filter 1802, a representation converter 1902, a membership control 2002, and a plurality of clients 2008. Functionally, this arrangement may search (through the spider 1402) and aggregate content from the sources 1404, filter the results, and convert the filtered results into a suitable representation, which may include media conversion and/or arrangement into an HTML or other format. The results may be restricted to specific clients 2008 using access control 2002. It will be understood that some or all of the services described above may be core services 110. At the same time, some or all of the services described above may be ad hoc services accessed through programming interfaces at third-party network locations. However at least one core service 110 or metaservice 120 will be employed to reconstitute the underlying services as a composite service 1100.

**[00151]** The clients 2008 may be computer programs under the control of a human, such as a feed reader in a browser that is being interactively operated by the human. The clients

2008 may be automatic computer programs, such as the service software modules 220 or any other software modules of the logical building block 108. The lines between the elements depict operative couplings between services. The arrowheads generally depict the flow of data and instructions, and imply a corresponding client-server coupling. Although this suggests a pull-based methodology (i.e. clients request then servers respond), it will be appreciated that other embodiments exist. For example, the elements may be configured as a collection of peers in a peer-to-peer configuration and/or may employ a push-based methodology (i.e., where servers transmit to clients without receiving explicit requests). All of these arrangements, and other configurations of the logical elements described herein, may fall within the scope of the present disclosure. More general, Figure 11 serves as an example only, and in no way limits the scope of this disclosure. Further examples are provided below in narrative form.

**[00152]** As a more concrete and detailed example of how the core services 110 may be adapted to special purpose use, the elements described above may be deployed to provide an OPML server and database, with the core services 110 server, or another remote server, acting as a centralized access point. The OPML server may be configured for user manipulation of OPML content. The OPML server may provide services and content to clients 112 using, for example, a Web interface, an API, an XML processing interface, an RSS feed, an OPML renderer, and the like.

**[00153]** The OPML server may, for example, provide a search engine service to visitors. Output from the OPML server may be an OPML file. The file may, for example, be provided a name that explicitly contains the search query from which it was created, to facilitate redistribution, modification, recreation, synchronization, updating, and storage of the OPML file. A user may also manipulate the file, such as by adding or removing outline elements representing individual search results, or by reprioritizing or otherwise reorganizing the results, and the user may optionally store the revised search as a new OPML file. Thus in one aspect the OPML server creates new, original OPML content based upon user queries submitted thereto. In a sense, this function is analogous to the function of aggregators in an RSS syndication system, where new content may be dynamically created from a variety of different sources and republished in a structured form.

**[00154]** The OPML server may, more generally provide a front-end for an OPML database, which may operate from the database 111 of the core services 110 as described above,

that stores OPML content. The OPML database may store OPML data in a number of forms, such as by casting the OPML structure into a corresponding relational database where each OPML file is encapsulated as one or more records. The OPML database may also store links to external OPML content, or may traverse OPML content through any number of layers and store data, files, and the like externally referenced in OPML documents. Thus for example, where an OPML file references an external OPML file, the external OPML file may be retrieved by the database 111 and parsed and stored. The external OPML file may, in turn, reference other external OPML files that may be similarly processed to construct, within the database 111, an entire OPML tree. The OPML database 111 may also, or instead, store OPML files as simple text, or in any number of formats optimized for searching (such as a number of well-known techniques used by large scale search engines Google, AltaVista, and the like), or for OPML processing, or for any other purpose(s). In a sense, the OPML database may provide the coherency for formation of an OPML network among an array of clients 112 and computing facilities 102, where content within the network is structured according to user-created OPML outlines.

**[00155]** The OPML database may, for example, operate through the OPML server to generate, monitor, and/or control spiders (deployed using, e.g., core services or ad hoc services) that locate OPML content. A spider may, upon identification of a valid OPML file, retrieve the file and process it into the database 111. A spider may also process an OPML file to identify external references, systematically traversing an entire OPML tree. A spider may be coordinated using known techniques to identify redundant references within a hierarchy. A spider may also differentiate processing according to, e.g., structure, content, location, file types, metadata, and the like. The user interface described below may also include one or more tools for configuring spiders, including a front end for generating initial queries, displaying results, and tagging results with any suitable metadata.

**[00156]** By way of example, and not of limitation, medical records may be stored as OPML files, either within the database 111, or in a distributed fashion among numerous locations across a network. Thus, for example, assorted X-ray data may be maintained in one location, MRI data in another location, patient biographical data in another location, and clinical notes in another location. This data may be entirely decoupled from individual patients (thus offering a degree of security and privacy), and may optionally include references to other



content, such as directories of other types of data, directories of readers or interpretive metadata for understanding or viewing records, and the like. Separately, OPML files may be created to provide structure to the distributed data. For example, a CT scan OPML master record may index the locations of all CT scan records, which may be useful, for example, for studies or research relating to aggregated CT scan data. This type of horizontal structure may be captured in one or more OPML records which may, themselves be hierarchical. Thus, for example, one OPML file may identify participating hospitals by external reference to OPML records for those hospitals. Each hospital may provide a top-level OPML file that identifies OPML records that are available, which may in turn identify all CT scan records maintained at that hospital. The CT scan master record may traverse the individual hospital OPML records to provide a flattened list of CT scan records available in the system. As another example, an OPML file may identify medical data for a particular patient. This OPML file may traverse records of any number of different hospitals or other medical institutions, or may directly identify particular records where, for example, concerns about confidentiality cause institutions to strip any personally identifying data from records. For certain applications, it may be desirable to have a central registry of data so that records such as patient data are not inadvertently lost due to, for example, data migration within a particular hospital.

**[00157]** Thus in one embodiment there is generally disclosed herein a pull-based data management system in which atomic units of data are passively maintained at any number of network-accessible locations, while structure is imposed on the data through atomic units of relationship that may be arbitrarily defined through OPML or other grammars. The source data may be selectively pulled and organized according to user-defined OPML definitions. The OPML server and OPML database may enable such a system by providing a repository for organization and search of source data in the network 100. Operations (such as traversing OPML trees to fully scope an outline composed of a number of nested OPML outlines) may be performed by a client 112, or may be performed by the OPML server, either upon request from a client 112 for a particular outline, or continually in a manner that insures integrity of external reference links.

**[00158]** In another aspect, there is disclosed herein a link maintenance system for use in an OPML network. In general, a link maintenance system may function to insure integrity of external references contained within OPML files. Broken links, which may result for example

from deletion or migration of source content, may be identified and addressed in a number of ways. For example, a search can be performed using the OPML server and OPML database for all OPML files including a reference to the missing target. Additionally, the OPML server and/or OPML database may include a registry of content sources including an e-mail contact for managers or administrators of outside sources. Notification of the broken link may be sent to all owners of content including a reference to the content. Optionally, the OPML server may automatically modify content to delete or replace the reference, assuming the OPML server has authorization to access such content. The OPML server may contact the owner of the missing content. The message to the owner may include a request to provide an alternative link, which may be forwarded to owners of all content that references the missing content. If the referenced subject matter has been fully indexed by the OPML server and/or OPML database, the content may be reconstructed, and a replacement link to the location of the reconstructed content provided. Various combinations of reconstruction and notification, such as those above, may be applied to maintain the integrity of links in OPML source files indexed in the database 111. In various embodiments the links may be continuously verified and updated, or the links may be updated only when an OPML document with a broken link is requested by a client 112 and processed or traversed by the client 112 or the OPML server in response.

**[00159]** The functionality of this OPML network, or more specifically, the medical OPML network, may be implemented using a combination of core services and metaservices. Thus as disclosed herein, the core services may be configured as a special purpose server, such as an OPML server and database, using pre-defined core services 110 and ad hoc services available as programming interfaces on a network.

**[00160]** Referring now to Fig. 12, an interface 2202 between logical building blocks 108 (alternatively referred to herein as functional elements, or simply elements) may include a Web URL 2204, a description URL 2208, a feed URL 2210, and a kernel URL 2212. This general format may be used for core services 110 and/or ad hoc services. As a significant advantage, ad hoc services conforming to this architecture by providing a description URL 2208 may be more readily discovered and exploited within the metaservices framework described herein.

**[00161]** The Web URL 2204 may provide an interface to a functional element using an HTTP server, which employs HTML-based representations of the services provided by the

element. This optional interface may be employed to provide access to services of the element for a web-only client such as a traditional Web browser.

**[00162]** The description URL 2208 may refer to a location where a client puts or gets a description or configuration file for the interface to the element, including aspects such as formats or syntax for accessing functionality of the element, alternative locations for accessing the element, parameters that may be passed to the element, and interpretation of any result from the service, such as format, structure, return codes, and so forth. The configuration file may be represented as an OPML file, or using any other suitable format.

**[00163]** The feed URL 2210 may provide a location where a client can retrieve a feed-based representation of the objects provided by the functional element. Effectively, this provides an output or response from the service that is accessible using an HTTP Get to, for example, an RSS feed of results. In an embodiment, the feed-based representation is provided according to the RSS 2.0 format, but any suitable format, such as a variety of syndication or outlining formats, XML, plain text, or the like may be used.

**[00164]** The kernel URL 2212 provides a location where a client may access the services of the kernel, its aspects, and the service software components built thereon. In the preferred embodiment, the services are accessed via HTTP Get and HTTP Post, however any suitable protocol may be used. Through this URL, a client may access the services of a functional element, or other functional elements associated with that functional element. More generally, the kernel URL provides a general and adaptable interface through which a client can access any service that the logical block implements, or that that the logical block has access to. Conversely, the other URLs of the interface may provide static pathways to corresponding content.

**[00165]** While HTTP is one useful protocol for use with the systems described herein, other embodiments may be usefully employed. For example, a client may access the services at the kernel URL via the SMTP protocol. In this case, the services at the kernel URL may accept inputs and provide outputs in the form of SMTP e-mail messages. In embodiments, the logical building block may include a plurality of kernel URLs, each of which implements a different protocol. Thus, the logical building block may have an HTTP kernel URL and an SMTP kernel URL. Numerous other examples will be appreciated and are intended to fall within the scope of the present invention.

**[00166]** In embodiments, one element may provide a service that crawls or spiders an environment to generate a description file for the environment, or resources (e.g., other elements) available in the environment. In one aspect, the results may be stored in a database, and the element may present this as a searchable database of functional elements within the environment, such as by indexing the results according to elements of the description file. In another aspect, the element may configure itself to communicate with other elements according to their description files, and the element may further modify its own description file to reflect any new services or remote elements accessible therefrom. It will be appreciated that such an automatically configuring element or group of elements may take many forms. For example, the element may incorporate any identified methods so that they operate within the element. As another example, the element may present references to external or remote methods so that they may be located, but not directly accessed, through the description file of the element. Some embodiments may run in one computing facility, others may operate over a plurality of computing facilities. Some embodiments may automatically provide redundancy, failover, logging, and the like, either by default or optionally through an interface described within the description file.

**[00167]** In one embodiment, the interface 2202 of Fig. 12 may be used to provide the composite service 1100 described in reference to Fig. 11. In another embodiment, the interface 2202 may be used to provide the OPML server and database described above. In general, any composite service that can be created using the core services 110, metase services 120, and ad hoc services described above may be provided as a network-accessible service using the interface 2202 of Fig. 12.

**[00168]** In one aspect, the architecture described above may be employed to provide an interface, such as an HTTP-based, put/get interface for a variety of syndication, outlining, and related functions. In embodiments, aspects of such a system may be presented to an application developer in the form of an Application Programming Interface ("API"). This API may include software interfaces allowing an application developer to access one or more syndication services within an operational kernel or description file of a server. This may include, without limitation, syndication services such as create, publish, and/or subscribe; semantic services such as outlining, listing, adding, deleting, tagging, labeling, analyzing, filtering, sorting, and the like; database functions such as read, write, search, retrieve, and the

like; security services such as encryption, decryption, authentication, access, and the like; infrastructure services such as traffic management, routing, redundancy, logging, and so forth; and any other services that might be usefully employed within an enhanced syndication context as described herein and in the documents incorporated herein by reference.

**[00169]** An application developer may use the API to develop an application that uses one or more of the syndication services and any other services in the enhanced syndication environment, as well as any number of ad hoc services available on a network. The syndication services may be implemented in an operating system, in a database management system, in a user-level process on a client, in a user-level process on a server, as a Web service, and so forth. While in one aspect, the API presented by a server may operate exclusively using on protocol (or combination of protocols), it will be understood that the API may access other services that communicate using a variety of different protocols or communications media, including ad hoc services available through programming interfaces on remote sites. For example, one service may have an API implemented in a user-level process on a client, and the interface between the application and the user-level process may be a socket through which one or more messages may be passed. As another example, one API may be implemented as a Web service, where the interface between a user (which may be another service) and the Web service is an HTTP session over which one or more messages may be passed via SOAP. The application programming interface may employ a TCP/IP socket over which remote procedure calls are passed. The API may be implemented in a database management system. The interface between the application and the user-level process may include XQUERY messages. Alternatively, the database management system may include an integral implementation of the API, which may without limitation be accessed as a Web service. Thus, a simple interface employing HTTP-based gets and puts may expose a variety of services within a networked environment in a manner transparent to a user. Further, this interface may be extended to provide access to services using other programming interfaces.

**[00170]** It will be appreciated that this approach to deploying and integrating services and functions offers significant advantages. The use of HTTP-based gets and puts offers effectively universal accessibility, while URL's offer a commonly accepted platform for addressing elements of an (extensible) API. Similarly, the description file may employ OPML or similar outlining structures for a standardized grammar for describing the interface. The

result may be a highly distributed, multi-user environment of variably-structured services and functional blocks. The system may employ any degree of data typing, and accommodate an ever expanding collection of cooperating elements which may be recursive, self-referential, and recombinant. The collective system may perform a wide variety of syndication-related, as well as non-syndication-related, functions at varying degrees of complexity. Thus, for example, an interface of an element may combine, index, access, move, convert, filter, or otherwise manipulate content. In addition, the interface may be employed to trigger other operations from other building blocks, or to display or transmit data.

**[00171]** In one application, the platform above may support a semantic computer that offers a family of functions organized around processing of content available on a network. This semantic computer may provide any number of core functions for processing, and optionally may provide extensibility as described above for additional functions that are, for example, user-created and endorsed by a user community. For example, the semantic computer may include a programming interface that includes an interface for membership/sign-in, spider configuration and deployment, aggregation or storage of spider results, parsing, organizing (using, e.g., OPML), and output or display of results.

**[00172]** By way of example and not of limitation, a programming interface for performing these functions may include the following core elements:

Get

In: A URL of an element, object, pool, relation; zero or more constraints

Out: The element, object, pool, relation, according to the constraints

Constraints: "current," "in date range <start> <end>," "matches <keyword>," "index is <index>."

Post

In: An element, object, pool, relation; a URL of a service

Actions: Posts the input to the service at the URL

Put

In: element, object, pool, or relation; a URL

Actions: Puts the input at the URL

Delete

In: A URL of an element, object, pool, or relation

Actions: Deletes the thing at the URL

GetIndexes

In: A URL of an element, object, pool, relation

Out: The indexes of the thing at the URL

PostIndex

In: An index; a URL of an element, object, pool, relation

Actions: Posts the index to the thing at the URL

PutIndex

In: An index; a URL

Actions: Puts the input at the URL

DeleteIndex

In: An index; a URL of an element, object, pool, or relation

Actions: Deletes the index of thing at the URL

**[00173]** The programming interface may also, or instead, include RSS-specific methods, such as:

ImportRSSFeed

In RSS Feed URL; import format (e.g. a schema, "HTML," etc.)

Out The feed formatted according to the import format

GetRSSFeedsByTag

In: A tag

Out: All RSS feeds tagged with the tag

**[00174]** The programming interface may also, or instead, include OPML-specific methods such as:

GetRelatedOPMLs

In URL of OPML file

Out OPML of all related OPML files

HostOPML\*

In OPML, opmlname

Out Hosted OPML URL

HostOPMLFromURL\*

In OPML URL, opmlname

Out Hosted OPML URL  
TagOPML  
In Hosted OPML URL, tag  
GetOPMLByTag  
In tag  
Out All OPML tagged with the tag  
CreateUserAccount  
In Username, Password  
Out Encrypted password key for username  
GetMyOPMLFiles  
In Username/Encrypted password  
Out OPML file of all hosted OPML files for username  
GetMyOPMLFile  
In Hosted OPML URL, encrypted password  
Out OPML  
ValidateOPML  
In OPML  
Out Valid/Invalid  
Validate OPMLByURL  
In OPML URL  
Out Valid/Invalid  
GetOPMLRSS  
In OPML URL  
Out RSS of all posts of contained feeds  
ReadOPML  
In OPML URL  
Out HTML of all posts of contained feeds  
IsOPMLContained  
In Parent OPML URL, Child OPML URL  
Out True/False of if the parent contains the child  
IsOPMLRelated



In Parent OPML URL, Child OPML URL  
Out True/False of if the two URLs are related

GetRelatedOPML

In OPML URL  
Out All OPML related to this one  
GetContainingOPML

In OPML URL  
Out OPML file of all parents of this OPML  
GetContainedOPML

In OPML URL  
Out OPML file of all descendents of this URL

[00175] The programming interface may also, or instead, include OPML search methods such as:

GetHTMLOPMLKeywordSearch

In keywords  
Out HTML Outlines containing keywords in the text

GetRSSOPMLKeywordSearch

In keywords  
Out RSS containing outlines containing keywords in their text

GetOPMLKeywordSearch

In keywords  
Out OPML containing outlines containing keywords in their text

GetOPMLPodcastSearch

In keywords, podcasts  
Out OPML with podcast feeds containing keyword(s) in their text

GetOPMLRSSSearch

In keywords, rss  
Out OPML with RSS feeds containing keyword(s) in their text

GetOPMLOPMLSearch

In keywords, opml  
Out OPML containing OPML outlines that contain keywords(s) in their text

GetCategoryByKeyword

In keyword, category

Out OPML of outlines containing keywords in their text that are not tree nodes

OPMLSearchFilterByUrl

In keyword, URL term

Out all the above results filtered to only show URLs containing the URL term

SearchOPMLByAttribute

In attribute name, keyword

Out return OPML of all OPML files containing the keyword in the value of attribute passed

**[00176]** It will be understood that the above methods are representative only, and that variations of the above methods may be suitably employed, including removal from or addition to the methods identified above. All such variations are intended to fall within the scope of this disclosure.

**[00177]** Thus, it will be appreciated that one general aspect of a system described herein includes a plurality of atomic functions for manipulation of OPML and RSS including search, presentation, navigation, publication, syndication, and so forth. These atomic functions may be exposed as individual services, as described generally above, or integrated into an OPML system, with a customized, web-based (or other) user interface for structured access to and use of OPML data. It will similarly be appreciated that the functionality described herein may be encapsulated in hardware such as a network server, a client computer, an integrated circuit, or a chip set.

**[00178]** A more generalized example of a useful arrangement of atomic functions for an OPML-based system is described below. In this example, atomic functions (or groupings of functions into atomic tool sets) are arranged around OPML creation, OPML validation, OPML publication, OPML search, OPML browsing, OPML reading, and subscription, which may be deployed using the architecture described above, or may form a set of core services 110 for an OPML-based metaservices system. These functions/groupings are discussed in greater detail below.

**[00179]** OPML creation: An OPML editor may be provided for creating and editing OPML files. An OPML manager may be provided for managing collections of OPML content

distributed across multiple files. OPML creation tools may include, for example, tools for migrating content into and out of OPML format, as well as reader/browser type tools for viewing OPML content. These OPML functions may be encapsulated in a functional module accessible to end users separately, or within an integrated OPML environment.

**[00180]** Publication/Validation: OPML content may be published at an OPML site, or directly from a client device. Publication may be in native OPML format, and/or may be suitably formatted and handled for syndication. For syndication purposes, a publication source may independently configure its own polling frequency or else use a remote, hosted ping API to notify other locations of content updates. The ping interface may be an XML-RPC standard API. A corresponding spider for related search and indexing may, for example, employ robots.txt conventions to flag content in the root domain of a source. The source URL may be identified to never be auto-pollled at a user's discretion. Auto-polling may occur at any suitable regular or irregular frequency, such as every 24 hours. Another tool that may be combined with publication tools or provided separately may validate OPML content for proper format, etc. These OPML functions may be encapsulated in a functional module accessible to end users separately, or within an integrated OPML environment.

**[00181]** Search: An OPML search engine may provide search capability across published OPML using, for example, the OPML search API's described above. A user may specify, for example, RSS, OPML, Podcasts, Categories, or the like. Once a user locates these types, the user may, through the interface, render the search results, as indicated in the interface with hyperlinks such as Read, Listen, View RSS feeds (this is so that you can preview a feed before you subscribe to it), and so forth. A user may also navigate to the OPML outline and content, such as using an OPML browser or an OPML reader. In addition, OPML files can be bookmarked within the interface to permit a user to return to bookmarked pages. These OPML functions may be encapsulated in a functional module independently accessible to end users, or within an integrated OPML environment.

**[00182]** Browse/Read: A browser interface and functionality may be provided for OPML files and content. In the browser, a user may navigate up and down a hierarchy of interrelated OPML content, and render leaf nodes containing, e.g., text, audio, video, and the like. Rendering engines may be provided for various media types. A user may also, or instead, directly read an OPML file, and navigate between OPML files through embedded references,

using, for example, a client-side or server-side OPML renderer. These OPML functions may be encapsulated in a functional module independently accessible to end users, or within an integrated OPML environment.

**[00183]** Subscription: A subscribe feature may allow a user to select a default reader for OPML search results. Subscription to a feed of RSS or other content identified in an OPML file may be encapsulated within that interface as a one-click operation with, e.g., a hyperlink or icon. OPML source files may also, or instead, be subscribed to through a one-click operation. These OPML functions may be encapsulated in a functional module independently accessible to end users, or within an integrated OPML environment. Reading lists may also be integrated into an OPML system. Reading lists may be OPML documents that point to RSS feeds. Rather than a typical RSS subscription, however, a reader or aggregator may subscribe directly to an OPML Reading list (or other document) itself. When the author of the OPML document adds a feed, the aggregator may automatically check that feed in its next scan, and when a feed is removed, the aggregator may stop checking that feed. The editor of an OPML file can thus update all subscribers by updating the OPML file. These OPML functions may be encapsulated in a functional module independently accessible to end users, or within an integrated OPML environment.

**[00184]** Each of these functions or functions sets, such as create, validate, publish, search, browse, and read may be deployed independently, e.g., as a web service, a client program, or a hosted service encapsulated within, e.g., a web page user interface or Application Interaction Interface. Each function or function set may, in certain embodiments, be accessed individually by end users, and groups of functions or function sets may be combined into an integrated interface for use by end users, either locally or hosted at a remote network location.

**[00185]** For example, the core services 110 and metaservices 120 may be adapted for use with medical records. A hospital directory may, for example, be constructed around OPML, with numerous data types and levels of hierarchy, all of which may be deployed in a conditional access environment for limited or controlled use of data and functions. The OPML metaservices system may be employed to permit custom interfaces for various users on top of the entire pool or environment of hospital data. This may include general information interfaces for the general public, patient interfaces with conditional access to records for a particular patient, physician interfaces with conditional access to data and functions (e.g., prescription

ordering) for groups of patients under treatment by a physician, administrator interfaces with conditional access to financial and payment data, and so forth.

**[00186]** The integrated system may be deployed as a private machine with access controlled by the creator using, e.g. password access to functions, function sets, source content, or integrated interfaces. Similarly, an OPML chip or physical device may integrate the functions and function sets into hardware. In another aspect, functions and function sets may themselves be deployed in a social network, as generally described above with respect to web superservices.

**[00187]** In a more general aspect, the platform described above may support a single point of contact for fixed services, extensible services, and/or ad hoc services. This generalized platform may be used to deploy new composite services created from various sources. The platform may, for example, be used to deploy a large-scale public aggregator that provides access control, searching, filtering, and clustering of content, or to deploy the OPML server/database described above.

**[00188]** In another aspect, the platform may provide an integrated system for managing semantic reference networks that arise from community-based, interactive collaboration and communication on a network such as the Internet. The integrated system may include classification schemes for naturally occurring structures such as labels, links, keywords, and so forth. In addition, the system may support conditional access, instrumentation to provide metrics for traffic and usage, security, and any of the semantic functions or other functions described above.

**[00189]** The term “service” and related terminology is now discussed in some detail as it relates to the systems disclosed herein. The term “superservice” is used in Application Integration and Middleware rubric to describe services that provide an API as a common service that replaces or masks other existing APIs. More generally, superservices may be understood as atomic, possibly canonical services that are released in a scalable, efficient, globally available form for re-use, combination, and re-composition into other services in a manner that requires no special activity by a user other than calling the superservice. Common superservices have evolved from, for example, various special purpose software that implements CRM, SCM, B2B, and other internal operational applications. These products usually operate across two or more operating systems, transaction processing systems, database management systems, application

servers, and/or networking layers. Examples include adapters for accessing ERP, CRM, or other third-party application packages. More generally, a superservice may be any highly scalable atomic function that can be exposed as a service. In one aspect, the system described herein provides a common platform and/or central point of contact for deploying new superservices formed from other services and superservices that exist as ad hoc, remote programming interfaces. Superservices are often recognized by decomposition of large, special-purpose software systems, or evolved by users who identify and address needs for services in a network environment, or may be derived from any other source. Superservices, along with other services and/or data or content sources, may be combined into composite services using, for example, a metaservice that provides a service for managing and combining services.

**[00190]** As used herein, the term “superservice” is intended to include the web superservices generally described above, as well as any other highly scalable, networked service that provides a front end for proprietary data and/or software such as enterprise systems. The term “metaservice” is intended to include a service for managing other services including for example storing, accessing, executing, testing, cataloging, indexing, discovering, searching, annotating, characterizing, combining, and/or publishing services and specifying interfaces therefore. The term “composite service”, as used herein, is intended to refer to a combination of services. As used herein, the term “service” is intended to more generally include any network-available service, including but not limited to the superservice, metaservice, and composite services described above, along with any other resource that might be cast as a service and made available through a network such as web services, search engines, mapping utilities, geolocation services, databases, dictionaries, RSS aggregators, spiders, and so forth, as well as mashups and other combinations of any of the foregoing.

**[00191]** Thus while a metaservice may be employed to arrange ad hoc services and core services into a specific application, such as an OPML database and/or server, the metaservice may more generally be used to provide any services such as web services, ad hoc services, superservices, composite services (combinations of superservices and/or other services, released as a new service) and/or metaservices (services for managing services, superservices, and/or composite services), thus enabling fully customizable, user-specified web services that combine any network-connected content and/or services. Thus any service, such as a front end for a database, may be combined with any other service, such as a mapping

system with an API, to provide unique user services. A server built from the core services 110 and database 111 may provide a metaservice for organizing and combining these services, and sharing new, composite services with others.

[00192] This type of services composition can be observed in mashups, including popular combinations of mapping and other information services currently available on the web. Consistent with this trend, many entities are releasing increasingly low level interfaces to web services, such as Google's maps, or eBay's auction site. However, mashups remain a fully custom, one-off technique for creating composite services from these low level interfaces. It will be appreciated that any number of useful combinations of these and other third party services may be created using the metaservices described herein. This may include combinations of any of the following services that are present on the World Wide Web, as well as any other services amenable to structured access: mapping, auctions, telephone directories, patent databases, Edgar/SEC corporate filings, online want-ads (such as Craigslist), search engines, location services for cellular phones, services directories (restaurants, hotels, museums, etc.), RSS syndicated content, news feeds, stock quotes, sports scores, dictionaries, real estate listings, electronic commerce, legal databases (statutes and case law), multi-player games, IRC/online chat, and instant messaging. This may also, or instead, include new (typically commercial) services which may be increasingly decomposed and liberated for use by the public, such as: ticket sales, reservation systems, equities trading, supply chain management, customer ordering, customer relationship management, inventory management, financial reconciliation, tax preparation, and human resources. This may also include new superservices emerging on the web. Virtually any service that is maintained within an enterprise or otherwise provided by computers may be decoupled from its environment and offered as a stand-alone superservice for combination with other services on the Internet. Interfaces to such services may be through scripting or other programmatic access to URLs or URIs with command line interfaces, RSS, OPML, XML, APIs (including SOAP), and/or any other input/output mechanisms through which such services may be rendered.

[00193] As described generally above, the metaservices 120 or core services 110 may provide for metering of access to services that are commercially available so that these services can be incorporated into a composite service on a pay-as-you-go basis. The core services 110

may also provide reports on usage, and may support automated or manual payment for usage of such commercial services.

**[00194]** In one aspect, there is disclosed herein a metaservice system for locating, manipulating, combining, and publishing services, web services, superservices, and/or content sources. Thus the OPML server, for example, may also function as a service server or metaservice platform through which individual, decomposed services are located, registered, and made available alone or in combination with other services through a server such as the server hosting the core services 110. A user interface may be provided for searching for services (or searching an index of services), for selecting and combining services, and for manually or automatically generating scripts or other portable instantiations of composite services which may be published, such as through the core services 110, for use by clients 112 connected to a network. In another aspect, the core services 110 and database 111 may operate as a search engine for location of services, superservices, composite services, or other metaservice servers. It will be noted that through a metaservice, composite services may themselves be created and syndicated, i.e., published for subscription and use by third parties.

**[00195]** In a social aspect, users of the core services 110 may communicate with one another and share various services, superservices, and combinations of services, which may, in turn, be layered into additional composite services. Additionally, the community of users may identify new services that are needed, which may be contributed by community members, or constructed from existing services. Thus, in an alternative embodiment to structuring content through the use of OPML outlines, the core services 110 may provide a metaservice platform for structuring services that combine services and/or other superservices for use through a network. Similarly, decomposed services may be canonically arranged and registered or stored within the database 111. The tools for supporting this type of social networking may be provided as the core services 110 or metaservices 120 described above, such as through the social network methods 810 of the application-aspect interface 802 thereof.

**[00196]** In another aspect, an interface provided by the metaservice platform may provide for user submissions of new services, and may provide a sandbox for testing new services, superservices, and composite services. The testing may ensure, for example, end-to-end integration and/or compatibility across various platform, hardware, and/or software. Thus, for example, the validation may ensure timeliness of updates or information, compatibility with



known web browsers, responsiveness of remote application programming interfaces, or compatibility with certain hardware for uploads (e.g., iPods, BlackBerry e-mail devices, Treos, cellular phones, etc.). The metaservice platform may also maintain a reference library of validated superservices meeting some performance criteria. Since the superservices themselves may have independent commercial value in such an environment, access to libraries of superservices may be fee based, using any number of known business models for electronic commerce or software licensing. These interfaces may be provided, for example through the infrastructure-aspect interface 1302 and/or the program method 804 of the application aspect interface 802 described above.

**[00197]** Significant advantages may be realized from a structured, human-readable approach to creating and deploying composite services that aggregate a number of different services to achieve a new service. As an example, OPML may provide a useful structure for describing an interrelationship of services to achieve a new composite service. More generally, any XML-based, plain text based, command line oriented, or other syntax capable of capturing hierarchy, chronology, structure, and the like in an outline or other suitable format may be usefully employed.

**[00198]** As a general example, services may be arranged in an outline that describes the manner in which they are combined. For example:

```
<COMPOSITE SERVICE>
  Item 1
    Subitem a
      Sub-sub item (i)
  Item 2
    Subitem b
    Subitem c
</COMPOSITE SERVICE>
```

**[00199]** The outline, or more generally, the conceptual structure within the outline, may also be expressed implicitly as a sequence of terms made available within a URL/URI. For example, the composite service described above can be written as an extension of or substitute for a URL with an in-line syntax to delimit components. Using, for example, an ampersand, the above expression may be stated as <COMPOSITE SERVICE>&Item 1&Subitem a&Sub-sub

item (i)&Item 2&Subitem b&Subitem c</COMPOSITE SERVICE>. Similarly (and consistent with IETF RFC 1738), variables for local action by a browser may be demarcated by a ‘#’ symbol.

**[00200]** Each element of the outline of a composite service may include a “name” and associated “value” or values. A name in this instance refers to a pre-defined variable and/or to a sub-action or sub-service that is to be invoked by that element. The value refers to but is not limited to a character, number, letter, word, term, list, array, cluster, object or any other kind of data element. The value may be inserted into the variable and/or used to condition the invocation of the action or service that is to be carried out by the element in the outline. The name and/or value may include elements of outlines, URI/URLs, and/or file names. For example, the name “search” might be associated with “General Electric” to invoke a search of a given data set for information related to General Electric Corporation. Additionally, name and value may be extracted from a file, a URL, and/or an element of text or other data stream, and this in turn may help condition the action or actions being invoked. For example, an image file or recorded music file or video file may have metadata encoded within the file itself, as is the case with ID3 data in music files.

**[00201]** A composite service may generate an outline as output from its action and/or as an output of any of its sub-actions. Outlines produced in this manner may in turn be used to invoke other services and to condition their action, and to direct the input of data into, and/or output of data from, the process or processes. Where none of the services provide persistence for this interim data, the data may be stored in the database 111 for the core services 110 (including, by way of example, as an RSS feed), or locally at a client device as discussed above. The data storage may be specified in the outline created, or may be specified along with the outline that specifies the composite service. In one aspect, services may each independently find a location to store interim data.

**[00202]** Each element of the outline of a composite service may refer to a specific service available on the network. The order may imply flow control for composition into a new service. Again referring to the example above, Item 1 may be performed by applying the results of Subitem a (which are in turn derived from Sub-sub item (i) to the service defined by Item 1. The output of Item 1 may be passed to the service defined by Item 2, which may receive an input that is the output of a sequential pre-processing by the services defined by Subitem b and

Subitem c. In addition to sequencing, flow control may be provided with additional delimiters for, e.g., concatenation or combination of outputs, branching, looping, conditional statements, exit conditions, return codes, and the like. Each item may be further defined using any number of required or optional parameters. For example:

Item

<service location>

<input(s)>

(e.g., source(s) or specific value(s), format, etc.)

<service parameters>

(specification of service(s))

(any optional or required command parameters for the service(s))

<output format>

1. data typing

2. output location

a) another web service

b) display (with, e.g., formatting/layout parameters)

c) remote data site

d) local machine folder

3. file type

a) audio

b) video

c) presentation document (Word, PowerPoint, etc.)

**[00203]** Where the service defined by the item is registered with a metaservice such as one of the metaservices described above, the parameters may be automatically reviewed, filtered, corrected, supplemented, or otherwise interpreted before invoking the service identified by the item. Thus a layer of intelligence may be provided by a metaservice for registered, or otherwise known or recognized services. Optionally, an unregistered service may be called blind, i.e., invoked by reference to a location with one or more strings of commands that are unconditionally passed to the identified location.

**[00204]** In one aspect, a composite service may take the form list/logic or attribute/value pairs. In a list/logic pair, the composite service grammar may specify locations

or lists coupled with logic. A list could, for example, include a URL, a source, a folder, a file, HTML code, HTML permalinks, source code, and so forth. More generally, the list may be any data or content at any location. The logic may specify one or more operations to perform on the list, or optionally, a service to receive the list. The logic may further be parameterized according to any corresponding capabilities of the service or logical operation receiving the list. This may include switches, parameters, options, and the like such as are conventionally found in a command line syntax or the like.

**[00205]** The core services 110 or metaservices 120 described above may include a metaservices engine such as a parsing service for analyzing and processing composite services, whether expressed as list/logic pairs or any other suitable syntax or grammar. Thus for example, the metaservices 120 may include a service for parsing, choreographing, and executing a composite service, and for post-processing any results therefrom. This metaservices engine may be invoked directly by passing a suitably formatted outline, or may be invoked internally by a user interface provided by the core services 110 (or metaservices 120), or some combination of these. In embodiments, a browser or similar program at a client 112 may be locally configured to provide a human-usable interface for accessing the metaservices engine. This interface may be generalized, or may be specific to a certain task, service, or function.

**[00206]** In one aspect, a virtual machine may be formed by a master list of tasks for the core services 110. The master list may include a list of tasks or elements, each of which may contain actual logic (e.g., code of any form) or abstract functional descriptions, or references to external sources of the foregoing. The master list may organize and schedule tasks. Much as a computer program executing on a computer, the core services 110 may parse and execute (as appropriate) elements of the master list in programmatic fashion to achieve a design objective. Using the techniques described above, the core services 110 may call remote services that consist of nothing more than an application programming interface available through a network. The core services 110 may orchestrate presentation of a suitably formatted request to the programming interface and retrieval of any output from the service. Thus, widely distributed and unstructured data and services may be marshaled to one or more programmatic objectives of the core services 110, which may be hosted at a metaservices server that provides a central point of contact for accessing and managing network services. As a significant advantage, this general architecture may accommodate various distributions of data and

processing, which may be optimized according to constraints such as data mobility or processing resources.

**[00207]** Other generalized computing concepts may be realized within the framework described above. Where a plurality of remote sites offer the same programming interface and services, the master list may employ parallelism and/or pipelining. Similarly, the master list may employ redundancy for important tasks. In various configurations, new tasks or logic may be expressed in the master list for execution, or deployed as a new service that can be invoked by a metaservices engine processing the master list. All such uses and variations are intended to fall within the scope of this disclosure.

**[00208]** In various embodiments, a composite service may execute locally on a client device that parses the structure of the master list, or the service may be created by a metaservice that orchestrates execution and provides any explicit or implicit flow control along with any required transient storage. Optionally, the metaservice may simply coordinate connections among the services without handling inputs and outputs except for a final result. Of course, this latter embodiment would require services that permit connections to be created among each other from a remote location, or otherwise provide for transient storage to support data persistence. A composite service may, for example, overlay or combine multiple outputs into a single, end-user display or data set.

**[00209]** Thus there is generally provided herein a programming language or syntax for creating, managing, invoking, searching, and syndicated composite services. The syntax may be expressed in OPML, or any other suitable grammar, and may provide for flow control, input/output management, parameterization of service calls, and the like for orchestration of a number of remote services into a composite service. The OPML (or other grammar) may be shared with others, who may use a particular composite service in combination with other services, or modify the composite service, or some combination of these. The composite service may be registered with a metaservice, or syndicated for third party use.

**[00210]** In another aspect, the metaservice may provide a forms-based system for creating composite services by providing a searchable database of registered services, along with forms that structure inputs, outputs, parameterization of service calls, and the like.

**[00211]** Unlike Universal Description, Discovery and Integration (UDDI), which relies on standardized protocols for a Web services registry, the systems described herein can

accommodate registration as a technique to simplify user access to registered services, while permitting reference to arbitrary services regardless of their relationship to a metaservice or other registration site. Further, the systems described herein may permit a human-readable document to describe the interrelationship and flow control of a number of separate services. In particular, a syntax such as OPML, or OPML supplemented by a flow control syntax, may be used to embody a composite service of remote services. In addition, whether intended for local or remote execution, the composite service description itself may be shared through direct transfer or syndication using, for example, RSS or any other syndication techniques described herein.

**[00212]** The composite services, and techniques for creating same, as described herein may be used to deploy a wide array of new services. For example, using a Wireless Access Protocol for a portable device, a user may provide location information, status, and the like, which may be converted to an OPML output that may be made available to OPML search engines. A composite service may be configured, for example, to retrieve information for specific individuals (such as through a search or filter of location/status OPML files) and present location information on a map (such as through the Google Map API) along with a link or call out to status information. Optionally, the output or display of data may be configured to show multiple user locations, along with buttons to select individuals on a pick list for whom location may be displayed. Or, if the user information includes a group or affiliation, the output may also include a control to select a particular group for display within a map. Where location data is also syndicated, a data feed may be used to reconstruct not only a location, but a map showing changes in location over time, or a path taken by the located individual.

**[00213]** Composite services may be created for enhanced aggregator functions. For example, a composite service may be configured to render an outline of syndicated feeds by displaying the outline structure, and rendering within that outline groups of items from each feed (such as an RSS feed) identified.

**[00214]** Composite services may be created for managing multimedia content. For example, a composite service may be configured to search for podcast content within an OPML data structure, and render the content as a list of podcast items, along with an address of a location for the item. As another example, a composite service may be configured to identify podcast content, apply a filter (which may be a remote service or program logic within the

composite service description) for suitable content, convert any responsive items into a single file format, and storing the converted items in a folder, which may be a remote storage folder or a folder on a user's local machine. A scheduler may re-execute the composite service either by prompting a user for a refresh or re-executing on a fixed timetable. As another example, the composite service described above may be applied to video content. The system may be extensible. For example, the description above mentions a filter, which would presumably be a filter for metadata associated with multimedia. However, an independent developer may develop a content filter that analyzes, e.g., audio content and creates a feature vector useful for measuring perceived similarity to other audio content. If the developer provides this functionality as a network-accessible programming interface, the content-based filter may be integrated into the podcast filter to identify, e.g., music that a user would probably like. As another example, a developer may create a technique for embedding media with a digital watermark that encodes data into the media. This may be employed to certify, identify, or log media as it is processed. This watermarking may be incorporated into any media processing through suitable incorporation of the corresponding remote, unstructured service.

**[00215]** Composite services may be used to prepare summary documents. For example, a spider or search engine may be applied to traverse an OPML tree structure, with results output to a file format such as PowerPoint, Word, or Excel. Parameters for such a service may include outline levels to be displayed, and an outline level at which pagination occurs (e.g., new page at each change in outline level 2).

**[00216]** Fig. 13 shows an architecture for user interaction with various distributed content. In general, the architecture 2300 includes an OPML-based content collection 2302, a rendering and conversion system 2304, an abstraction layer 2308, and a number of user interfaces 2310 on client devices 2312. While generally depicted as a bottom-up process of rendering content and services on client devices, it should be understood that, consistent with this disclosure, the process may move from client devices 2312 to the collection 2302, such as where a client device 2312 publishes new content or creates and releases a new service. Similarly, while the description below focuses on obtaining customized client views of unstructured content and services, these new views may themselves form new services or content that can be published to the collection 2302 for use by others. In the following description, the term "content" is used generally to refer to data in any format, as well as

services such as any of the services described above, and any other network-accessible, or potentially network-accessible resources.

**[00217]** The OPML-based content collection 2302 includes content of any type that might be organized into relationships using, for example, OPML. Other content may include documents (such as e-mail, calendar entries, spreadsheets, word processing documents, PDF, presentation documents (such as power point), and the like), services such as any of the services described above, multimedia (audio, video, animation, etc.), RSS or other syndicated formats, databases (including search engines) and any other electronic content, as well as additional OPML structures which may or may not be interrelated with one another. It will be appreciated that, while OPML is one convenient language for interrelating content into knowledge structures, any other suitable technologies may be employed such as other outlining languages, directory structures, relational databases, and so forth. More generally, the collection 2302 may be understood as a set of network-accessible content, along with an infrastructure for accessing and manipulating same. In one embodiment, this encompasses all content available on the Internet. In other embodiments, the content may be all content within an enterprise, or a subset of publicly available resources defined by access-control restrictions, individual preference, or the like.

**[00218]** The rendering and conversion system 2304 may be provided to accommodate the various content types available to the system. This may include rendering engines for various content using proprietary and open formats, as well as any number of conversion engines for converting content into a suitable form for end use. In addition, significant advantages may be realized by providing bi-directional converters for OPML (or any other language used to interrelate the content) so that knowledge structures may be readily ported into or out of the system. Thus, for example, an OPML structure may be converted into a power point presentation for purposes of communicating to others, or a Word document may be converted into an OPML outline. More generally, bi-directional converters may be usefully employed to enhance content (including services) creation options for client devices 2312. Suitable converters may be provided, for example, through the interpretation method 1108 of the semantic-aspect interface 1102, the media viewer method 808 of the application-aspect interface 802, the format-display method 904 of the client-aspect interface 902, or the data



transformation method 1010 of the data-aspect interface 1002, or any combination of these appropriate for a particular conversion type.

**[00219]** The abstraction layer 2308 may be employed to translate content between its native, distributed format and a form suitable for interaction in a user interface (such as the interfaces described with reference to Figs. 14 and 15 below). Thus for example, documents, services, multimedia and other content may be represented as icons or other user-operable objects that can be positioned within a user interface. OPML relationships may be converted to lines interconnecting such objects. Where the underlying content has associated metadata or properties, these may be presented within the object as buttons, drop-down menus, text, and so forth. Thus, for example, where a database object is presented through the abstraction layer, a table structure or a query interface may be presented directly within the object in a user interface, or in a manner accessible through the interface. Multimedia may be presented within a frame that includes controls such as volume, progress bar, forward, back, pause, color, and so forth. Services may be presented in an object that lists or provides access to API components.

**[00220]** In a navigation mode, a user interface 2310 may present content to a user in its abstracted form, with relationships shown within the interface 2310. The interface 2310 may enable management of the content by, for example, showing a certain number of layers within a hierarchy, and permitting a user to jump from node to node within a hierarchy. A user interface 2310 may also be preconfigured for certain types of data. For example, a health care information user interface might automatically provide a directory of hospitals, a topical map of high level categories, and a link to data restricted to access by treating physicians. The health care interface might also place conditional access controls at a top menu level for ready access, and may provide access to functional aspects of health care systems for suitably authorized users (for example, a prescription ordering system). By contrast, a general news interface might place filtering controls at a top level, along with an area for configuring paid content subscriptions. Each interface might default to specific OPML data sets or hierarchies. A user may also configure the interface according to personal preferences for rendering modes and tools that are provided within a menu hierarchy. In one aspect, the interface and elements thereof may be provided by the core services 110 and/or metaservices 120 described above.

**[00221]** In a manipulation mode, the user interface 2310 may permit a user to alter content. Thus for example, a new document may be added to the content by, for example,

dragging and dropping an icon into the interface, associating the icon with a local document, and connecting the icon (within the interface) to an OPML structure or another document already in the interface. The resulting document and association(s) may then be automatically passed through the abstraction layer (uploaded), passed through any appropriate rendering/conversion steps, and placed into the content collection 2302. Alternatively, an explicit publish command may be provided by the user. In a services example, a number of services within the network may be interconnected within the user interface to create a composite service as generally described above. The user interface may also include tools for validating and publishing such composite services. In another example, a database may be queried, with results passed to a service that outputs content which may be stored in a spreadsheet, which may be provided to a user through the user interface or republished into the OPML-based content. Any number of permutations are possible. Generally, the architecture provided herein contemplates access to and use of all such resources within a user interface which may be customized for various users and use types. Search engines, media converters, outlines, syndicated content, Web pages, and any other content, whether document-based or functional, may be viewed and manipulated.

**[00222]** Thus in one aspect there is provided herein a visualization tool for Internet content. The visualization tool provides a medium for viewing, manipulating, interrelating, and viewing relationships among various content. The tool may also provide configurable access to services. Views may be configured for different data types (e.g., health care, financial, news, sports, etc.), different professions (doctor, lawyer, accountant), and different data structures (e.g., OPML, structured databases, etc.). These views may be expressed as composite services that can be processed by a metaservice, and may be customized for individual use, and may be shared or published for third parties.

**[00223]** In another aspect, provided herein are visual design tools for manipulating web-accessible services.

**[00224]** In another aspect, provided herein is a design environment for functionally interconnecting web-based content.

**[00225]** In another aspect, provided herein is a visually oriented OPML manager providing tools for visualizing and manipulating OPML-based relationships and content.

**[00226]** In one embodiment, an additional functional layer may be added for post-processing content. In one implementation, data such as RSS data or other documents, may be processed to create organizational metadata such as an index, table of contents, list of figures/multimedia, bibliography, and the like, and this may be converted into an OPML structure that may be navigated using, e.g., the viewers described above. Thus in one aspect, disclosed herein is a system for automatic conversion of syndicated content or other data into OPML structures. The conversion may include searching, filtering, and clustering of syndicated content according to user parameters, as described generally, for example, in U.S. App. No. 11/223,826 and the documents referred to therein.

**[00227]** Fig. 14 shows a spreadsheet paradigm that may be used as a user interface 2400 for the navigation, management, and visualization tools described herein. It will be understood that the term spreadsheet, as used herein, refers to the visual organization of elements rather than functionality. At the same time, the spreadsheet interface 2400 may employ conventional spreadsheet nomenclature and operations to manipulate data and other items within the spreadsheet.

**[00228]** In the spreadsheet, a service such as a search may be parameterized using, e.g., values entered into cells of the spreadsheet. Cells may also, or instead, contain functional specifications, such as descriptions of Boolean operators, aggregation, filtering, output formats, conversions, mathematical operators, conditional statements, and so forth. These may be, for example written in a programming language specifically adapted for spreadsheet visualization, or using an existing programming language or syntax, by a creator of the spreadsheet or, they may contain interim or final results copied and pasted from other locations. In other embodiments, a cell may simply contain a reference to an external location where the desired service, function, parameter, or the like is present. Thus each cell may carry local content, or be defined with respect to other content. Similarly, each cell carries a global reference unique to the spreadsheet, so that it may be referenced from within the spreadsheet. These cells may also, or instead, be globally unique if the name of the spreadsheet can be uniquely identified within a global name hierarchy.

**[00229]** Thus, as depicted in Fig. 14, a user may create an area for filters within the spreadsheet. The top line of the filters area carries a title, which may or may not have functional significance within the spreadsheet. That is, the title line may simply carry text

inserted by the user to denote the significance of additional content within a column, row, or, as depicted, a bordered box of cells. In other embodiments, the word “Filter” appearing in a column may indicate that related adjacent cells are filter descriptions having a certain predetermined syntax. Similarly, a user may populate cells with spidering tools, formatting tools, source lists, media types, keywords, or any other relevant parameters for a spreadsheet. For example, each cell within the spidering tools list may include a title that is visible within the spreadsheet, along with an associated specification for a spider which may either be executed locally or deployed as a service accessible at a defined network address. Thus the functionality of a particular spider may be abstracted to a title meaningful to the user. A user may, as another example, populate “format outputs” cells within various format conversions, which may, in turn be parameterized with exception handling, pagination, viewer capabilities, and so forth. Each cell may, for example, be parameterized with content from other cells, such as by using media types to parameterize a filter or spider, or by using an output type to parameterize the format for results.

**[00230]** In one aspect, a composite function may be formed from other functions within the spreadsheet. Thus, for example, a current view may be constructed by parameterizing a search and a filter operation, and sending the output to, in this case, a region within the current page, using an output format designated in another cell. An output format may, for example, designate a content conversion, an output format, and related parameters. For example, an RSS-to-CSV formatted output may specify that only a source, content hyperlink, and title are to be presented. Thus the output of a spreadsheet may be as depicted in Fig. 14 – a list of relevant items, along with hyperlinks (including text and/or icons) to underlying content. This list may also be referenced by additional functions, such as a sorting function in another cell, which sorts according to some user-defined criteria and presents only the top five results. These results may also be used to populate a pre-defined region of the spreadsheet, or may be output to another medium such as an OPML document, a Word document or, where the content is multimedia content, to a portable device such as a cellular phone or iPod.

**[00231]** While a search is depicted, it will be appreciated that this methodology may be applied to any combination of services that combines databases, RSS feeds, OPML, web pages, web services, unstructured services, maps, API's, and any other resources that might be

available on a network, such as the services described above, and may be used to specify complex, composite services within an intuitive user environment. Also, while the structure of the “Current View” is depicted as a command line, it will be appreciated that the structure may be graphically depicted using a flow chart, state diagram, or other process-oriented graphical language.

**[00232]** In addition, the view itself may be constructed within a graphical user interface using drag-and-drop components, each of which may be user-defined and/or user parameterized. One example of a suitable graphical user interface is described below with reference to Fig. 15.

**[00233]** Fig. 15 depicts a user interface for a graphical environment for managing composite services as described herein. As depicted, the interface 2502 may provide drag-and-drop elements such as functions 2504 (which may include manually defined mathematical, Boolean, conditional, and/or flow related aspects, as well as references to external functions), inputs 2506 (which may include local data items, key-value or attribute-value pairs, lists, OPML, data sources (optionally with defined queries), and the like, as well as references to remote or external data) and outputs 2508 (which may include data formats, database commands, display types, document or media formats, device types, and the like, as well as addresses for output where an output file is created), all of which may be placed and interconnected by a user. Each object may be active, and may be activated such as by double clicking to invoke a menu, lists, command lines, or other elements that may be used to configure the selected item, such as providing data or pointers to local or remote data and/or services that might be used to form the composite service. In one aspect, the system may employ internal feedback, such as by supplementing a user-defined source list (“Source List 1”) with additional sources identified during a search, or by supplementing the Keywords with highly correlated terms.

**[00234]** The interface may also provide operational data, such as the last time and/or date that the output was updated, or a most recent date for inputs or externally referenced functions (e.g., remote services). A refresh command may be provided to permit a manual refresh of output. In addition, a user may configure the service to refresh periodically. Where the composite service created within the interface 2502 is to be published, a user may also provide, through the interface 2502, a description of the syntax for invoking the service, such as

the order and format of inputs. This description may be presented to external users through a variety of means, including without limitation the description URL described above. The interface 2502 may also provide a publication tool that permits the composite service, once designed and tested to the satisfaction of the author, to be published along with an automatically or manually generated API for accessing the composite service. In other embodiments, the publication tool may permit publication as a web application adapted for human use through a web browser or the like.

**[00235]** Thus, there is disclosed herein a graphical user interface for managing composite services. The interface may provide for creation, visualization, editing, and publication of composite services in web application or programming interface form. The interface may provide GUI access to any of the core services 110 or metaservices 120 described above, as well as other content and services, and may provide accompanying tools for validation and so forth.

**[00236]** The systems and methods described above may be used to deploy a wide array of network-based services. A number of specific examples of such services are provided below by way of illustration and not of limitation. It will be understood that, while the techniques described above may be usefully employed in the following examples, the systems and methods described below may also or instead use other techniques such as stand-alone applications, hosted web services, services-oriented architecture services, and so forth, as well as various combinations of these. All such combinations are intended to fall within the scope of this disclosure.

**[00237]** In one aspect, a calendar synchronization system uses syndicated data to transfer calendar items among calendars in disparate calendaring systems.

**[00238]** Fig. 16 shows a method for calendar synchronization, which may be deployed, for example, as one or more computer programs. The method 2600 may begin by accessing a first calendaring supported by a first calendaring system as shown in step 2602. The first calendaring system may be any calendaring system including local calendar applications such as Microsoft Outlook or any of a variety of open source or other client-side calendar applications, as well as calendars associated with other application programs that use time-oriented data such as financial programs (e.g., QuickBooks), scheduling and project management programs, and so forth. A client-side calendar application may also, or instead,

include a calendar application on a portable device such as a cellular phone, wireless electronic mail client, palmtop computer, iPhone, iPod, or the like. The calendaring system may instead be a web-based calendaring system such as Google calendar or iCalendar or an enterprise application that uses time-oriented data such as Salesforce. A calendar supported by such a system may include a variety of information.

**[00239]** Access to the calendar system may include providing a username and password for authenticated access. Access may also include establishing a network connection to the calendar system where the calendar system is remote from the system accessing the calendar. In one embodiment, a web-based service receives user information including a location of the calendar (e.g., network location, path, username, password, etc.), and establishes access using the user-provided information.

**[00240]** As shown in step 2604, the method 2600 may proceed to retrieve calendar items. This may include retrieving all calendar items from a calendar. The retrieval may also be parameterized so that only items within a specified date/time range are retrieved. The retrieval may also be filtered in any number of ways supported by calendar metadata, e.g., by location, by participant, or by any other data within the calendar system and/or calendar items. As used herein, the term calendar item is intended to refer to items such as appointments, due dates, and the like however represented by a calendaring system within a particular calendar. A calendar may have metadata applicable to the entire calendar such as an owner, file location information, access right information, and so forth. Each calendar item may also include calendar information. For example, the calendar information may include temporal information such as a time, a time period, a date, or a date range that specifies the time for a particular calendar item. The calendar information may include metadata such as a location, a description, participants, and so forth. Each item of metadata may have associated text, numbers, dates, or the like. The calendar information may also specify recurrence such as for birthdays, anniversaries, and the like. The calendar information may also include attachments such as documents, media (e.g., audio, video, pictures, etc.), and the like, as well as hyperlinks to relevant information. The calendar information may include reminder data specifying when and/or how to provide reminders concerned scheduled events. Still more generally, calendar information may include any information or links to information that might be relevant to an item in a calendar.

**[00241]** As shown in step 2606, the retrieved calendar items may be converted into a syndicated format such as RSS to provide a calendar feed. In one aspect, the various information types within the calendar system and/or calendar items may be tagged within the RSS feed for identification and transfer to other calendaring systems. For example, dates and times may be provided with corresponding element types within the RSS feed to facilitate automated processing. Similarly, attachments and other metadata may be similarly tagged.

**[00242]** As shown in step 2608, the calendar feed may be processed. This may include a variety of processing steps such as filtering, sorting, aggregating, and the like according to automated or user-provided criteria. Thus, processing may include filtering a calendar to remove unwanted items. For example, a user may request all sales leads having a dollar value greater than a specified amount, or may request all sales leads created in the last week, or all sales leads open for the coming week. As another example, a user may specify all meeting having certain individuals or relating to certain topics. As another example, processing may include supplementing the data feed with additional calendar items such as items of possible user interest (e.g., based on a user profile or explicitly specified or implicitly derived user interests), or advertisements which may be included as calendar items within the data feed.

**[00243]** In another aspect, processing may include aggregating a number of calendars by synchronization calendar data with data from one or more other calendars and/or calendaring systems (which may be similarly retrieved and converted to a syndicated feed as described above). In such an operation, conflicts among calendar items may be automatically or manually resolved. The processing may also, or instead retain conflicting information and present the information accordingly within the calendar system, such as in the form of two concurrent (and conflicting) appointments.

**[00244]** However processed, the resulting feed of calendar items may then be provided to a second calendar supported by a second calendaring system, as shown in step 2610. This operation may be supported by a syndication-to-calendar application connector that (with suitable user-provided credentials and location/path information) converts the calendar feed into calendar items for a target calendar system. The syndication-to-calendar application may be deployed as a hosted web service, which may be a free or fee-based service, or the syndication-to-calendar application may be deployed as a local application that subscribes to the processed calendar feed and locally re-formats items for insertion into a local calendar. In one



aspect, the second calendaring system may be different from the first calendaring system, such as where a Google calendar is migrated to an Outlook calendar or a Salesforce calendar is migrated to Google. In another aspect, the second calendaring system may be the same as the first calendaring system, such as for migrating from one Google calendar to another. In still another aspect, the first calendar and the second calendar may be the same, such as where a calendar is enriched by processing the syndicated calendar feed and returned to the source calendar.

**[00245]** In one aspect, the general method described above may be used to generate calendar items from non-calendar data. For example, one or more syndicated data feeds may be subscribed to and received at a location where the method 2600 is performed. The feed may be processed to identify items of interest, which may include time-sensitive or time-related data (e.g., new movie opening in theaters on a specified date, or conference about a topic of interest being held on a specific day) and to generate corresponding alerts. The alerts may be converted into calendar items and delivered to a calendar system for presentation within a calendar. In one aspect, the alerts may be provided in syndicated form, and added to a calendar using the syndication-to-calendar application connector described above. In another aspect, the calendar item may be time that is the current time so that the item is immediately due. The receiving calendar system may then generate any suitable reminders or other notifications to the calendar user.

**[00246]** In another aspect, the methods described herein may be applied more generally to perform a background data synchronization process across two or more data sets. In general, connecting to data sets may be performed as described above. The data sets may include any data from an enterprise system or application as noted above, including databases and database management systems. The processing may include synchronizing data between the data sets, which may be represented within the remote service as syndicated data feeds, or in any other XML or other suitable format. During this background synchronization, the process may identify a conflict among the at least two data sets requiring manual resolution, and may create a human-readable message requesting resolution of the conflict. This message may be communicated to a user over a human-readable communication channel such as a channel carrying cellular phone text messages (e.g., the short messaging service (SMS)), a channel carrying instant messages, a channel carrying electronic mail, or a channel carrying telephone

communications. The user may respond through the same communication channel, e.g., by replying to an electronic mail with an instruction that resolves the conflict, and the remote service may apply the response in the background synchronization process. In another embodiment, the user may respond through a different communication channel. By way of example, the remote service may generate a human-readable description and transmit the description to a user via e-mail. The e-mail may specify that a response is to be provided by calling a telephone number and responding to a numeric-menu based system to provide a selection among one or more conflict resolution options.

**[00247]** The background data synchronization process may synchronize data for a calendar, a contacts database, a mailing list, or more generally any database or collection of databases whether operated by an application, a web service, an enterprise system, or any other platform. The process may, for example, synchronize data between two web applications such as Google calendar and a calendar in Salesforce or a calendar in Microsoft Office Live. Similarly, the process may synchronize data between a web application and a client-based application, such as Google calendar and Microsoft Outlook.

**[00248]** In one aspect, a calendar-based advertising system may insert advertisements into a calendar.

**[00249]** In general, the advertising system (or method) may access calendars using the techniques described above. This may, for example, be performed during processing of syndicated data in the calendar synchronization process described above. More generally, any system for accessing a calendar within a calendar system, including any of the calendars and calendar systems described above, may usefully employ the following advertisement systems, generally in conjunction with access credentials willingly supplied for this purpose by the user of the calendar. In one embodiment, access to the calendar may be supplied in exchange for (otherwise) free use of the web-based calendar synchronization service described above. Using this access, the system or method may insert an advertisement into the calendar as a calendar item.

**[00250]** The calendar item may have a time associated with the occurrence of an advertised event, such as a television program, a show, a movie opening, a lecture, a presentation, a concert, a sports event, a web seminar, a web presentation, a sale, or any other event. In another aspect, the calendar item may have a time that is the current time so that the

calendar item appears for immediate attention and/or disposition by the user. The advertisement may include a video presentation, an audio presentation, a Flash media presentation, a hyperlink to a web page, or any other descriptive, promotional, or otherwise explanatory material. A calendar system that maintains the calendar may generate reminders responsive to the calendar item, with reminder characteristics of the calendar item selected according to a desired affect on the user. For example, if the advertisement indicates a time when tickets go on sale, a reminder may be provided fifteen minutes before the scheduled time. By contrast, if the advertisement indicates a movie opening, a reminder might more usefully be generated several hours in advance.

**[00251]** In one embodiment, the advertisements may be inserted into the calendar by subscribing the calendar or a supporting calendar system to a syndicated feed of advertisements, such as an RSS feed. Advertisements may be selected from a plurality of available advertisements according to user preferences, user demographics, user history, or any other known advertisement targeting techniques.

**[00252]** In one aspect, a password management method is provided for handling passwords at a remote service that operates as an intermediary between a user and a web service.

**[00253]** Fig. 17 shows a method for managing passwords at a remote service.

**[00254]** The method 2700 may begin with providing a remote service including an RSS interface to a secure enterprise system, as shown in step 2702. This service may operate as an intermediary between a user and an enterprise system that authenticates users. The service may access enterprise data on one hand, and provide RSS or other syndicated data feeds of the enterprise data on the other hand. While one example of such a service is the calendar synchronization system described above, it will be understood that the techniques described below may be more generally applied to any secure enterprise system or application from which data might usefully be extracted in a syndicated format. This includes, for example, contacts organizers, mailing lists, electronic mail systems, calendar systems, customer relationship management systems, enterprise resource planning systems, document management systems, content management systems, and so forth. In one aspect, the secure enterprise system may include any web application, web service, application integration interface or other resource that is accessible through the web using login credentials.

**[00255]** While the method described herein may be embodied as a remote service that provides an RSS interface to a secure enterprise system, it will be understood that the techniques described herein may more generally be applied in any context where a proxy is used between a client and a network service that receives login credentials.

**[00256]** As shown in step 2704 the service may receive login credentials for accessing the secure enterprise system. The credentials may include, for example a user name and password. The password may be encrypted for secure communication between the client and the remote service. The client may provide other information, such as a uniform resource locator (URL), internet protocol (IP) address, or other information that specifies a location of the enterprise system on a network. More finely grained information such as a path or other directory information may also be provided. In another aspect, the remote service may store location information associated with a username so that a user need only provide the login credentials to the remote service when accessing the secure enterprise system through the remote service.

**[00257]** As shown in step 2706, the remote service may validate a subscription to the remote service with the username. This may include, for example, accessing a database of subscription data to ensure that the associated user has paid for a current subscription to the remote service. Where no subscription exists, or an existing subscription has expired or is about to expire, corresponding messages may be transmitted to the client, either in human-readable or machine-readable form, for display to a user. In this manner, the remote service may operate as a fee-based service, such as for converting data from a customer relationship management system to a syndicated data feed, that independently validates subscribers.

**[00258]** As shown in step 2708, the remote service may then log in to the secure enterprise system using the username and password provided by the client. During this process, the password may be decrypted, and then encrypted for secure communication between the remote service and the secure enterprise system (which would not be able to decrypt the originally encrypted password from the client using conventional techniques). The service may locally cache a session identifier for the secure enterprise connection in order to maintain this connection across a number of client/user-side interactions.

**[00259]** In one aspect, the method may mitigate repetitive logins with bad credentials and other behavior that might cause the enterprise system to lock out a user associated with the

username. For example, before the remote service attempts to login with user-provided credentials, the service may determine whether the credentials have been previously and unsuccessfully used to log in to the enterprise system (or other network service). A compact method for maintaining this data at the remote service includes maintaining a database of usernames and corresponding hashes of passwords for each unsuccessful login attempt. Entries may be selectively added and removed according to the success of login attempts.

**[00260]** More specifically, when a username is not present, the service may proceed to attempt a login to the enterprise system. If this login attempt fails, the username and a hash of the password may be stored in the database. If the username is present in the database, a hash of the user-provided password may be compared to the stored password hash. If the stored password hash is the same as a hash of the user-provided password, then the service may provide a responsive human-readable or machine-readable error message to the client/user and cease further attempts to log in to the secure enterprise system. If the stored password hash is different from the hash of the user-provided password, then the system may attempt a new login. If successful, then the database entry corresponding to the username may be deleted. If unsuccessful, the database entry corresponding to the username may be updated to include the hash of the new user-provided password.

**[00261]** Stated differently, a method for managing passwords at a proxy to a secure network service may include the following conditional steps: when a username is re-submitted with the same password, a human-readable message may be created and returned to the user/client. When a username is located in the database with a different password (as determined by comparison of hashes), the remote service may remove the username and password hash from the database and attempt a new log in to the secure enterprise system. Using this technique to determine when login credentials have been previously and unsuccessfully applied to log in, appropriate responsive action may be taken by the remote service. When the login credentials have not been previously and unsuccessfully applied to log in, the remote service may transmit the login credentials to the network service. Conversely, when the login credentials have been previously and unsuccessfully applied to log in, the remote service may transmit an error message to the user.

**[00262]** As shown in step 2710, the remote service may then provide an RSS interface (or other syndicated format interface) for the secure enterprise system to the user at the

client. The RSS interface may employ the secure connection to the enterprise system on one hand and the client connection on the other hand to exchange requests and responses between the client and the secure enterprise system, with the client-side connection using a syndicated data format. This interface may be used in a number of ways. For example, the user may, through the interface, obtain a data feed of information from the secure enterprise system such as contact information, news items, scheduled events, or any other data maintained within the secure enterprise system. In addition, the syndicated data feed may be processed as generally described above to filter, enrich, sort, or otherwise manipulate data within the data feed for use by the user. This may include, for example, adding ordered OPML lists of items, such as expenses in an expense report, or any other data. In another aspect, the remote service may receive RSS or other syndicated data and submit the data to the secure enterprise system through the RSS interface. It will be appreciated that, using the techniques described herein, a number of platform independent application connectors may be provided. By way of example and not of limitation, a local electronic mail client may be synchronized with a user's enterprise electronic mail account. As another example, a local time keeping client may be coupled to an enterprise time tracking system for use by professional services personnel such as consultants, lawyers, accountants, engineers, and the like. In a customer relationship management system, the client may send and retrieve sales leads or other sales information through the RSS interface of the remote service. Where the secure enterprise system includes a calendar, calendar entries (or "items" as described above) may be exchanged with the system using the RSS interface of the remote service. Where the secure enterprise system includes a database or database management system, database operations such as queries, reads, writes, joins, and so forth may be conducted through the RSS interface of the remote service.

**[00263]** In another aspect, the systems and methods described above may be employed to provide dynamic feed generation from enterprise system data. The process may begin by receiving a request at the remote service from a client to access an enterprise system. The remote service may then access the enterprise system using, for example, login credentials provided by the user. Where the system is not secure, no login credentials would be required. The remote service may query the enterprise system using any parameters provided by the user, thus obtaining enterprise data, and the enterprise data may be converted by the remote service into a syndicated feed. The syndicated feed may then be provided to the client such as by

publishing the feed to a predetermined or user-specified location. In general, converting the enterprise data into a syndicated feed may include additional processing such as filtering, sorting, categorizing, formatting, tagging, or other processing of the enterprise data.

**[00264]** In this method, object identifiers in the enterprise data may be dynamically interpreted by the remote service. For example, the enterprise data may include an object identifier having a type such as a name, a customer name, a customer number, a customer contact, a date, an address, and so forth. The remote service may respond to the object identifier such as by performing a supplemental request to the enterprise system for additional data relating thereto and inserting the additional data into the syndicated data feed. The supplemental request may be based on the type of the object identifier or the content of the object identifier, or both. In this manner, non-human readable information may be converted into human-readable form before presentation of a data feed to the user. For example and not by way of limitation, a ten digit customer code may be converted into a textual company name along with address and/or other information, or an internal numerical date representation may be converted into a mm/dd/yy or other representation of a date for downstream processing and display.

**[00265]** In another aspect, the remote service may interpret user requests before presenting them to the enterprise system. For example, where the request from the client includes at least one relative term, the remote service may convert the relative term into a specific term for presentation to the enterprise system. This may include, for example, a range of times (e.g., this year, this month, last month, today, next three days, etc.).

**[00266]** In another aspect, the remote service may incorporate a syntax for interpretation during presentation of the feed. For example, the syntax may reference a local variable determined from the client (e.g., display size, user name, report title, etc.). The syntax may identify a value as a date which may be interpreted for presentation according to user preferences. Similarly, the syntax may identify dollar amounts, foreign currencies, percentages, or any other values having one or more formatting options.

**[00267]** In another aspect, the remote service may convert enterprise system and/or remote server error messages (e.g., failed login attempt, server busy, server unavailable, lapsed subscription to a service) into human-readable form for textual incorporation into an output data feed. Where the error message relates to a failure that prevents the remote service from

obtaining enterprise data, the message may also, or instead, be communicated through a separate medium such as an electronic mail, an instant message, or a cellular phone text message.

**[00268]** In another aspect, converting the enterprise data into a syndicated feed may include mapping the enterprise data to at least one RSS extension. For example, the enterprise data may be encoded into custom namespace elements of an RSS output that correspond to input to another application or enterprise system. In this manner, the resulting syndicated data feed may be constructed to facilitate data transfer between applications, such as by selecting namespace elements that correspond to a metadata or data structure for a target system. The target system may in general be a second enterprise system or a local application.

**[00269]** Other systems and methods may also be deployed using the techniques described above. For example, a syndication platform may provide a usage-based prioritization of feed items. As another example, surge-based advertising may deploy advertisements in response to substantially real-time changes in user activity. These methods are now described in greater detail.

**[00270]** Fig. 17 shows a usage-based post prioritization process. The process 2800 provides an automated, popularity-based filter with very low computational overhead.

**[00271]** The process 2800 may begin by publishing a first feed with a plurality of items as shown in step 2802. The items may contain any syndicated content including for example, aggregated content from a variety of syndication sources, posts from various sources such as music by individual artists submitted as audio attachments to RSS items, content from a single syndication source such as news items from a media outlet, or any other items of combination of items suitable for publication as a data feed.

**[00272]** Access to the plurality of items may be monitored as shown in step 2804. This step may employ a variety of monitoring techniques. For example, access may be monitored by direct access to logs maintained by the feed publisher, or access may be determined using any available third party usage tracking services, or a monitoring or traffic management method associated with the infrastructure aspect interface 1302 described above. The monitoring may relate to a number of views for each item, such as total views or unique views. In another aspect, the monitoring may relate to the number of citations to or from the post. For example, this may include locating and analyzing links to each of the items, or



monitoring clickthroughs on links within the item. More generally, any usage or popularity metric may be usefully employed to monitor access to the items.

**[00273]** The remote accesses may be analyzed to determine a most-viewed subset of the plurality of items as shown in step 2806. The most-viewed subset may be determined for a certain time period such as one hour, one day, one week, or one month. The subset may be a fixed number such as ten items, twenty items, or one-hundred items. Analyzing the items may include filtering the items to restrict the most-viewed subset. This filter may, for example restrict the subset to items containing multimedia attachments (e.g., audio, video, etc.), or may remove unwanted or inappropriate posts before selecting the most-viewed subset.

**[00274]** The most-viewed subset of the plurality of items may be published as a second feed as shown in step 2808. The feed may be sorted according to the selection metric (e.g., views) to provide an ordered feed of popular items.

**[00275]** Fig. 19 shows a surge-based advertising process. The process 2900 may be deployed using, for example a monitoring system, a presentation system, and an advertisement delivery system. The monitoring system may monitor traffic patterns at web logs and report usage statistics in substantially real time. The monitoring system may analyze traffic patterns to identify a surge in activity. A number of traffic monitoring techniques are known in the art and may be usefully employed with the systems described herein. A presentation system may provide a human-readable display of the surge, such as by showing an ordered list of access to various web logs, or displaying a color-coded network or geographic map that displays increases in activity through color differentiation. An advertisement delivery system may provide manual and/or automated response to a surge by selecting and deploying relevant advertisements. Each of these systems is described in greater detail below.

**[00276]** It is generally known to gather statistics that measure and characterize online traffic such as web page views or downloads. Commercial services provide related usage statistics in real time or near real time, which may in turn be used to target advertisements. However, conventional targeted advertising uses metrics such as aggregate traffic that may fail to capture valuable targeting opportunities in environments such as the blogosphere, which exhibits highly dynamic usage patterns marked by rapid increases and decreases around certain locations, or certain subject matter. Current online advertising techniques exhibit a number of deficiencies with respect to such activity. For example, the value of an ad placement varies

dynamically. As another difficulty, changes in value are not conveniently visible to advertisers. In addition, advertisers cannot respond readily to changes in value of advertisements. There remains a need for improved advertising systems for use in a weblog environment.

**[00277]** A system described herein employs tracking of surges in online activity to target advertisements to rapid changes in online activity. As used herein, a surge may be generally understood as a rapid increase in activity at a location, or related to a particular topic or keyword. This type of swarming activity is common around weblog topics of current interest, and may focus on a current event, a book, an individual, a political issue, and so forth. A platform for delivering advertisements within a surge of online activity may include a real-time (or near real time) monitoring system, a presentation system, and an advertisement delivery system.

**[00278]** The monitoring system may use any service or combination of services that provides monitoring and analysis of both traffic patterns and semantic patterns on the web. The value of an internet ad placed by a particular advertiser on a particular site or cluster of sites goes up and down depending on multiple factors of what may be referred to as surge, which may reflect one or more of traffic, subscriptions, demographics, referral activity, reader influence, reader location, reader navigation patterns, momentary popularity, and semantic characteristics such as relationship to a topic, persistence, innovation trajectory, importance of a topic to other topics, sites, or readers, and so forth.

**[00279]** The presentation system may provide specialized presentation of patterns and their implications to advertisers and web publishers. In particular, an advertiser may wish to assess the immediate value of a site to his or her intentions. This may depend, for example on factors such as the rate of change in usage, or the semantic content of a current surge of activity. Thus the presentation system may usefully present information concerning one or more of the following: momentary information about surge; a “surge map” of one or more major internet sites; specialized surge maps customized for a particular advertiser based upon, for example, semantic content of a current surge; tools for drilling into particular surges and doing specialized analysis including but not limited to examining the semantics and the references involved, the people and organizations involved, and their intentions and collaboration in the current situation.

**[00280]** The delivery system may provide tools for advertisers and publishers to respond to surges in real-time by modifying placement of ads and the content of ads. The advertiser may wish to adjust his or her placement of advertising across sites and clusters of sites, as well as modify messages to respond to momentary opportunities. Thus, the media-buying platform for delivering advertisements should permit adjustments and response at a speed corresponding to the speed of the dialogues in the blog space and the associated news space and any other part of the associated media. A suitable delivery system may include, for example, the ability for the advertiser to purchase ads on a momentary basis, the ability for the advertising market platform to manage reserve advertising capacity in order to enable yield management, like done by airlines, so that there is typically some advertising capacity available at premium prices at the last minute, at least on sites likely to be part of surges, and the ability for advertisers to change their messages in real-time, according to what they believe will be effective given the current surge.

**[00281]** The process 2900 may begin by monitoring traffic patterns and semantic patterns on the web, as shown in step 2902 and as generally described above. Traffic patterns may include, for example, traffic, subscriptions, identity, navigation, rate of change in traffic, demographics, and change in demographics. Semantic patterns may include identity of referrals, source popularity, content, persistence of a topic, trajectory of a topic, relationship of topic to location of readers, and relationship of topic to web location, and any other patterns that can be identified on the basis of semantic content of communications.

**[00282]** As shown in step 2904, a surge may be detected in web activity based upon the monitoring. The surge may include, for example, a rapid increase in activity around a current event, a book, an individual (such as a politician or celebrity), a political issue, a natural disaster, and so forth. It will be appreciated that a number of techniques for quantitatively characterizing change are known in the art, and may usefully be employed to identify surges as described herein. This may include, for example, continuous calculation of a second derivative (change in the rate of change) of traffic to identify acceleration in activity or to spot discontinuities in activity. This may also, or instead, include relative adjustments so that a change from one view per day to five views per day is not mis-interpreted as a surge based upon a mathematical 500% increase. At the same time, surges may be identified using a combination

of semantic and traffic activity, such as by detecting an increase in activity at some or all of the web sites or web logs relating to a particular topic.

**[00283]** Regardless of how a surge is quantitatively detected and displayed, advertisements may be deployed in response to the surge as shown in step 2906. This may include modifying the placement of the advertisement, such as by requesting a more prominent position within a page or other display, or by requesting a placement at a locus of the surge or a related site. Deploying an advertisement in response to the surge may also, or instead, include modifying content of the advertisement, such as by modifying text, graphics, or the like according to the semantic content of the surge. Deploying an advertisement in response to the surge may also, or instead, include selecting one of a plurality of ads according to semantic content of the surge. For example, if the surge relates to a specific topic such as automobile accidents, a user may select advertisements relating to automobile safety products, or the user may select from among a number of different ads for a product based on automobile safety themes within a particular ad.

**[00284]** In general, the order of steps in the methods described above may be varied unless a particular order is explicitly required or otherwise clear from the context. In addition, it will be understood that, while depicted as single, linear processes, the methods described above may be performed multiple times using multiple processing threads, processors, computers, servers, or the like, and may be repeated any number of time as appropriate for a particular application. It will further be understood that additional steps may be included, and in certain instances, steps may be excluded. All such variations as would be apparent to one of ordinary skill in the art are intended to fall within the scope of this disclosure.

**[00285]** It will be appreciated that the above processes, and steps thereof, may be realized in hardware, software, or any combination of these suitable for a particular application. The hardware may include a general purpose computer and/or dedicated computing device. The processes may be realized in one or more microprocessors, microcontrollers, embedded microcontrollers, programmable digital signal processors or other programmable device, along with internal and/or external memory. The processes may also, or instead, be embodied in an application specific integrated circuit, a programmable gate array, programmable array logic, or any other device that may be configured to process electronic signals. It will further be appreciated that the process may be realized as computer executable code created using a

structured programming language such as C, an object oriented programming language such as C++, or any other high-level or low-level programming language (including assembly languages, hardware description languages, and database programming languages and technologies) that may be stored, compiled or interpreted to run on one of the above devices, as well as heterogeneous combinations of processors, processor architectures, or combinations of different hardware and software. At the same time, processing may be distributed across a camera system and/or a computer in a number of ways, or all of the functionality may be integrated into a dedicated, standalone device or other hardware. All such permutations and combinations are intended to fall within the scope of the present disclosure.

[00286] It will also be appreciated that means for performing the steps associated with the processes described above may include any of the hardware and/or software described above. In another aspect, each process, including individual process steps described above and combinations thereof, may be embodied in computer executable code that, when executing on one or more computing devices, performs the steps thereof.

[00287] While the invention has been disclosed in connection with the preferred embodiments shown and described in detail, various modifications and improvements thereon will become readily apparent to those skilled in the art. Accordingly, the spirit and scope of the present invention is not to be limited by the foregoing examples, but is to be understood in the broadest sense allowable by law.

## CLAIMS

What is claimed is:

1. A method comprising:  
accessing a first calendar, the first calendar supported by a first calendaring system and the first calendar containing one or more calendar items;  
converting the one or more calendar items into a syndicated format to provide a calendar feed;  
processing the calendar feed; and  
providing the calendar feed to a second calendar supported by a second calendaring system.
2. The method of claim 1 wherein providing the calendar feed to the second calendaring system includes passing the calendar feed to a syndication-to-calendar application connector.
3. The method of claim 1 wherein processing the calendar feed includes adding a calendar item.
4. The method of claim 1 wherein processing the calendar feed includes filtering the calendar feed.
5. The method of claim 1 wherein processing the calendar feed includes aggregating at least two calendars into a single feed.
6. The method of claim 1 wherein processing the calendar feed includes converting additional calendar items from the second calendar system into a second calendar feed and resolving one or more conflicts between the calendar feed and the second calendar feed.

7. The method of claim 1 wherein the first calendaring system includes a system selected from the group consisting of Microsoft Outlook, Google calendar, iCalendar, and a Salesforce calendar.

8. The method of claim 1 wherein the second calendaring system includes a system selected from the group consisting of Microsoft Outlook, Google calendar, iCalendar, and a Salesforce calendar.

9. The method of claim 1 wherein the first calendaring system is different from the second calendaring system.

10. A computer program product comprising computer executable code that when executing on one or more computer devices performs the steps of:

accessing a first calendar, the first calendar supported by a first calendaring system and the first calendar containing one or more calendar items;

converting the one or more calendar items into a syndicated format to provide a calendar feed;

processing the calendar feed; and

providing the calendar feed to a second calendar supported by a second calendaring system.

11. The computer program product of claim 10 wherein providing the calendar feed to the second calendaring system includes passing the calendar feed to a syndication-to-calendar application connector.

12. The computer program product of claim 10 wherein processing the calendar feed includes adding a calendar item.

13. The computer program product of claim 10 wherein processing the calendar feed includes filtering the calendar feed.

14. The computer program product of claim 10 wherein processing the calendar feed includes aggregating at least two calendars into a single feed.
15. The computer program product of claim 10 wherein processing the calendar feed includes converting additional calendar items from the second calendar system into a second calendar feed and resolving one or more conflicts between the calendar feed and the second calendar feed.
16. The computer program product of claim 10 wherein the first calendaring system includes a system selected from the group consisting of Microsoft Outlook, Google calendar, iCalendar, and a SalesForce calendar.
17. The computer program product of claim 10 wherein the second calendaring system includes a system selected from the group consisting of Microsoft Outlook, Google calendar, iCalendar, and a SalesForce calendar.
18. The computer program product of claim 10 wherein the first calendaring system is different from the second calendaring system.
19. A method comprising:
  - receiving a syndicated feed;
  - processing the syndicated feed to generate an alert;
  - converting the alert into a calendar item; and
  - delivering the calendar item to a calendar system.
20. The method of claim 19 wherein the calendar item is a scheduled event having a time that is the current time.
21. A method comprising:
  - providing a remote service including an RSS interface to a secure enterprise system;



receiving login credentials for accessing the secure enterprise system, the login credentials including a username and a password, the password being encrypted for secure communication between a client and the remote service;

validating a subscription to the remote service with the username;

decrypting the password;

encrypting the password for secure communication between the remote service and the secure enterprise system;

logging in to the secure enterprise system; and

providing the client with the RSS interface to the secure enterprise system.

22. The method of claim 21 wherein the secure enterprise system includes a customer relationship management system.

23. The method of claim 22 further comprising providing sales leads to the client through the RSS interface.

24. The method of claim 21 wherein the secure enterprise system includes a calendar.

25. The method of claim 24 further comprising providing calendar entries to the client through the RSS interface.

26. The method of claim 21 wherein the secure enterprise system includes a database.

27. The method of claim 21 wherein the secure enterprise system includes a web application.

28. The method of claim 21 further comprising locally caching a session identifier at the remote service in order to maintain a session between the remote service and the secure enterprise system.

29. The method of claim 21 wherein the remote service includes a fee-based service for converting data from a customer relationship management system to a syndicated data feed.

30. The method of claim 29 wherein the syndicated data feed includes one or more calendar items.

31. A computer program product comprising computer executable code that when executing on one or more computer devices performs the steps of:

- providing a remote service including an RSS interface to a secure enterprise system;
- receiving login credentials for accessing the secure enterprise system, the login credentials including a username and a password, the password being encrypted for secure communication between a client and the remote service;
- validating a subscription to the remote service with the username;
- decrypting the password;
- encrypting the password for secure communication between the remote service and the secure enterprise system;
- logging in to the secure enterprise system; and
- providing the client with the RSS interface to the secure enterprise system.

32. The computer program product of claim 31 wherein the secure enterprise system includes a customer relationship management system.

33. The computer program product of claim 32 further comprising computer executable code that performs the step of providing sales leads to the client through the RSS interface.

34. The computer program product of claim 31 wherein the secure enterprise system includes a calendar.

35. The computer program product of claim 34 further comprising computer executable code that performs the step of providing calendar entries to the client through the RSS interface.

36. The computer program product of claim 31 wherein the secure enterprise system includes a database.

37. The computer program product of claim 31 wherein the secure enterprise system includes a web application.

38. A method comprising:

receiving login credentials from a user at a proxy for logging in to a network service, the login credentials including a username and a password;

determining whether the login credentials have been previously and unsuccessfully applied to log in to the network service;

transmitting the login credentials to the network service when the login credentials have not been previously and unsuccessfully applied to log in to the network service; and

transmitting an error message to the user when the login credentials have been previously and unsuccessfully applied to log in to the network service.

39. The method of claim 38 wherein determining whether the login credentials have been previously and unsuccessfully applied includes maintaining a database of usernames and corresponding hashes of passwords for each unsuccessful login attempt.

40. The method of claim 39 wherein maintaining the database includes removing a username and corresponding hash of a password when the username is submitted with a password having a hash that does not match the hash in the database.

41. A method comprising:

receiving a request from a client to access an enterprise system;

accessing the enterprise system;

querying the enterprise system to obtain enterprise data;

converting the enterprise data into a syndicated feed; and

providing the syndicated feed to the client.

42. The method of claim 41 further comprising:  
identifying an object identifier in the enterprise data;  
performing a supplemental request to the enterprise system for additional data relating to the object identifier; and  
inserting the additional data into the syndicated feed.
43. The method of claim 42 wherein the supplemental request is based upon a type of object specified by the identifier.
44. The method of claim 43 wherein the type includes one or more of a name, a customer number, a date, and an address.
45. The method of claim 41 wherein the request from the client includes at least one relative term, the method further comprising converting the relative term into a specific term for presentation to the enterprise system.
46. The method of claim 45 wherein the relative term specifies a range of times.
47. The method of claim 46 wherein the range of times is one or more of this year, this month, last month, today, and next three days.
48. The method of claim 41 wherein converting the enterprise data into a syndicated feed includes using a syntax for formatting during presentation of the feed.
49. The method of claim 48 wherein the syntax specifies one or more of a local variable, a date, a dollar amount, and a percentage.
50. The method of claim 41 wherein converting the enterprise data into a syndicated feed includes converting at least one error into a human-readable message.

51. The method of claim 50 wherein the human-readable message describes one or more of an inability to reach the enterprise system, a failure of login credentials, and a lapsed subscription to a service.
52. The method of claim 50 wherein the human-readable message is communicated through a medium including one or more of an electronic mail, an instant message, and a cellular phone text message.
53. The method of claim 41 wherein converting the enterprise data into a syndicated feed includes mapping the enterprise data to at least one RSS extension.
54. The method of claim 41 wherein converting the enterprise data into a syndicated feed includes encoding the enterprise query into one or more custom namespace elements of an RSS output, the custom namespace elements providing a machine-readable form to facilitate transfer of the enterprise data to an application.
55. The method of claim 54 wherein the application is one or more of a local application and a second enterprise system.
56. The method of claim 41 wherein converting the enterprise data into a syndicated feed includes one or more of filtering, sorting, categorizing and formatting the enterprise data.
57. A computer program product comprising computer executable code that when executing on one or more computer devices performs the steps of:
- receiving a request from a client to access an enterprise system;
  - accessing the enterprise system;
  - querying the enterprise system to obtain enterprise data;
  - converting the enterprise data into a syndicated feed; and
  - providing the syndicated feed to the client.

58. The computer program product of claim 57 further comprising computer executable code that performs the steps of:

- identifying an object identifier in the enterprise data;
- performing a supplemental request to the enterprise system for additional data relating to the object identifier; and
- inserting the additional data into the syndicated feed.

59. The computer program product of claim 58 wherein the supplemental request is based upon a type of object specified by the identifier.

60. The computer program product of claim 59 wherein the type includes one or more of a name, a customer number, a date, and an address.

61. A method comprising:

- performing a background data synchronization process across at least two data sets;
- identifying a conflict among the at least two data sets requiring manual resolution;
- creating a human-readable message requesting resolution of the conflict;
- communicating the human-readable message over a human-readable communication channel;
- receiving a response to the human readable-message resolving the conflict over the human-readable communication channel; and
- applying the response in the background data synchronization process.

62. The method of claim 61 wherein the human-readable communication channel carries cellular phone text messages.

63. The method of claim 61 wherein the human-readable communication channel carries instant messages.

64. The method of claim 61 wherein the human-readable communication channel carries electronic mail.

65. The method of claim 61 wherein the human-readable communication channel carries telephonic communications.

66. The method of claim 61 wherein the background data synchronization process synchronizes data for one or more of a calendar, a contacts database, and a mailing list.

67. The method of claim 61 wherein performing the background data synchronization process includes synchronizing data for two or more web applications.

68. The method of claim 61 wherein performing the background data synchronization process includes synchronizing data between a web application and a client-based application.

69. A method comprising:

performing a background data synchronization process across at least two data sets;  
identifying a conflict among the at least two data sets requiring manual resolution;  
creating a human-readable message requesting resolution of the conflict;  
communicating the human-readable message over a first human-readable communication channel;

receiving a response to the human readable-message resolving the conflict over a second human-readable communication channel, the second human-readable communication channel being different from the first human-readable communication channel; and

applying the response in the background data synchronization process.

70. The method of claim 69 wherein at least one of the first human-readable communication channel and the second human-readable communication channel carries cellular phone text messages.

71. The method of claim 69 wherein at least one of the first human-readable communication channel and the second human-readable communication channel carries instant messages.

72. The method of claim 69 wherein at least one of the first human-readable communication channel and the second human-readable communication channel carries electronic mail.

73. A computer program product comprising computer executable code that when executing on one or more computer devices performs the steps of:

- performing a background data synchronization process across at least two data sets;
- identifying a conflict among the at least two data sets requiring manual resolution;
- creating a human-readable message requesting resolution of the conflict;
- communicating the human-readable message over a human-readable communication channel;

- receiving a response to the human readable-message resolving the conflict over the human-readable communication channel; and

- applying the response in the background data synchronization process.

74. The computer program product of claim 73 wherein the human-readable communication channel carries cellular phone text messages.

75. The computer program product of claim 73 wherein the human-readable communication channel carries instant messages.

76. The computer program product of claim 73 wherein the human-readable communication channel carries electronic mail.

77. The computer program product of claim 73 wherein the human-readable communication channel carries voice traffic.

78. The computer program product of claim 73 wherein the background data synchronization process synchronizes data for one or more of a calendar, a contacts database, and a mailing list.



79. The computer program product of claim 73 wherein performing a background data synchronization process includes synchronizing data for two or more web applications.
80. The computer program product of claim 73 wherein performing a background data synchronization process includes synchronizing data between a web application and a client-based application.
81. A method comprising:  
accessing a calendar through a programming interface; and  
inserting an advertisement into the calendar as a calendar item.
82. The method of claim 81 wherein the calendar item has a time associated with an occurrence of an advertised event.
83. The method of claim 82 wherein the advertised event is one or more of a sale, a television program, a show, a presentation, and a web presentation.
84. The method of claim 81 wherein accessing the calendar includes obtaining access in exchange for use of a free web-based calendar synchronization service.
85. The method of claim 81 wherein the advertisement includes one or more of a video presentation, an audio presentation, a Flash media presentation, and a hyperlink to a web page.
86. The method of claim 81 wherein the calendar item has a time that is the current time.
87. The method of claim 81 further comprising generating a reminder for the advertisement from a calendar system that maintains the calendar.
88. The method of claim 81 wherein the calendar is maintained by a web-based calendar system.

89. The method of claim 81 wherein accessing the calendar includes subscribing the calendar to an RSS feed of advertisements.

90. The method of claim 81 wherein the advertisement is selected from a plurality of advertisements according to one or more user preferences.

91. A computer program product comprising computer executable code that when executing on one or more computer devices performs the steps of:  
accessing a calendar through a programming interface; and  
inserting an advertisement into the calendar as a calendar item.

92. The computer program product of claim 91 wherein the calendar item has a time associated with an occurrence of an advertised event.

93. The computer program product of claim 92 wherein the advertised event is one or more of a sale, a television program, a show, a presentation, and a web presentation.

94. The computer program product of claim 91 wherein accessing the calendar includes obtaining access in exchange for use of a free web-based calendar synchronization service.

95. The computer program product of claim 91 wherein the advertisement includes one or more of a video presentation, an audio presentation, a Flash media presentation, and a hyperlink to a web page.

96. The computer program product of claim 91 wherein the calendar item has a time that is the current time.

97. The computer program product of claim 91 further comprising computer executable code that performs the step of generating a reminder for the advertisement from a calendar system that maintains the calendar.

98. The computer program product of claim 91 wherein the calendar is maintained by a web-based calendar system.

99. The computer program product of claim 91 wherein accessing the calendar includes subscribing the calendar to an RSS feed of advertisements.

100. The computer program product of claim 91 wherein the advertisement is selected from a plurality of advertisements according to one or more user preferences.

101. A method comprising:

publishing a first feed, the first feed including a plurality of items;

monitoring remote accesses to the plurality of items;

analyzing the remote accesses to identify a most-viewed subset of the plurality of items;

and

publishing a second feed, the second feed consisting of the most-viewed subset of the plurality of items.

102. The method of claim 101 wherein the most-viewed subset is determined for a time period of one hour.

103. The method of claim 101 wherein the most-viewed subset is determined for a time period of one day.

104. The method of claim 101 wherein the most-viewed subset is determined for a time period of one week.

105. The method of claim 101 wherein the most-viewed subset is determined for a time period of one month.

106. The method of claim 101 wherein the most-viewed subset consists of ten items.

107. The method of claim 101 wherein monitoring remote accesses to the plurality of items includes measuring a number of views.

108. The method of claim 101 wherein monitoring remote accesses to the plurality of items includes measuring a number of citations.

109. The method of claim 101 wherein analyzing the remote accesses includes filtering the plurality of items to identify one or more of the plurality of items containing multimedia attachments.

110. The method of claim 109 wherein the second feed consists of items having at least one multimedia attachment.

111. A computer program product comprising computer executable code that when executing on one or more computer devices performs the steps of:

- publishing a first feed, the first feed including a plurality of items;
- monitoring remote accesses to the plurality of items;
- analyzing the remote accesses to identify a most-viewed subset of the plurality of items;

and

- publishing a second feed, the second feed consisting of the most-viewed subset of the plurality of items.

112. The computer program product of claim 111 wherein the most-viewed subset is determined for a time period of one hour.

113. The computer program product of claim 111 wherein the most-viewed subset is determined for a time period of one day.

114. The computer program product of claim 111 wherein the most-viewed subset is determined for a time period of one week.

115. The computer program product of claim 111 wherein the most-viewed subset is determined for a time period of one month.

116. The computer program product of claim 111 wherein the most-viewed subset consists of ten items.

117. The computer program product of claim 111 wherein monitoring remote accesses to the plurality of items includes measuring a number of views.

118. The computer program product of claim 111 wherein monitoring remote accesses to the plurality of items includes measuring a number of citations.

119. The computer program product of claim 111 wherein analyzing the remote accesses includes filtering the plurality of items to identify one or more of the plurality of items containing multimedia attachments.

120. The computer program product of claim 119 wherein the second feed consists of items having at least one multimedia attachment.

121. A system comprising:

a monitoring system that monitors traffic patterns at a plurality of web logs in substantially real time and provides an analysis of the traffic patterns including an identification of a surge in activity;

a presentation system for displaying the surge in human-readable form; and  
an advertisement delivery system for deploying an advertisement in response to the surge.

122. The system of claim 121 wherein the surge includes a rapid increase in activity around one or more of a current event, a book, an individual, and a political issue.

123. The system of claim 121 wherein the traffic patterns include one or more of traffic, subscriptions, identity, navigation, rate of change in traffic, demographics, and change in demographics.

124. The system of claim 121 wherein the analysis of traffic patterns includes an analysis of semantic patterns that include one or more of identity of referrals, source popularity, content, persistence of a topic, trajectory of a topic, relationship of topic to location of readers, and relationship of topic to web location.

125. The system of claim 121 wherein deploying the advertisement includes modifying a placement of the advertisement.

126. The system of claim 125 wherein modifying the placement of the advertisement includes directing the advertisement at a locus of the surge.

127. The system of claim 121 wherein deploying the advertisement includes modifying content of the advertisement.

128. A method comprising:  
monitoring traffic patterns and semantic patterns on the web;  
detecting a surge in web activity in substantially real time based upon the traffic patterns and the semantic patterns; and  
deploying an advertisement in response to the surge.

129. The method of claim 128 wherein the surge includes a rapid increase in activity around one or more of a current event, a book, an individual, and a political issue.

130. The method of claim 128 wherein the traffic patterns include one or more of traffic, subscriptions, identity, navigation, rate of change in traffic, demographics, and change in demographics.

131. The method of claim 128 wherein the semantic patterns include one or more of identity of referrals, source popularity, content, persistence of a topic, trajectory of a topic, relationship of topic to location of readers, and relationship of topic to web location.

132. The method of claim 128 wherein deploying the advertisement includes modifying a placement of the advertisement.

133. The method of claim 132 wherein modifying the placement of the advertisement includes directing the advertisement at a locus of the surge.

134. The method of claim 128 wherein deploying the advertisement includes modifying content of the advertisement.

135. A computer program product comprising computer executable code that when executing on one or more computer devices performs the steps of:

- monitoring traffic patterns and semantic patterns on the web;
- detecting a surge in web activity in substantially real time based upon the traffic patterns and the semantic patterns; and
- deploying an advertisement in response to the surge.

136. The computer program product of claim 135 wherein the surge includes a rapid increase in activity around one or more of a current event, a book, an individual, and a political issue.

137. The computer program product of claim 135 wherein the traffic patterns include one or more of traffic, subscriptions, identity, navigation, rate of change in traffic, demographics, and change in demographics.

138. The computer program product of claim 135 wherein the semantic patterns include one or more of identity of referrals, source popularity, content, persistence of a topic, trajectory of a topic, relationship of topic to location of readers, and relationship of topic to web location.

139. The computer program product of claim 135 wherein deploying the advertisement includes modifying placement of the advertisement.

140. The computer program product of claim 139 wherein modifying placement of the advertisement includes directing the advertisement at a locus of the surge.



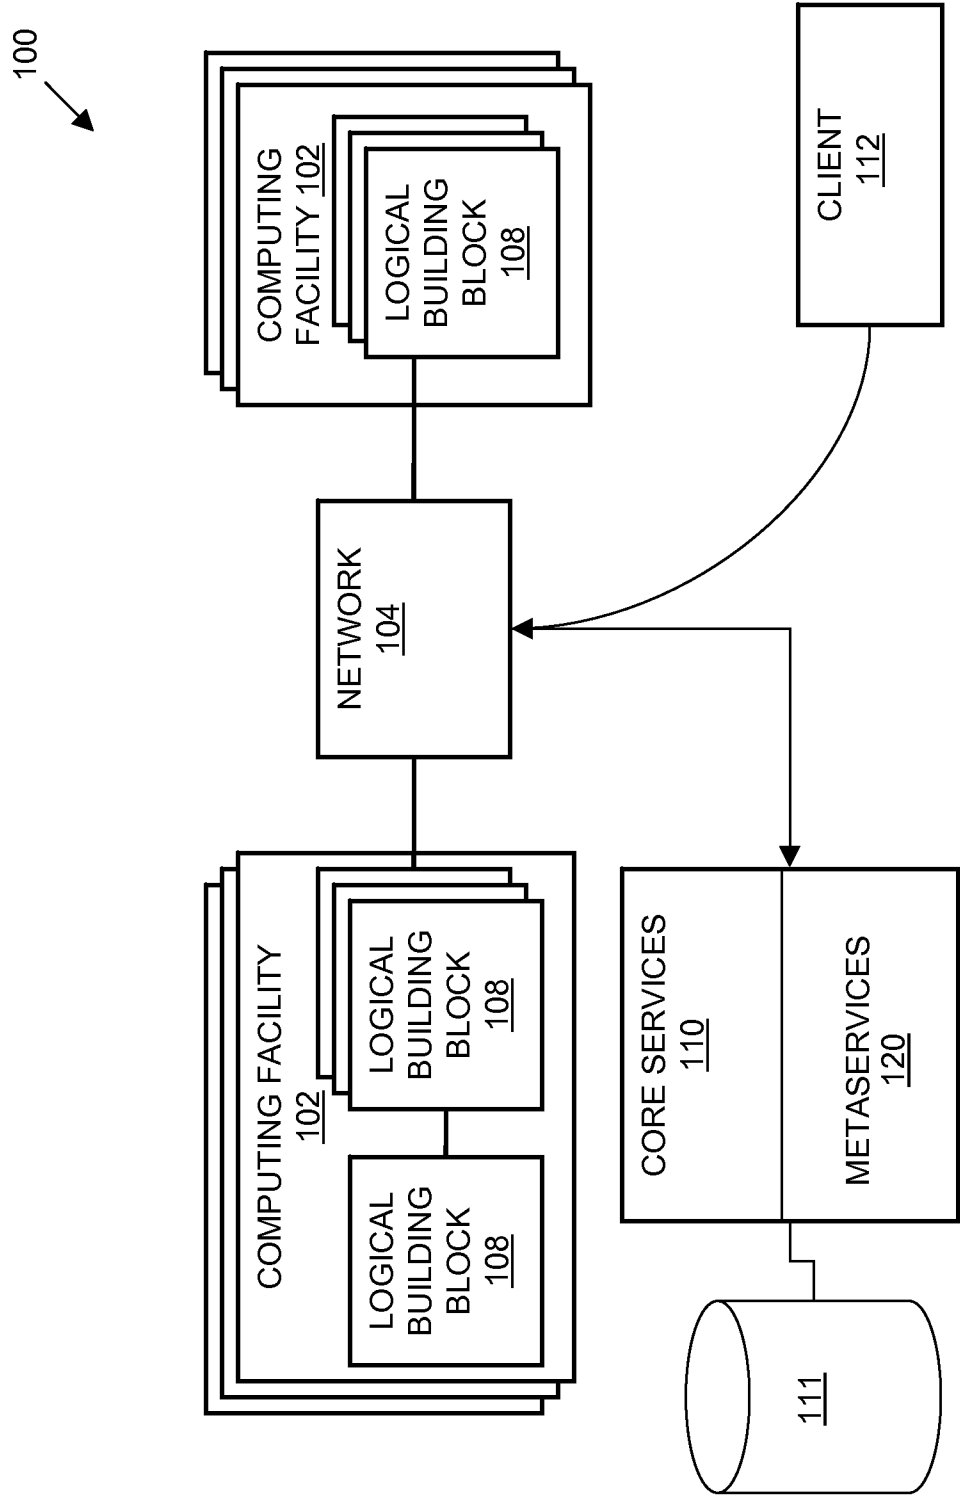


Fig. 1

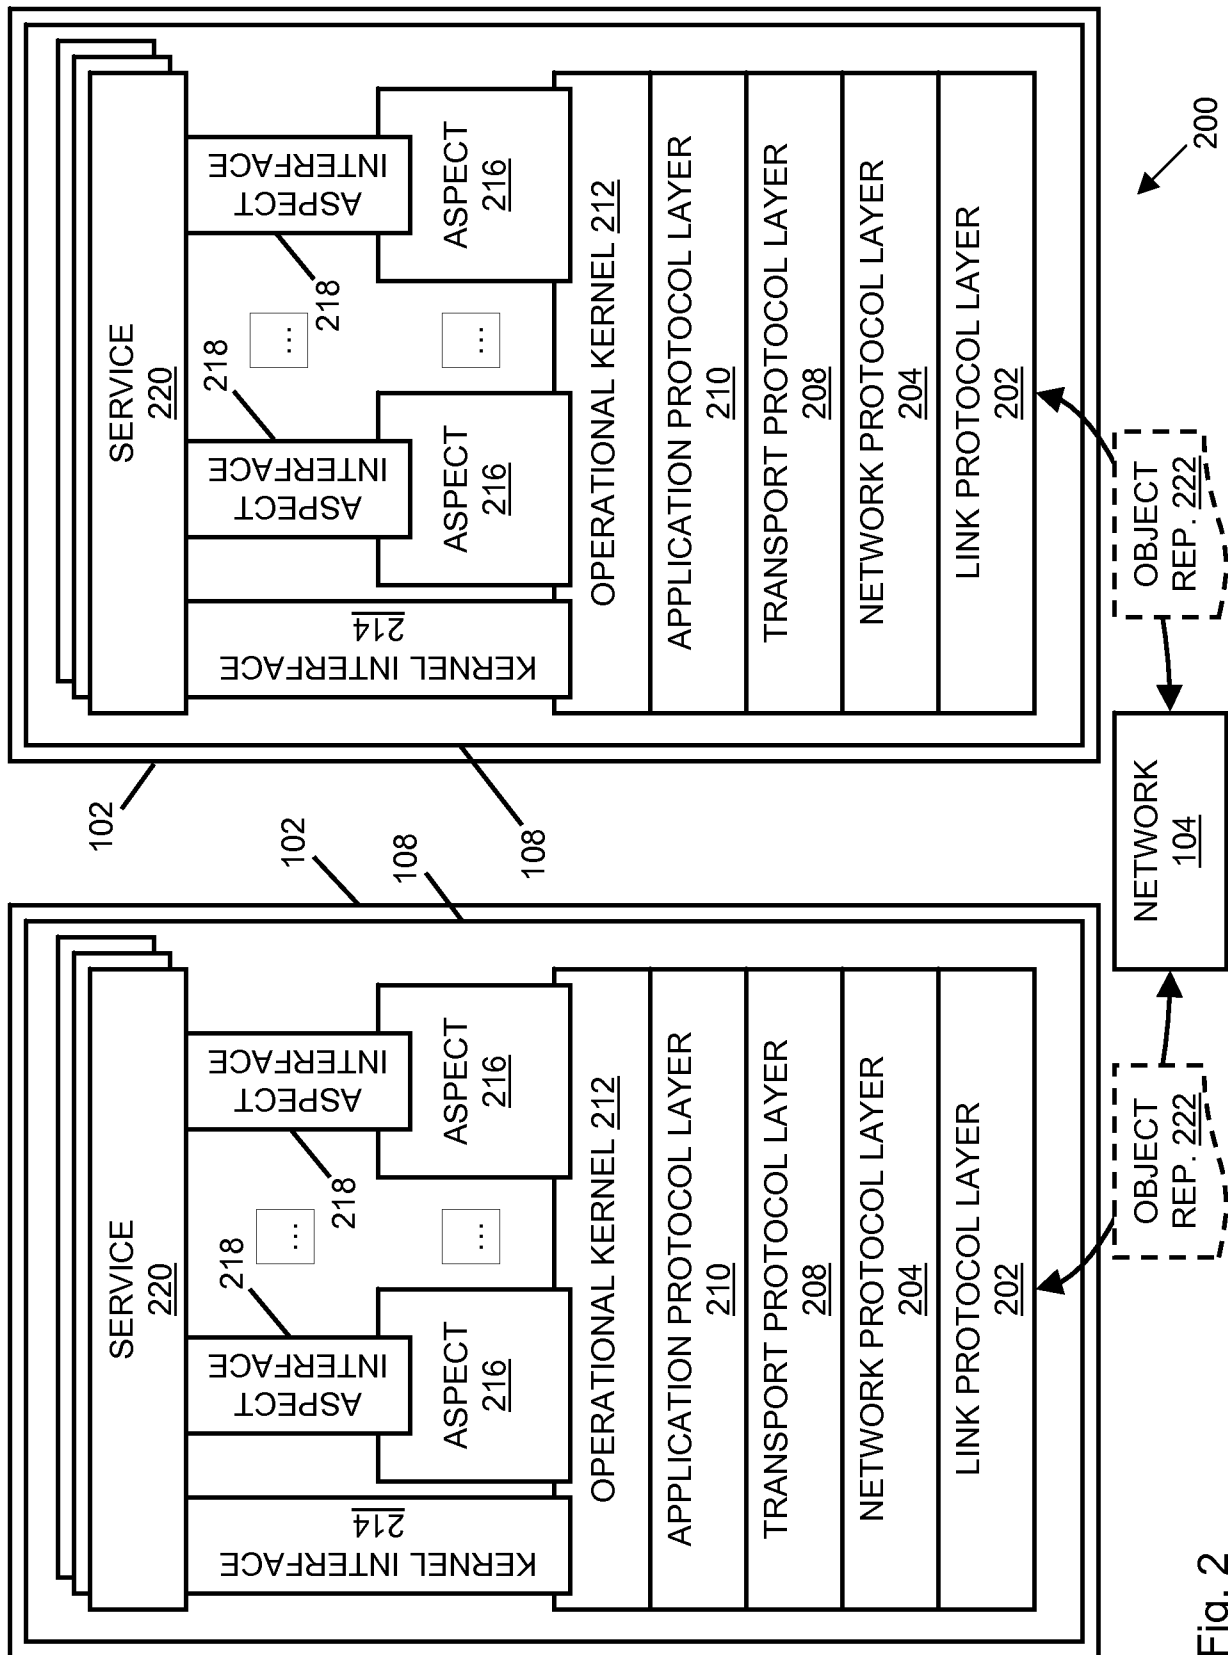


Fig. 2

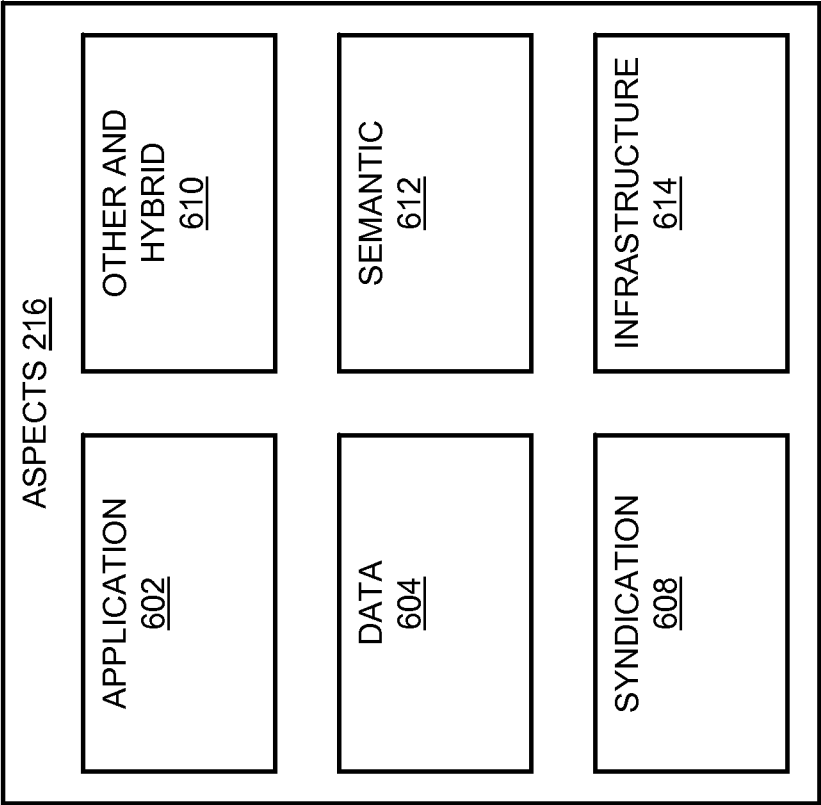


Fig. 3

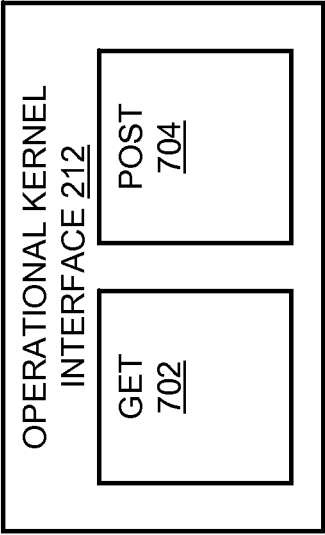


Fig. 4

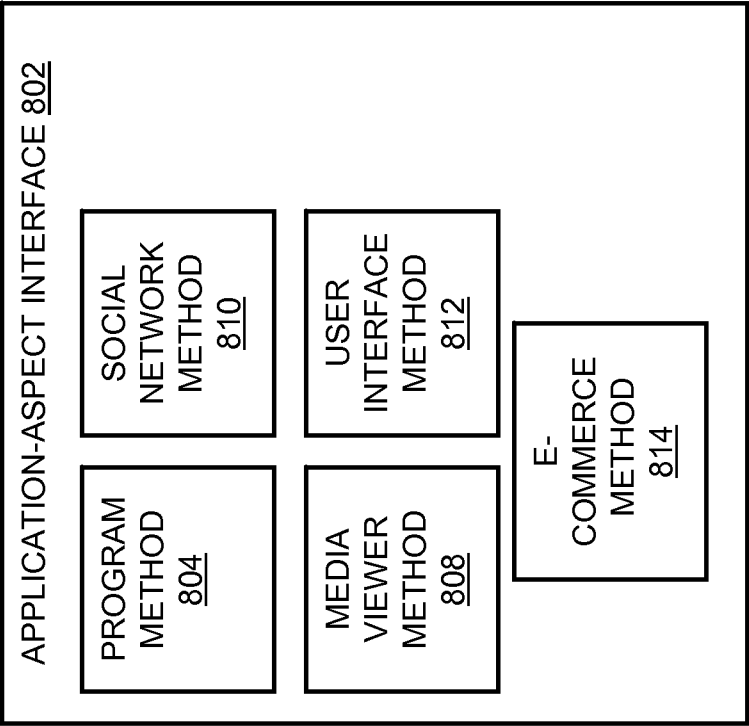


Fig. 5

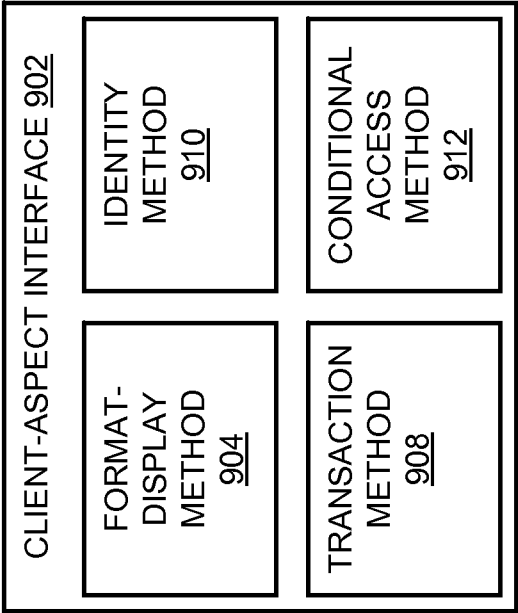


Fig. 6

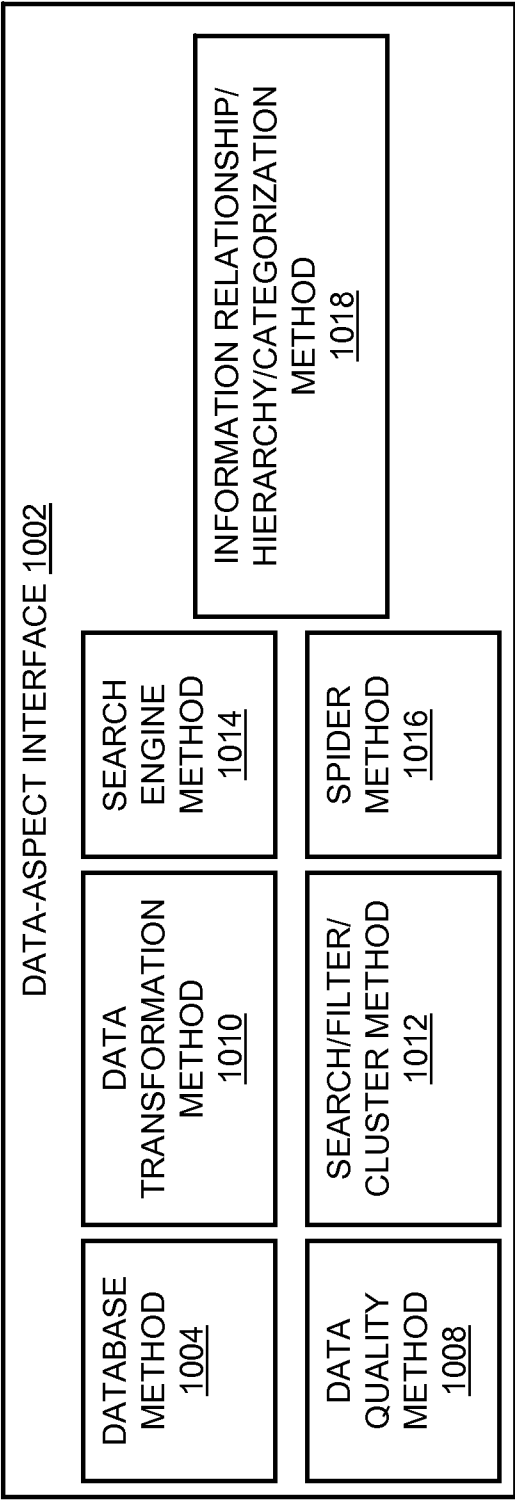


Fig. 7

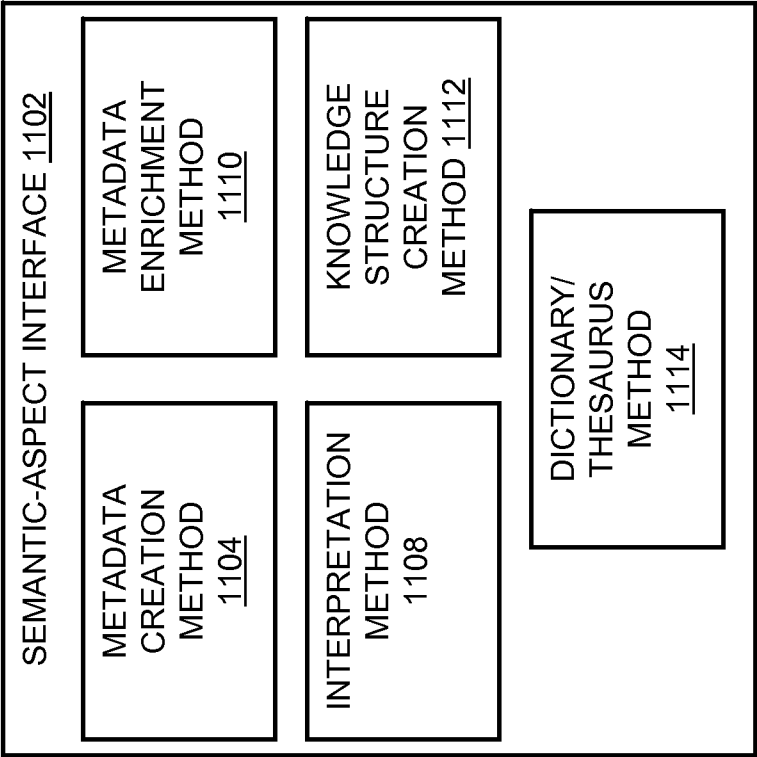


Fig. 8

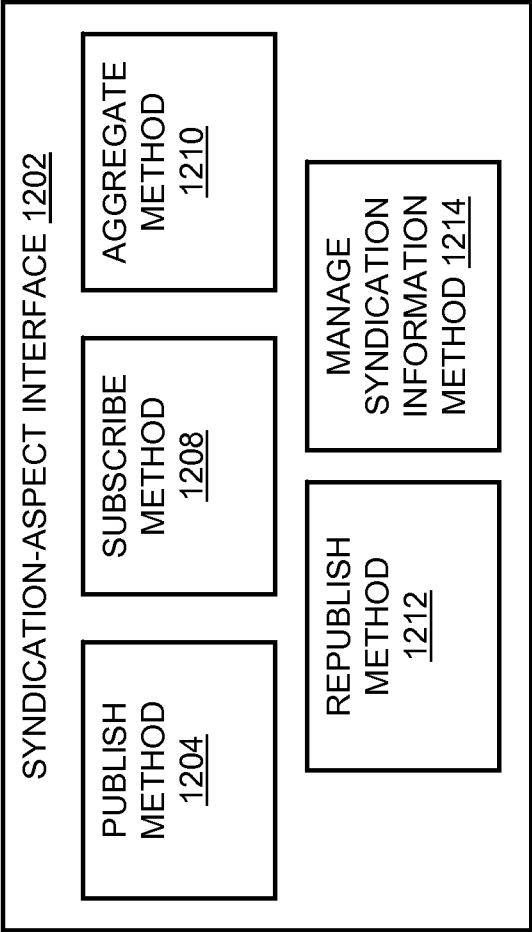


Fig. 9

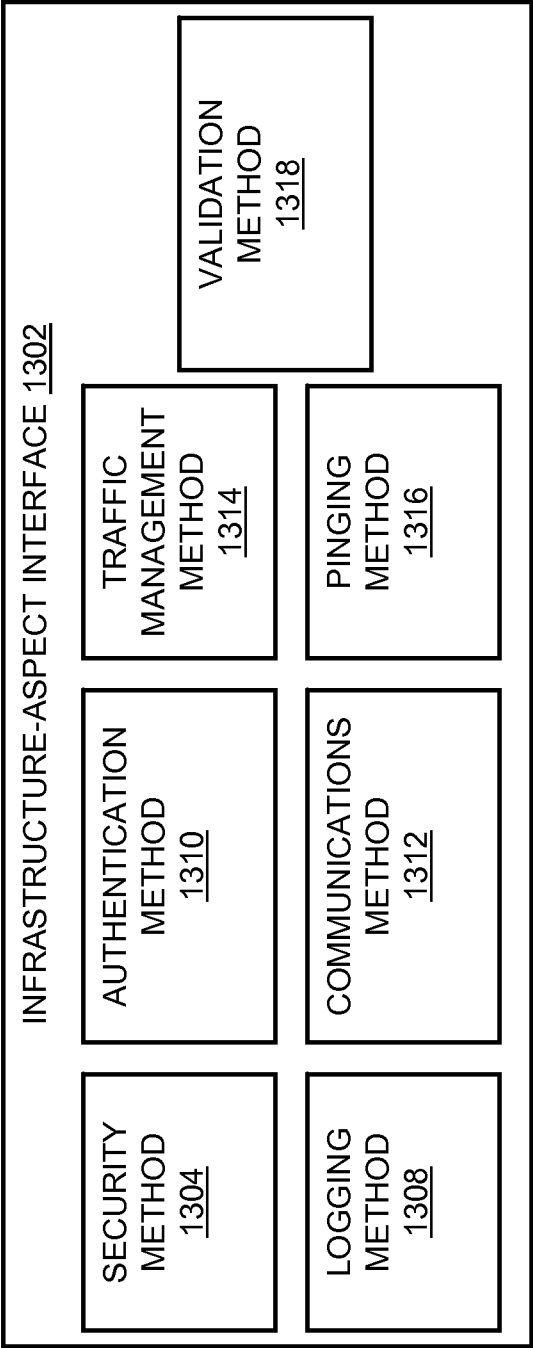


Fig. 10

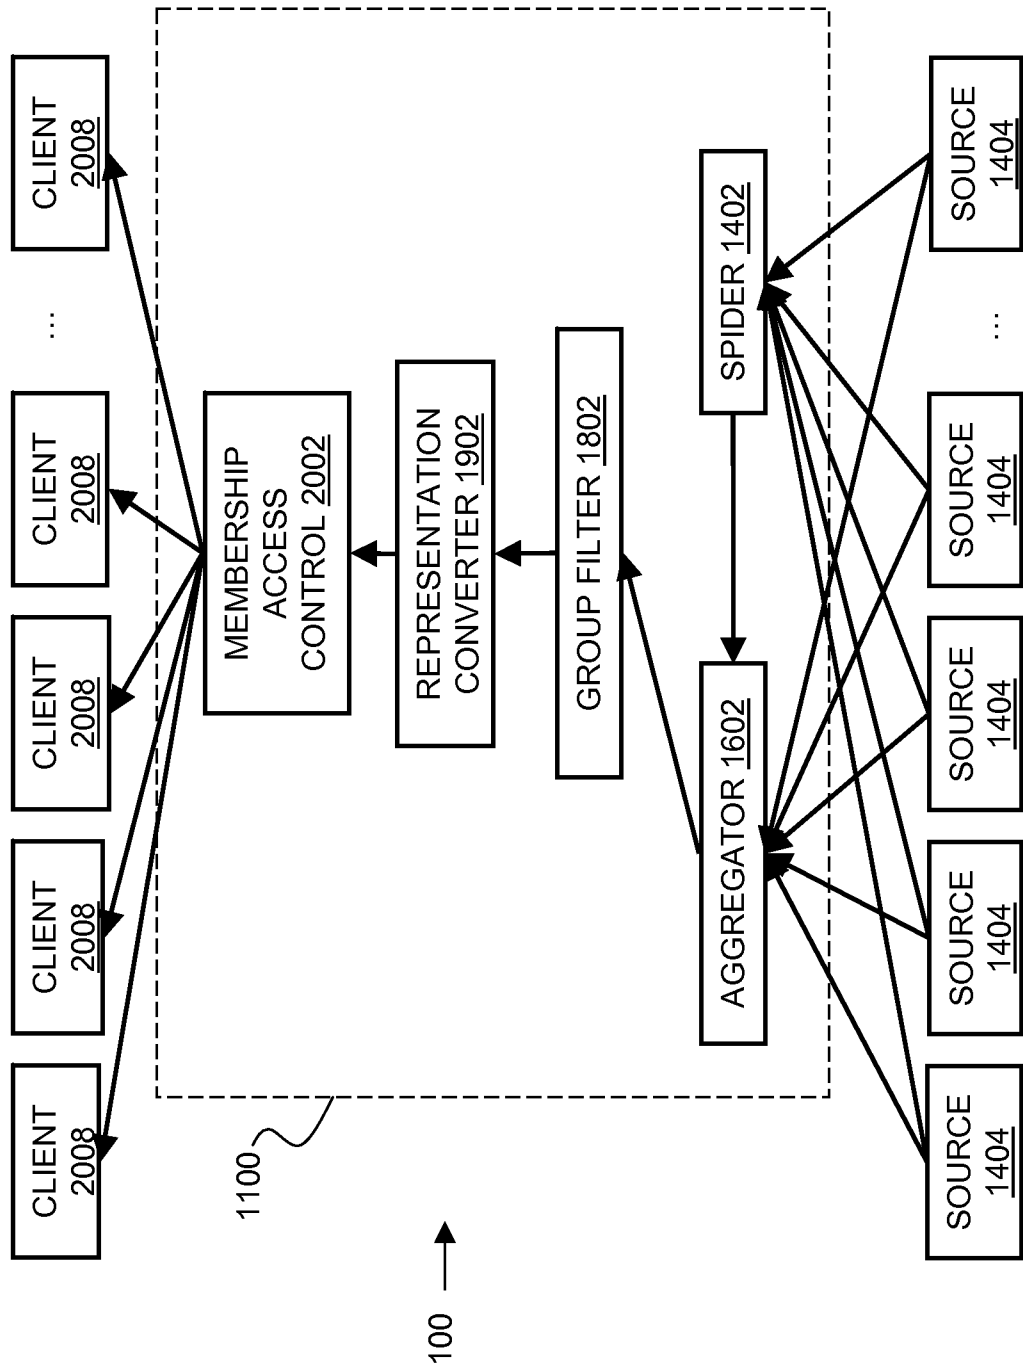


Fig. 11



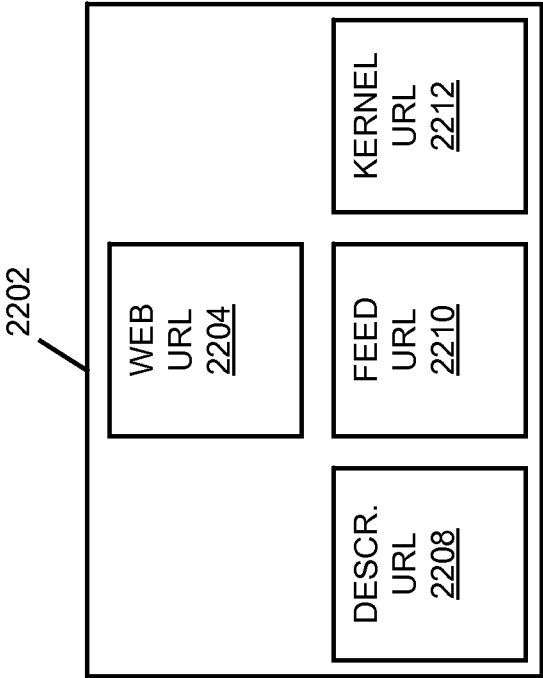


Fig. 12

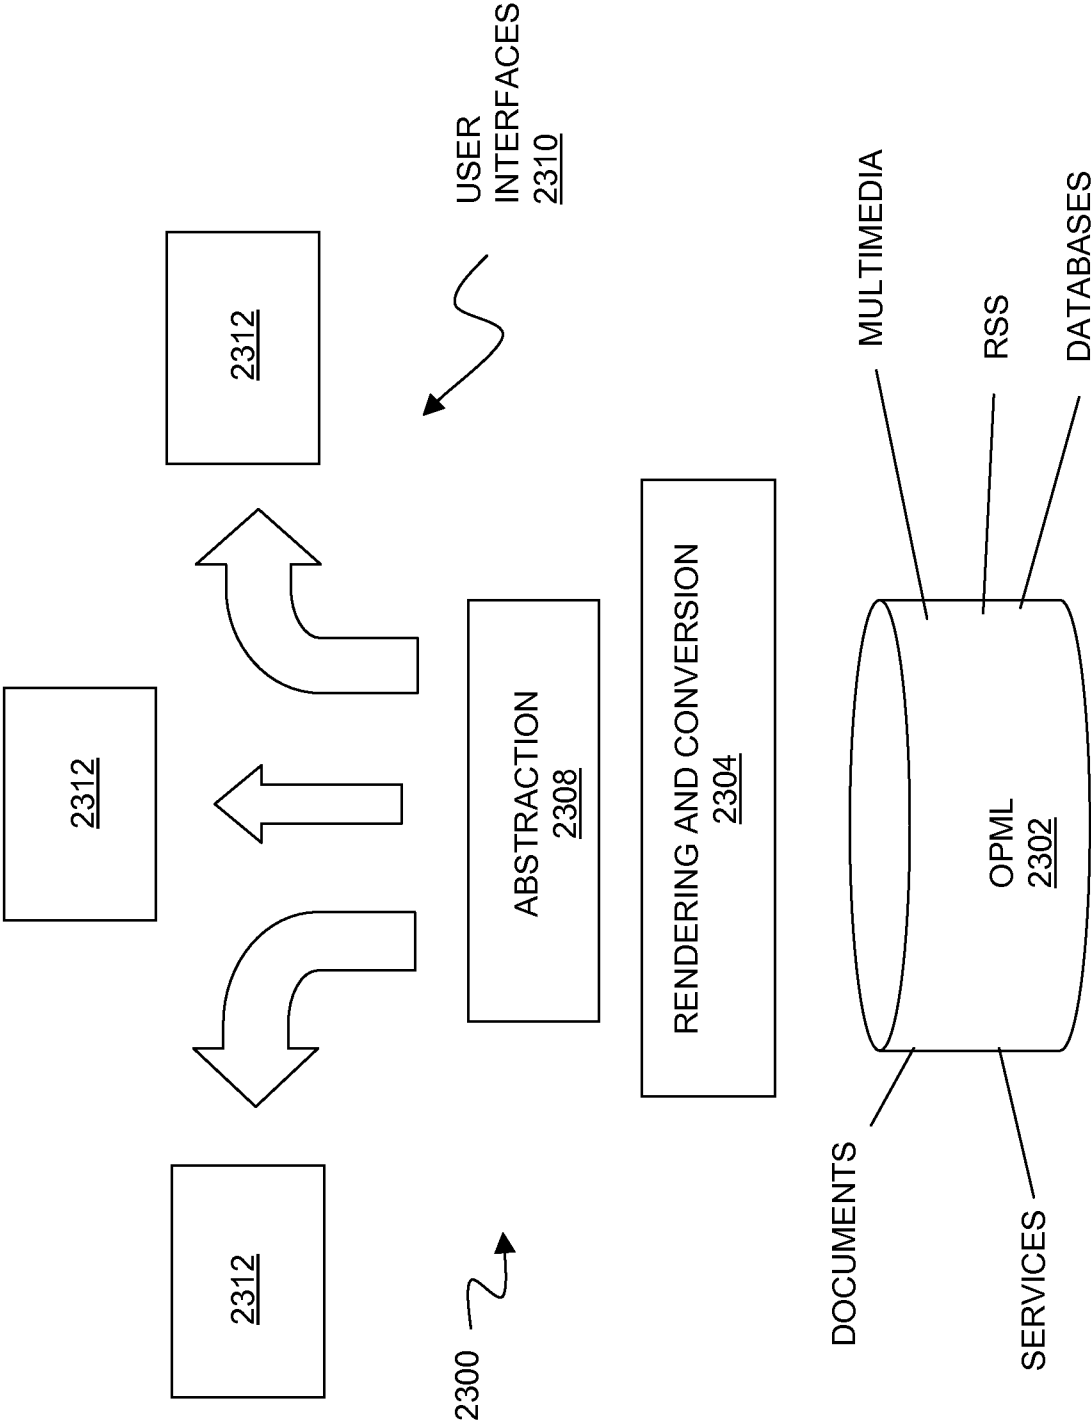


Fig. 13

[illegible]

Fig. 14

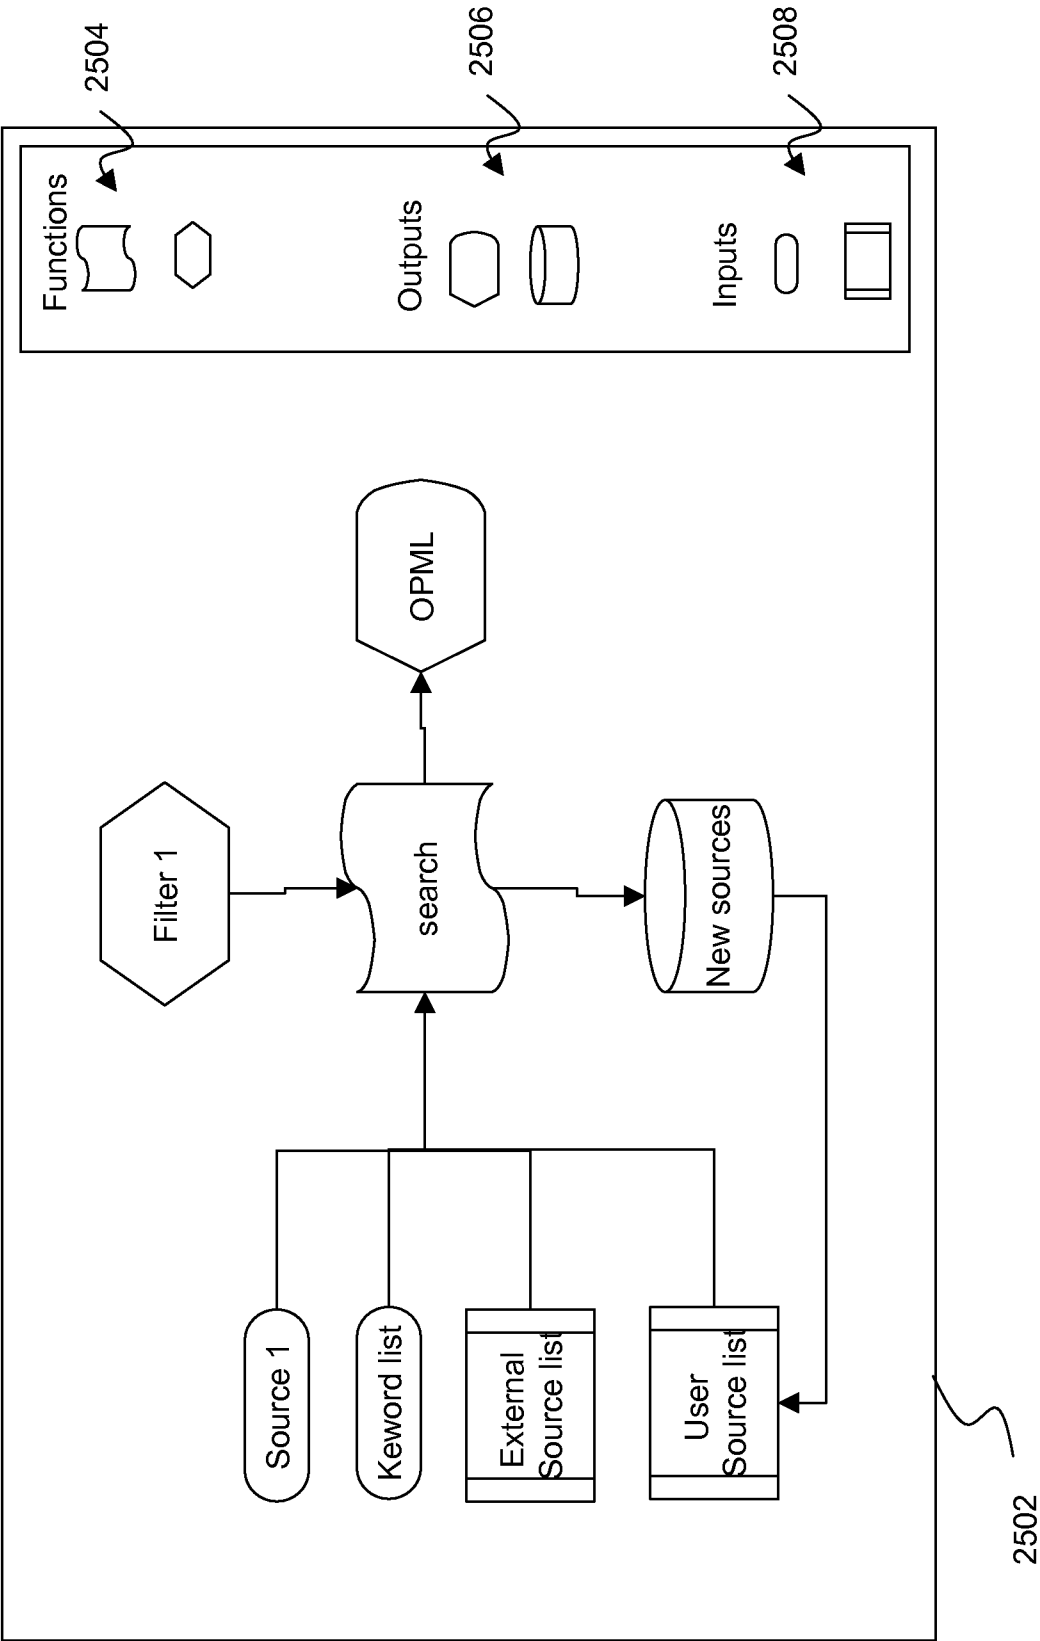


Fig. 15

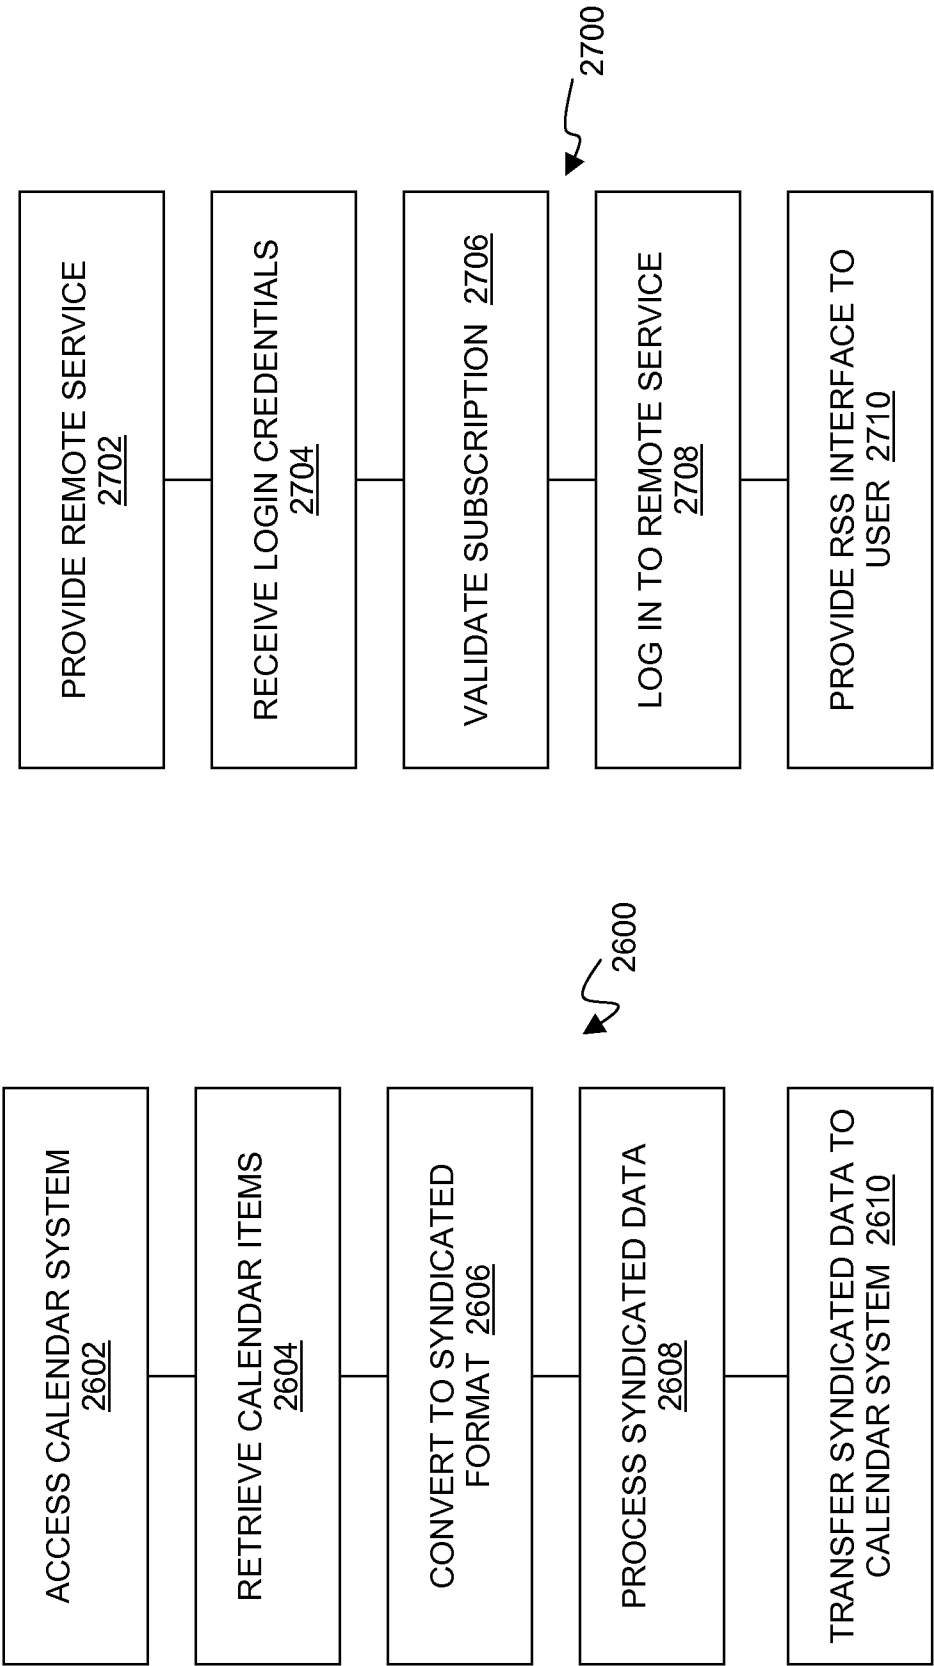


Fig. 16

Fig. 17

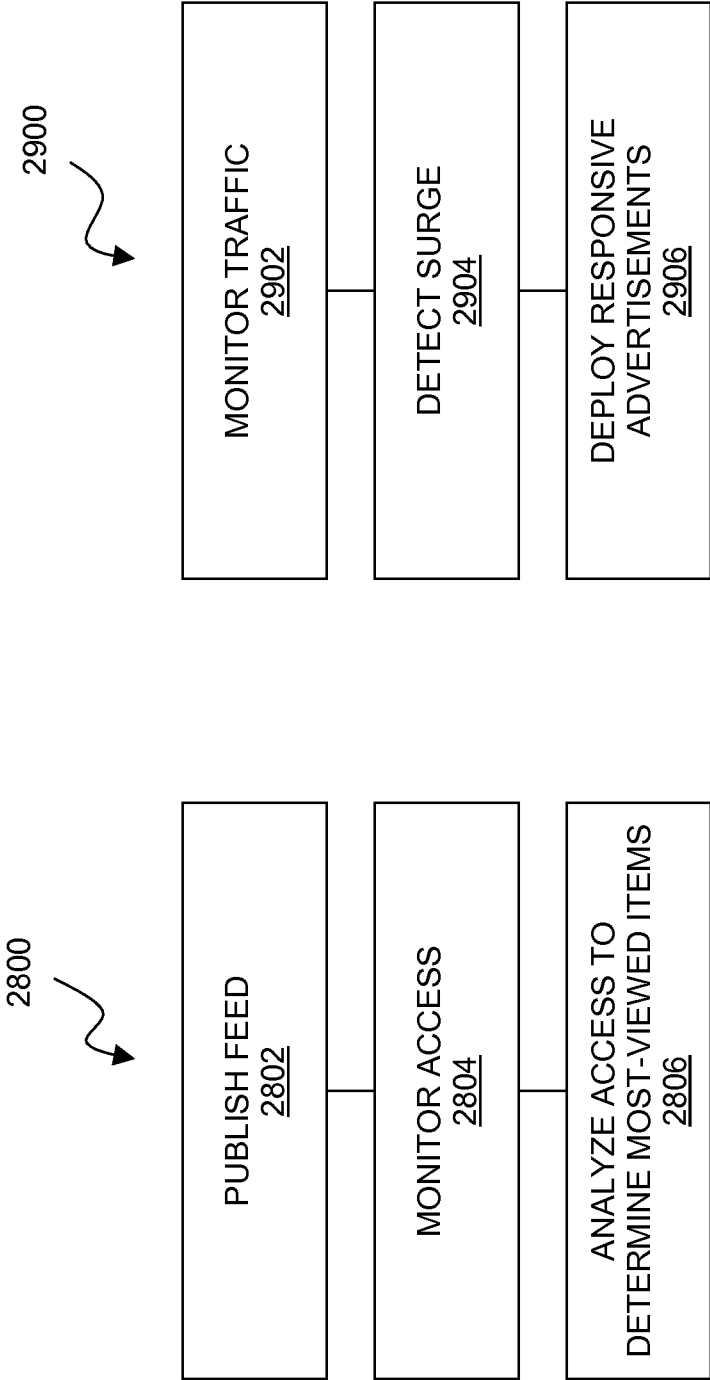


Fig. 18

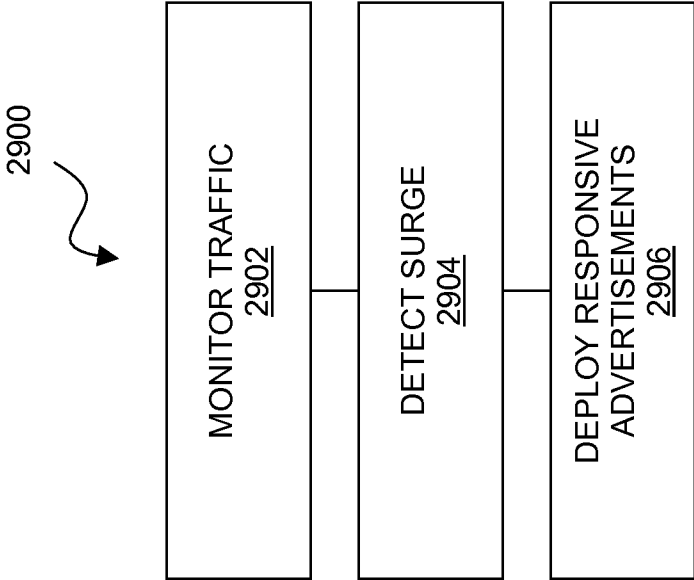


Fig. 19