

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
4 January 2007 (04.01.2007)

PCT

(10) International Publication Number
WO 2007/001864 A1

(51) International Patent Classification:
G06F 9/44 (2006.01)

(21) International Application Number:
PCT/US2006/023184

(22) International Filing Date: 14 June 2006 (14.06.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/158,936 21 June 2005 (21.06.2005) US

(71) Applicant (for all designated States except US): **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: **GANDHI, Amar, S.**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **PRAITIS, Edward, J.**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **KIM, Jane, T.**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **LYNDERSAY,**


Sean, O.; One Microsoft Way, Redmond, Washington 98052-6399 (US). **VON KOCH, Walter, V.**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **GOULD, William**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **MORGAN, Bruce, A.**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **KWAN, Cindy**; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

[Continued on next page]

(54) Title: CONTENT SYNDICATION PLATFORM



(57) Abstract: A content syndication platform, such as a web content syndication platform, manages, organizes and makes available for consumption content that is acquired from the Internet. In at least some embodiments, the platform can acquire and organize web content, and make such content available for consumption by many different types of applications. These applications may or may not necessarily understand the particular syndication format. An application program interface (API) exposes an object model which allows applications and users to easily accomplish many different tasks such as creating, reading, updating, deleting feeds and the like.

WO 2007/001864 A1



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, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, 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).

Published:

— with international search report

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

CONTENT SYNDICATION PLATFORM

BACKGROUND

RSS, which stands for *Really Simple Syndication*, is one type of web content syndication format. RSS web feeds have become more and more popular on the web and numerous software applications with RSS support are being developed. These numerous applications can have many varied features and can lead users to install several different RSS-enabled applications. Each RSS application will typically have its own list of subscriptions. When the list of subscriptions is small, it is fairly easy for a user to enter and manage those subscriptions across the different applications. As the list of subscriptions grows, however, management of the subscriptions in connection with each of these different RSS-enabled applications becomes very difficult. Thus, it is very easy for subscription lists to become unsynchronized.

In addition, web feeds come in several different file formats, with the popular ones being RSS 0.91, 0.92, 1.0, 2.0 and Atom. Each RSS-enabled application has to support most of these formats and possibly even more in the future. Implementing parsers for use in the RSS context for some applications is more difficult than for others. Given that not all application developers are RSS experts who possess experience and knowledge with regard to the intricacies of each format, it is unlikely that all application developers will implement the parsers correctly. Hence, it is likely given the rich number of file formats that some application developers will opt to not develop applications in this space or, if they do, the applications will not be configured to fully exploit all of the features that are available across the different file formats.

Another aspect of RSS and web feeds pertains to the publishing of content. For example, the number of users with blogs (weblogs) is increasing. There are many publicly available services that provide free blog services. Publishing content to a blog service, however, can be rather cumbersome since it might involve opening a browser, navigating to the blog service, signing in, and then typing the entry and submitting it.

1 Many application developers would prefer to be able to publish from within their
2 particular application, without breaking the user flow by having to go to a website. In
3 addition, there are many different types of protocols that can be used to communicate
4 between a client device and a particular service. Given this, it is unlikely that
5 application developers will implement all protocols. As such, the user experience will
6 not be all that it could be.

8 **SUMMARY**

9 A content syndication platform, such as a web content syndication platform,
10 manages, organizes and makes available for consumption content that is acquired from a
11 source, such as the Internet, an intranet, a private network or other computing device, to
12 name just a few. In some embodiments, the platform can acquire and organize web
13 content, and make such content available for consumption by many different types of
14 applications. These applications may or may not necessarily understand the particular
15 syndication format. An application program interface (API) exposes an object model
16 which allows applications and users to easily accomplish many different tasks such as
17 creating, reading, updating, deleting feeds and the like.

18 In addition, the platform can abstract away a particular feed format to provide a
19 common format which promotes the useability of feed data that comes into the platform.
20 Further, the platform processes and manages enclosures that might be received via a
21 web feed in a manner that can make the enclosures available for consumption to both
22 syndication-aware applications and applications that are not syndication-aware.

24 **BRIEF DESCRIPTION OF THE DRAWINGS**

25 Fig. 1 is a high level block diagram that illustrates a system that includes a web
content syndication platform in accordance with one embodiment.

Fig. 2 is a block diagram illustrates aspects of an object model in accordance with one embodiment.

Fig. 3 is a block diagram that illustrates a feed synchronization engine in accordance with one embodiment.

Fig. 4 illustrates an exemplary feed store in accordance with one embodiment.

Fig. 5 illustrates an exemplary user's profile in accordance with one embodiment.

Fig. 6 illustrates exemplary objects in accordance with one embodiment.

Fig. 7 illustrates exemplary objects in accordance with one embodiment.

DETAILED DESCRIPTION

Overview

A content syndication platform, such as a web content syndication platform, is described which is utilized to manage, organize and make available for consumption content that is acquired from a source, such as the Internet, an intranet, a private network or other computing device, to name just a few. In the context of this document, the platform is described in the context of an RSS platform that is designed to be used in the context of RSS web feeds. It is to be appreciated and understood that the RSS context constitutes but one example and is not intended to limit application of the claimed subject matter to only RSS contexts. The description below assumes some familiarity on the part of the reader with RSS. For background on RSS, there are a number of publicly available specifications that provide information that may be of interest to the reader.

In this document, certain terminology will be used in the context of the RSS embodiment that is described. An *item* is a basic unit of a feed. Typically, an item represents a blog entry or a news article/abstract, with a link to the actual article on the website. An *enclosure* is similar to an email attachment, except that there is a link to actual content. A *feed* is a list of items in a resource, usually only the most recent

1 additions. A *system feed list* is a list of feeds to which a user is subscribed. A
2 *subscription* refers to the act of signing up to receive notifications of new feed items.

3 In the various embodiments described in this document, the platform can acquire
4 and organize web content, and make such content available for consumption by many
5 different types of applications. These applications may or may not necessarily
6 understand the particular syndication format. Thus, in the implementation example,
7 applications that do not understand the RSS format can nonetheless, through the
8 platform, acquire and consume content, such as enclosures, acquired by the platform
9 through an RSS feed.

10 The platform comprises an application program interface (API) that exposes an
11 object model which allows applications and users to easily accomplish many different
12 tasks such as creating, reading, updating, deleting feeds and the like. For example,
13 using the API, many different types of applications can access, manage and consume
14 feedlists which includes a list of feeds.

15 In at least one embodiment, the platform provides multiple different feed parsers
16 each of which can parse a particular format in which a web feed may be received. The
17 parsed format is then converted into a *common format* which can then be leveraged by
18 applications and users. The common format is utilized to abstract away specific notions
19 embodied by any one particular format in favor of a more universal, understandable
20 format.

21 Further, the platform processes and manages enclosures that might be received
22 via a web feed in a manner that can make the enclosures available for consumption to
23 both syndication-aware applications and applications that are not syndication-aware. In
24 at least some embodiments, the APIs allow for discovery of the relationship between an
25 enclosure and its associated feed item.

In the discussion that follows, an exemplary platform and its components are first
described under the heading "Web Content Syndication Platform". Following this

1 discussion, an implementation example (under the heading "Implementation Example")
2 is provided and describes a set of APIs that expose an object model that enables
3 applications and users to interact with the platform in a meaningful and robust way.

4 5 **Web Content Syndication Platform**

6 Fig. 1 shows an exemplary system in accordance with one embodiment, generally
7 at 100. Aspects of system 100 can be implemented in connection with any suitable
8 hardware, software, firmware or combination thereof. In at least one embodiment,
9 aspects of the system are implemented as computer-readable instructions that reside on
10 some type of computer-readable medium.

11 In this example, system 100 comprises a content syndication platform 102 and a
12 collection of applications 104 individual ones of which can be configured to utilize the
13 platform in different ways, as will become apparent below. In at least some
14 embodiments, the content syndication platform comprises a web content syndication
15 platform. In the discussion that follows, the platform 102 is described in the context of
16 an RSS platform. It is to be appreciated and understood that this is intended as but an
17 example and is not intended to limit application of the claimed subject matter to only
18 RSS environments. Rather, principles of the described embodiments can be utilized in
19 other syndication environments without departing from the spirit and scope of the
20 claimed subject matter.

21 In this example, platform 102 comprises an object model 106 that is exposed by a
22 set of APIs that enable applications 104 to interact with the platform. A synchronization
23 engine 108 is provided and is configured to, among other things, acquire web content
24 and, in at least some embodiments, convert the web content into a so-called *common*
25 *format*, which is described in more detail below.

A publishing engine 110 permits users to publish content, such as blogs, in a
manner that abstracts away, via the APIs, the communication protocol that is utilized to

1 communicate between the user's application or computing device and the server or
2 destination software that is to receive the content.

3 In addition, in at least one embodiment, platform 102 includes a feed store 112
4 that stores both feed lists 114 and feed data 116. Further, platform 102 utilizes, in at
5 least one embodiment, file system 118 to store and maintain enclosures 120. Using the
6 file system carries with it advantages among which include enabling applications that do
7 not necessarily understand the syndication format to nonetheless consume enclosures
8 that may be of interest. Further, platform 102 includes a post queue 122 that holds post
9 data 124 that is to be posted to a particular web-accessible location.

10 As noted above, platform 102 can enable applications to access, consume and
11 publish web content. Accordingly, the collection of applications 104 can include many
12 different types of applications. In at least some embodiments, the types of applications
13 can include those that are syndication-aware and those that are not syndication-aware.
14 By "syndication-aware" is meant that the application is at least somewhat familiar with
15 the syndication format that is utilized. Thus, in the RSS context, a syndication-aware
16 application is one that may be configured to process data or otherwise interact with
17 content that is represented in an RSS format. This can include having the ability to
18 parse and meaningfully interact with RSS-formatted data. Similarly, an application that
19 is not syndication-aware is typically not configured to understand the syndication
20 format. Yet, through the platform, as will become apparent below, applications that are
21 not syndication aware can still access and consume content that arrives at the platform in
22 a syndication format.

23 Looking more specifically at the different types of applications that can interact
24 with the platform, collection 104 includes a web browser application 122, an RSS reader
25 application 124, a digital image library application 126, a media player application 128
and a blog service 130. In this example, RSS reader application 124 is a syndication-
aware application, while media player 128 may not necessarily be a syndication-aware

1 application. Further, web browser application 122 may or may not be a syndication-
2 aware application. Of course, these applications constitute but examples of the different
3 types of applications that can interact with the platform. As such, other types of
4 applications that are the same or different from those illustrated can be utilized without
5 departing from the spirit and scope of the claimed subject matter. By way of example
6 and not limitation, these other types of applications can include calendar applications for
7 event feeds, social networking and email applications for contact feeds, screen saver
8 applications for picture feeds, CRM for document feeds, and the like.

9 In the discussion that follows, aspects of the individual components of the
10 platform 102 are described in more detail, each under its own heading.

11 12 Object Model

13 Fig. 2 illustrates individual objects of object model 106 in accordance with one
14 embodiment. The object model about to be described constitutes but one example of an
15 object model that can be utilized and is not intended to limit application of the claimed
16 subject matter to only the object model that is described below. As noted above, the
17 object model is exposed by an API, an example of which is described below.

18 In this particular object model, a top level object 200 called *feeds* is provided.
19 The feeds object 200 has a property called *subscriptions* of the type folder. Subscription
20 or folder objects 202 are modeled as a hierarchy of folders. Thus, in this particular
21 example, subscription or folder objects have properties that include *subfolders* 204 of
22 the type folder and *feeds* 206 of the type feed. Underneath the feeds object 206 is an
23 item object 208 of the type item, and underneath the item object 206 is an enclosure
24 object 210 of the type object.

25 The individual objects of the object model have properties, methods and, in some
instances, events that can be utilized to manage web content that is received by the
platform. The above-described object model permits a hierarchical structure to be

1 utilized to do such things as manage feedlists and the like. For example, using a folder
2 structure, the platform can execute against a set of feeds. As will be appreciated by the
3 skilled artisan, this makes it easier for the application developer. For example,
4 executing against a set of feeds provides the ability to refresh all of the "news" feeds,
5 located within the news folder.

6 As an example, consider the following. Assume that a user wishes to interact
7 with or consume data associated with a feed to which they are not actually subscribed.
8 For feeds that are subscribed to, i.e. those that are represented inside the root level
9 subscription folder, the synchronization engine 108 (Fig. 1) will pick up the feed and
10 start to, on an appropriate interval, fetch data associated with the feed. There are cases,
11 however, when an application that uses the platform does not wish to be subscribed to a
12 particular feed. Rather, the application just wants to use the functionality of the
13 platform to access data from a feed. In this case, in this particular embodiment,
14 subscriptions object 202 supports a method that allows a feed to be downloaded without
15 subscribing to the feed. In this particular example, the application calls the method and
16 provides it with a URL associated with the feed. The platform then utilizes the URL to
17 fetch the data of interest to the application. In this manner, the application can acquire
18 data associated with a feed in an adhoc fashion without ever having to subscribe to the
19 feed.

20 Considering the object model further, consider item and enclosure objects 208,
21 210 respectively. Here, these objects very much reflect how RSS is structured itself.
22 That is, each RSS feed has individual items inside of which can optionally appear an
23 enclosure. Thus, the structure of the object model is configured to reflect the structure
24 of the syndication format.

25 From an object model perspective, there are basically two different types of
methods and properties on an item. A first type of method/property pertains to data

1 which is read only, and a second type of method/property pertains to data which can be
2 both read and written.

3 As an example of the first type of method property, consider the following. Each
4 feed can have data associated with it that is represented in an XML structure. This data
5 includes such things as the title, author, language and the like. Data such as this is
6 treated by the object model as read only. For example, the data that is received by a
7 feed and associated with individual items is typically treated as read only. This prevents
8 applications from manipulating this data. Using an XML structure to represent the feed
9 data also carries with it advantages as follows. Assume that the synchronization engine
10 does not understand a new XML element that has been added. Nonetheless, the
11 synchronization engine can still store the element and its associated data as part of the
12 feed item data. For those applications that do understand the element, this element and
13 its associated data are still available for the application to discover and consume.

14 On the other hand, there is data that is treated as read/write data, such as the name
15 of a particular feed. That is, the user may wish to personalize a particular feed for their
16 particular user interface. In this case, the object model has properties that are read/write.
17 For example, a user may wish to change the name of a feed from "New York Times" to
18 "NYT". In this situation, the name property may be readable and writable.

20 Feed Synchronization Engine

21 In the illustrated and described embodiment, feed synchronization engine 108
22 (Fig. 1) is responsible for downloading RSS feeds from a source. A source can
23 comprise any suitable source for a feed, such as a web site, a feed publishing site and the
24 like. In at least one embodiment, any suitable valid URL or resource identifier can
25 comprise the source of a feed. The synchronization engine receives feeds and processes
the various feed formats, takes care of scheduling, handles content and enclosure
downloads, as well as organizes archiving activities.

Fig. 3 shows an exemplary feed synchronization engine 108 in a little more detail in accordance with one embodiment. In this embodiment, synchronization engine includes a feed format module 300, a feed schedule module 302, a feed content download module 304, an enclosure download module 306 and an archiving module 308. It is to be appreciated and understood that these module are shown as logically separate modules for purposes of clearly describing their particular functionalities. The logically separate modules are not intended to limit the claimed subject matter to only the particular structures or architectures described herein.

Feed Format Module -- 300

In the illustrated and described embodiment, feeds are capable of being received in a number of different feed formats. By way of example and not limitation, these feed formats can include RSS 1.0, 1.1, .9x, 2.0, Atom .3, and so on. The synchronization engine, via the feed format module, receives these feeds in the various formats, parses the format and transforms the format into a normalized format referred to as the *common format*. The common format is essentially a superset of all supported formats. One of the benefits of using a common format is that applications that are format-aware now need to only be aware of one format—the common format. In addition, managing content that has been converted into the common format is much easier as the platform need only be concerned with one format, rather than several. Further, as additional syndication formats are developed in the future, the feed format module can be adapted to handle the format, while at the same time permit applications that are completely unaware of the new format to nonetheless leverage and use content that arrives at the platform via the new format.

With regard to the common format, consider the following. From a format standpoint, the common format is represented by an XML schema that is common between the different formats. In a different format, certain elements may have different

names, different locations within the hierarchy of the XML format and the like. Accordingly, the common format is directed to presenting a common structure and syntax that is derived collectively from all of the different formats that are possible. Thus, in some instances, elements from one format may be mapped into elements of the common format.

Feed Schedule Module -- 302

Each feed can have its own schedule of when the synchronization engine 108 should check to ascertain whether there is new content available. Accordingly, the synchronization engine, through the feed schedule module 302, manages such schedules to respect a site's as well as a user's or a system's requirements and limitations.

As an example, consider the following. When a feed is first downloaded, an update schedule (i.e. a schedule of when the feed is updated) may be included in the feed's header. In this case, the feed schedule module 302 maintains the update schedule for this particular feed and checks for new content in accordance with the update schedule. If, however, no schedule information is included, then the feed schedule module can utilize a *default schedule* to check for new content. Any suitable default schedule can be used such as, for example, re-downloading the feed content every 24 hours. In at least some embodiments, the user may specify a different default work schedule.

In addition, in at least some embodiments, the feed schedule module can support what is referred to as a *minimum schedule*. The *minimum schedule* refers to a minimum update time that defines a period of time between updates. That is, the platform will not update a feed more often than what the minimum schedule defines. In at least some embodiments, the user can change the minimum time. In addition, the user can also initiate a manual refresh of any, or all feeds.

1 In addition to supporting default and minimum schedules, in at least some
2 embodiments, the feed schedule module can support *publisher-specified schedules*. As
3 the name implies, a publisher-specified schedule is a schedule that is specified by a
4 particular publisher. For example, the publisher-specified schedule can typically specify
5 how many minutes until the client should next update the feed. This can be specified
6 using the RSS 0.9x/2.0 "ttl" element. The synchronization engine should not fetch a
7 new copy of the feed until at least that number of minutes has passed. The publisher-
8 specified schedule can also be specified at different levels of granularity such as hourly,
9 daily, weekly, etc.

10 It should be noted that each copy of a feed document can have a different
11 publisher-specified schedule. For example, during the day, the publisher may provide a
12 schedule of 15 minutes, and then during the night, the publisher may provide a schedule
13 of 1 hour. In this case, the synchronization engine updates its behavior every time the
14 feed is downloaded.

15 In addition, in at least some embodiments, the synchronization engine, via the
16 feed schedule module 302, supports the notion of skipping hours and/or days.
17 Specifically, RSS 0.9 and 2.0 enable a server to block out certain days and hours during
18 which the client should not conduct an update. In this case, the synchronization engine
19 respects these settings, if provided by the server, and does not update the feed during
20 those times.

21 In addition to the default, minimum and publisher-specified schedules, in at least
22 some embodiments, the synchronization engine supports the notion of user-specified
23 schedules and manual updates. More specifically, on a per-feed basis, the user can
24 specify a schedule of their choice. From a platform perspective, the user-specified
25 schedule can be as complex as specified by a server. In this instance, the platform, via
the feed schedule module, maintains the most recent schedule extracted from the feed as
well as the user schedule. In at least some embodiments, the user schedule always

1 overrides the publisher's schedule. In addition, at any time, an application can initiate a
2 forced update of all feeds or individual feeds.

3 With regard to bandwidth and server considerations, consider the following. In
4 accordance with one embodiment, the synchronization engine can be designed in view
5 of two related issues. First, the synchronization should be considerate of the user's
6 bandwidth and CPU. Second, because of widespread use of the RSS platform, the
7 synchronization engine should be considerate of its impact on servers. These two issues
8 have an impact on both *when* and *how* feeds are downloaded.

9 From the perspective of *when* a feed is downloaded, synchronization engine can
10 be designed with the following considerations in mind. In the absence of a schedule
11 from the server, and any other instructions from the user, the synchronization engine
12 should be very conservative in how often it updates. Hence, in at least some
13 embodiments, the default schedule is set to 24 hours. Further, to protect the user's
14 resources from being adversely impacted by an inefficient server, a minimum schedule
15 can be enforced to keep the synchronization engine from updating too often, even if the
16 server specifies otherwise. In addition, updates at login time (and at common intervals,
17 e.g. each hour from the startup time) should be carefully managed. Feed updates should
18 be delayed until a specified period of time after user login has completed, and should be
19 staggered slightly to avoid large update hits each hour, on the hour. This can be
20 balanced against a user's desire to have all of the updates happen at once. Further, when
21 a server uses the skip hours or skip days feature described above, the client should not
22 immediately fetch an update as soon as the moratorium period is over. Instead, the client
23 should wait a random interval ranging up to 15 minutes before fetching the content.

24 To assist the synchronization engine in this regard, the feed schedule module 302
25 can maintain a state for each feed, such as fresh or stale. A "fresh" state means that,
based on the publisher schedule, the feed is fresh. A "stale" state means that the
publisher's schedule has indicated an update, but the synchronization engine has not yet

1 completed the update. Clients with an interest in the freshest content can request an
2 immediate update, and be notified when it is available. If this expectation is set, then
3 the synchronization engine can implement arbitrary delays in updating the content,
4 rather than rigorously following the schedule to the detriment of the user and the server.

5 With regard to *how* a feed is downloaded, consider the following. In one
6 embodiment, the synchronization engine can use a task scheduler to launch a
7 synchronization engine process at a pre-defined time. After the synchronization engine
8 has completed, it updates a task schedule with the next time it should launch the
9 synchronization engine again (i.e., NextSyncEngineLaunchTime).

10 When the synchronization engine launches, it queues up all "pending" feeds
11 whose NextUpdateTime is less or equal to the currentTime and then processes them as
12 follows. For each feed, the following properties are tracked: LastUpdateTime,
13 NextUpdateTime, Interval (specified in minutes) and LastErrorInterval.

14 At the end of successfully synching a feed, the feed's LastUpdateTime is set to
15 the current time and NextUpdateTime is set to LastUpdateTime plus an interval plus
16 randomness (1/10th of the interval). Specifically:

```
17 LastUpdateTime = currentTime  
18 NextUpdateTime = currentTime + Interval + Random(Interval * 0.1)  
19 ErrorInterval = 0
```

20 Random(argument) is defined to be a positive value between 0 and its argument.
21 For example Random(10) returns a float between 0..10.

22
23 If synching of a feed failed for one of the following reasons:

```
24  
25 HTTP 4xx response code;  
HTTP 5xx response code;  
Winsock/network error; or  
HTTP 200, but response body has a parsing error (not a recognized feed  
format)
```


then an exponential back off algorithm is applied as follows:

```
LastUpdateTime = <unchanged>
ErrorInterval = min( max(ErrorInterval * 2 , lmin), Interval)
NextUpdateTime = currentTime + ErrorInterval + Random(ErrorInterval *
0.1)
```

After synchronization of all “pending” feeds has completed, the synchronization engine determines if there are any feeds whose NextUpdateTime has passed (NextUpdateTime ≤ currentTime). If there are, then those “pending” feeds are queued and processed as if the synchronization engine just launched.

If there are no outstanding “pending” feeds, then the synchronization engine determines if there are any “soon-to-sync” feeds whose NextUpdateTime is within two minutes of the current time (currentTime + 2min ≥ NextUpdateTime). If there are any “soon-to-sync” feeds then the synchronization engine process continues to run, and it sets a timer to “wake up” at NextUpdateTime and process “pending” feeds.

If there are no “soon-to-sync” feeds then the NextSyncEngineLaunch is set to the NextUpdateTime of the feed with the soonest NextUpdateTime. Then the task scheduler is set to NextSyncEngineLaunchTime and the synchronization engine process ends.

In accordance with one embodiment, if there are several “pending” feeds in the queue, the synchronization engine can synchronize multiple feeds in parallel. However, the number of parallel synchronizations should be limited, as well as how many synchronizations are performed in a certain time period in order to not saturate network bandwidth and processor utilization. In accordance with one embodiment, feed synchronization shaping is provided via a token-bucket. Conceptually, the token bucket works as follows.

- A token is added to the bucket every $1/r$ seconds;
- The bucket can hold at most b tokens; if a token arrives when the bucket is full, it is discarded;
- When a feed needs to be synchronized, a token is removed from the bucket and the feed is synchronized;
- If no tokens are available, the feed stays in the queue and waits until a token becomes available.

This approach allows for bursts of feed synchronizations of up to b feeds. Over the long run, however, the synchronizations are limited to a constant rate r . In an implementation example, the synchronization engine uses the following values for b and r : $b = 4$ and $r = 2$.

Feed Content Download Module -- 304

In accordance with one embodiment, feed content download module 304 handles the process of downloading a feed and merging the new feed items with the existing feed data.

As an example of how one can implement a feed content download module, consider the following. At the appropriate time, the synchronization engine, via the feed content download module, connects to a server and downloads the appropriate content.

In accordance with one embodiment, the platform is configured to support different protocols for downloading content. For example, the synchronization engine can support downloading the feed document over HTTP. In addition, the synchronization engine can support encrypted HTTP URLs (e.g., SSL, https and the like). Likewise, the synchronization engine can also support compression using the HTTP gzip support, as well as support feed downloads from Universal Naming Convention (UNC) shares.

In addition, the synchronization engine via the feed content download module can support various types of authentication. For example, the synchronization engine

1 can store a username/password for each feed, and can use this username/password for
2 HTTP Basic authentication to retrieve the feed document.

3 With regard to updating a feed, consider the following. To determine if a feed
4 has new content, the synchronization engine keeps the following pieces of information,
5 for each feed:

- 6 • The last time the feed was updated as reported by the Last-modified
7 header on the HTTP response;
- 8 • The value of the Etag header in the last HTTP response; and
- 9 • The most recent pubDate value for the feed (i.e. the feed-level publication
10 date and time).

11 If the site supports Etag or Last-modified, then the synchronization engine can
12 use these to check if there is new content. The site can respond with an HTTP response
13 code 304 to indicate that there is no new content. Otherwise, the content is downloaded.
14 For example, if the site supports RFC 3229-for-feeds, the site can return only the new
15 content, based on the Etag passed by the client. Either way, the client then merges the
16 new content with the stored content.

17 As a more detailed description of how feed content can be downloaded in but one
18 implementation example, consider the following. To determine if a particular site has
19 changed, the synchronization engine will submit a request with:

- 20 • The *If-None-Match* header, if the client has a saved Etag;
 - 21 ○ The header A-IM with the values: feed, gzip (used for RFC 3229-
22 for-feeds);
- 23 • The *If-Modified-Since* header, if the client has a saved Last-modified
24 value.

25 If the server responds with an HTTP Response code 304, then the content has not
changed and the process may end here. If the server responds with content (i.e. HTTP
codes 200 or 206), then the downloaded content is merged with the local content (note:

code 206 means that the server supports RFC3229-for-feeds, and the content downloaded is only the new content).

If there is content available and if the synchronization engine has a pubDate stored, and the downloaded feed document contains a channel-level pubDate element, the two dates are compared. If the local pubDate is the same as the downloaded pubDate, then the content has not been updated. The downloaded feed document can then be discarded.

If the synchronization engine processes each item one at a time, each item's pubDate is compared against the pubDate that the synchronization engine has stored (if any) and older items are discarded. Each item is then compared against the items in the store. The comparison should use the guid element, if present, or the link element, if guid is not present. If a match is found, then the content of the new item replaces that of the old item (if both have a pubDate, then it is used to determine which is newer, otherwise, the most recently downloaded is new). If no match is found, then the new item is pre-pended to the stored feed content (maintaining a "most recent at the top" semantic). If any item is added or updated in the local feed, the feed is considered updated, and clients of the RSS platform are notified.

For error cases, consider the following. If the server responds with a code 500 or most 400 errors, the synchronization schedule is reset and the server tries again later. The HTTP error 410, however, should be treated as an indication to reset the update schedule to "no more updates."

HTTP-level redirects should be followed, but no changes should be made to the client configuration (there are several pathological scenarios where redirects are given accidentally).

If the server responds with an XML redirect, then the feed should be redirected, and the stored URL to the feed should be automatically updated. This is the only case where the client updates the feed URL automatically.

1 With regard to downloading the feed, the download should not interrupt ordinary
2 usage of the machine (e.g., bandwidth or CPU) when the user is engaged in other tasks.
3 In addition, the user should be able to get the content as fast as possible when in an
4 interactive application that relies on the content.

6 Enclosure Download Module -- 306

7 In accordance with one embodiment, enclosure download module 306 is
8 responsible for downloading enclosure files for a feed and applying the appropriate
9 security zone. At the time of downloading the feed content, the enclosures are
10 downloaded as well.

11 Downloading enclosures can be handled in a couple of different ways. First, a
12 *basic enclosure* is considered to be an RSS 2.0-style enclosure. For basic enclosures,
13 the synchronization engine, via the enclosure download module 306, will automatically
14 parse the downloaded feeds for enclosure links. The synchronization engine is
15 configured to support multiple basic enclosures. Using the enclosure link, the enclosure
16 download module can then download the enclosure. In at least some embodiments, for
17 any new feed, the default action is not to download basic enclosures. Using the API
18 which exposes the above-described object model, client can do such things as change
19 the behavior on a per-feed basis to, for example, always download enclosures or force
20 the download of a specific enclosure of a specific item in a specific feed.

21 Enhanced enclosure handling can be provided through the use of the common
22 format described above. Specifically, in at least one embodiment, the common format
23 defines additional functionality for enclosures. Specifically, the common format enables
24 multiple representations of a particular piece of content. This includes, for example,
25 including standard definitions of preview content and default content, as well as the
ability to indicate whether an enclosure should be downloaded or streamed. In addition,
the common format permits arbitrary metadata on an enclosure, and on representations

of the content. For any new feed, the default action is to download the “preview” version of any enclosure, subject to a default size limit of, for example, 10k per item.

Using the API, clients can do such things as change the behavior on the per-feed basis. For example, the behavior can be changed to always download the “default” version of the items in a feed or to always download any specific version that has a metadata element of a particular value. This can be done, for example, with a client callback which provides the “download this?” logic for each enclosure. In addition, using the API, clients can force immediate download of any specific representation of any specific enclosure of any specific item (or all items) in a specific feed.

With regard to providing security in the enclosure download process, consider the following.

In accordance with one embodiment, downloaded enclosures use the Windows XP SP2 Attachment Execution Service (SP2 AES) functionality. This functionality can provide file-type and zone based security. For example, provided with a file name and zone information (i.e. where an enclosure came from), AES can indicate whether to block, allow or prompt.

With regard to zone persistence, when saving a file, AES can persist the zone information so that, when it is subsequently opened, the user can be prompted.

The table just below describes AES risk-level/zone to action mapping:

Risk Levels	Restricted	Internet	Intranet	Local	Trusted
Dangerous, e.g. EXE	Block	Prompt	Allow	Allow	Allow
Moderate/Unknown, e.g. DOC or FOO	Prompt	Prompt	Allow	Allow	Allow
Low, e.g. TXT or JPG	Allow	Allow	Allow	Allow	Allow

1 In the illustrated and described embodiment, the synchronization engine will call
2 a method, for example ::CheckPolicy, for each enclosure that it downloads. Based on
3 the response, the synchronization engine can do one of the following:

- 4 • Block: Don't save (mark it as failed in the feed file);
- 5 • Allow: Save the enclosure
- 6 • Prompt: Save, but persist zone information. This means that if the user
7 double-clicks on the file, they'll get a "Run/Don't Run" prompt.

8 In accordance with one embodiment, the synchronization engine will first save an
9 enclosure to disk and will not download the enclosure in memory. Saving to disk
10 triggers filter-based antivirus applications and gives these applications an opportunity to
11 quarantine the enclosure if they choose.

13 Archiving Module -- 308

14 In accordance with one embodiment, archiving module 308 is responsible for
15 dealing with old feed data. By default, a feed will hold a maximum of 200 items. When
16 a feed exceeds the specified maximum, the older feed items are deleted by the archiving
17 module. The associated enclosures are not, however, deleted.

19 Feed Store

20 In accordance with one embodiment, feed store 112 (Fig. 1) holds two types of
21 information – a feed list 114 and feed data 116. As an example, consider Fig. 4. There,
22 feed list 114 is embodied as a hierarchical tree structure 400 of the list of feeds. The
23 feed data 116 comprises the data associated with a particular feed. In this example, the
24 feed data 116 is arranged on a per-feed basis to include a collection 402 of items and
25 enclosures.

1 There are many different ways that one might implement a feed store. In this
2 particular embodiment, the feed store comprises part of the file system. One reason for
3 this pertains to simplicity. That is, in this embodiment, the feed list is represented
4 simply as a regular directory under which there can be sub-directories and files. The
5 hierarchy is reflected as a normal file system hierarchy. Thus, each folder such as
6 “News” and “Blogs” is essentially a regular directory in the file system with
7 subdirectories and files.

8 In this particular example, there is a special file type that represents a feed
9 subscription. By way of example only, consider that this type of file has the following
10 format: “xyz.stg”. The .stg file stores all of the data for a feed. Thus, you have a feed
11 list, such as the list embodied in tree structure 400, and inside each feed (or file) is the
12 feed data.

13 In the illustrated and described embodiment, the .stg files are implemented using
14 *structured storage* technology. Structure storage techniques are known and will be
15 appreciated by the skilled artisan. As brief background, however, consider the
16 following.

17 Structured storage provides file and data persistence in COM by handling a single
18 file as a structured collection of objects known as storages and streams. The purpose of
19 structured storage is to reduce the performance penalties and overhead associated with
20 storing separate object parts in different files. Structured storage provides a solution by
21 defining how to handle a single file entity as a structured collection of two types of
22 objects—storages and streams—through a standard implementation called *compound*
23 *files*. This enables the user to interact with, and manage, a compound file as if it were a
24 single file rather than a nested hierarchy of separate objects. The storage objects and
25 stream objects function as a file system within a file, as will be appreciated by the
skilled artisan. Structured storage solves performance problems by eliminating the need
to totally rewrite a file to storage whenever a new object is added to a compound file, or

1 an existing object increases in size. The new data is written to the next available location
2 in permanent storage, and the storage object updates the table of pointers it maintains to
3 track the locations of its storage objects and stream objects.

4 Thus, in the illustrated and described embodiment, the .stg files are implemented
5 using structured storage techniques and an API on top of the feed store allows access to
6 the different streams and storages. In this particular example, each RSS item is written
7 into one stream. Additionally, a header stream contains information associated with a
8 particular feed such as the title, subscription, feed URL and the like. Further, another
9 stream stores index-type metadata that allows quick and efficient access to contents in
10 the file for purposes that include quickly marking something as read/unread, deleting an
11 item and the like.

12 13 File System – Enclosures

14 In the illustrated and described embodiment, enclosures are not stored in
15 structured storage or as part of the feed data, as indicated in Fig. 1. Rather, enclosures
16 are recognized as being items, such as a picture or pictures, that other applications and
17 the user may want to access and manipulate.

18 Thus, in the illustrated and described embodiment, enclosures are written into a
19 user's particular profile. A link, however, is maintained between the enclosure and the
20 associated feed item.

21 As an example, consider Fig. 5. Once a user starts subscribing to a feed, the feed
22 content is stored locally under the user's profile, either in Application Data or in a
23 Knownfolder "feeds".

24 The feedlist and feeds are stored in Application Data to better be able to control
25 the format of the feedlist and the feeds. APIs are exposed (as will be described below)
such that applications can access and manage the feeds.

1 The feedlist is the set of feeds that the user is subscribed to. In this example, the
2 file that comprises the Feedlist is located at:

3
4 C:\Users\<Username>\AppData\Roaming\Microsoft\RSS\
5

6 The file contains the feed's properties, as well as items and enclosure properties
7 (a URL to the file that is associated to the item). For example, the file for feed "NYT"
8 is located at:

9
10 C:\Users\<Username>\AppData\Roaming\Microsoft\RSS\NYT.stg
11

12 In this example, the enclosures are grouped by feed and stored in the
13 Knownfolder "feeds". This enables the user and other applications to easily access and
14 use downloaded files.

15 For example, a user subscribes to the NPR feed and wants to make sure that their
16 media player application can automatically add those files. Making this a Knownfolder
17 enables the user to browse to it from the media player and set it as a monitored folder.
18 Enclosures have the appropriate metadata of the feed and post such that applications can
19 access the associated post and feed. Enclosures are located as follows:

20
21 C:\Users\<Username>\Feeds\<Feedname>\
22

23 Each enclosure that is written to the user's hard disk will have a secondary stream
24 (e.g., a NTFS stream) which contains metadata about this enclosure. The metadata can
25 include by way of example and not limitation, the feed that enclosure is from, author,
link to feed item, description, title, publish date, and download date as well as other
meta data as appropriate.

Publishing Engine/Post Queue

Many times when one writes a regular blog post, essentially what is being written is an RSS item. This RSS item is typically sent to some type of server, and this server maintains account information, the location of the blog, and the like. In this context, publishing engine 110 (Fig. 1) is configured to enable an application to make a posting or publish content, while at the same time abstract from the application the communication protocol that is utilized to communicate with the server. Hence, the application need only provide the data or content that is to be posted, and the publishing engine will handle the remaining task of formatting and communicating the content to the appropriate server.

As there can be several different protocols that are used, abstracting the protocols away from the applications provides a great deal of flexibility insofar as enabling many different types of applications to leverage the publishing functionality. In the illustrated and described embodiment, the publishing engine's functionality is implemented as an API that allows an application to post a blog without having to be knowledgeable of the protocol used to communicate with the server.

Hence, in this example, the API has a method to create a new post which, when called, creates an RSSItem object. This RSSItem object has a post method which, when called, stores the content—in this case a blog—in a temporary store, i.e. post queue 122 (Fig. 1). The content is stored in a temporary store because the user may not be on line at the time the blog is created. Then, when the user makes an on line connection, publishing engine 110 makes a connection to the appropriate server and uses the server-appropriate protocol to upload the blog to the server.

Implementation Example

In the description that follows, an exemplary set of APIs is described to provide but one example of how one might implement and structure APIs to implement the above-described functionality. It is to be appreciated and understood that other APIs can be utilized without departing from the spirit and scope of the claimed subject matter. The described APIs are typically embodied as computer-readable instructions and data that reside on some type of computer-readable medium.

The APIs that are described below can be used to manipulate the set of feeds that a user is subscribed to (System Feed List) and the properties on the feeds. In addition, feed data APIs (i.e., item and enclosures) provide access to feeds that are stored in the feed store, as well as ad-hoc download of feeds. Using the Feed APIs, applications such as web browsers, media players, digital image library applications and the like can then expose the feed data within their experience.

In the example about to be described, the APIs are implemented as COM dual interfaces which also makes the APIs useable from scripting languages, managed code as well as native Win32 (C++) code.

Fig. 6 illustrates a top level object or interface IFeeds and an IFeedFolder object or interface together with their associated properties, methods and events in accordance with one embodiment.

In this example, IFeeds has one property—subscriptions which is an IFeedFolder. This is a root folder for all subscriptions. There are a number of methods on the root object such as DeleteFeed(), DeleteFeedByGuid(), DeleteFolder() and the like.

Of interest in this example is the GetFeedByGuid() method. This method can be called by applications to access a particular feed by, for example, the feed's GUID. Thus, the application need not be knowledgeable of the hierarchical ordering of the feeds. Rather, the feed's GUID can be used by the application to enable the platform to fetch the feed.

1 In addition, the ExistFeed() method checks for the existence of a feed by name,
2 and the ExistFeedByGuid() check for a feed's existence by GUID. The GetFeed()
3 method gets a feed by name or by GUID. The IsSubscribed() method enables an
4 application or caller to ascertain whether a particular feed has been subscribed to.

5 In addition, the IFeeds object also has a SubscriptionsNotifications event which
6 allows for registration for notifications for changes on the system feed list.

7 As noted above, Subscriptions are of the type IFeedFolder. The IFeedFolder
8 object or interface essentially provides a directory and has similar kinds of properties
9 such as the Name, Parent, Path and the like. In addition, the IFeedFolder object has a
10 Feeds property of the type IFeed and a Subfolders property of the type IFeedFolder.
11 The Subfolders property pertains to a collection of the folders underneath the instant
12 folder (e.g., this is where the hierarchical structure derives) and Feeds property pertains
13 to the actual feeds in a particular folder. In addition, the IFeedFolder has a
14 LastWriteTime property which indicates the last time that anything was written to inside
15 the folder. This property is useful for applications that may not have been running for a
16 while, but yet need to look at the feed platform and ascertain its state so that it can
17 synchronize if necessary.

18 There are a number of methods on the IFeedFolder, at some of which pertain to
19 creating a feed (which creates a feed that the system does not have and adds it to a
20 particular folder), creating a subfolder, deleting a folder or subfolder and the like.

21 Fig. 7 illustrates additional objects and their associated methods in accordance
22 with one embodiment. Specifically illustrated are the IFeed, Item and IEnclosure
23 objects.

24 Starting first with the IFeed object, consider the following. Many of the
25 properties associated with this object come from the RSS feed itself, e.g. Title, Url,
Webmaster, SkipHours, SkipDays, ManagingEditor, Homepage, ImageURL and the
like, as will be appreciated by the skilled artisan. In addition, there is another set of

1 properties of interest, i.e. the *Items property* which is a collection that has all of the
2 items that are part of a feed and the *LocalEnclosurePath* property which provides the
3 actual directory to which all of the enclosures are written. Thus, for an application, the
4 latter property makes it very easy for an application to access the enclosures.

5 In addition, this object supports a small set of methods such as Delete() and
6 Download() which are used to manage particular feeds. Further, this object supports a
7 method XML(), which returns a feed's XML in the *common format*. The XML data can
8 be used for such things as creating a newspaper view of a feed. Clone() returns a copy of
9 the feed that is not subscribed to.

10 Moving to the Item object, this object has a set of properties that represent
11 regular RSS elements, e.g. Description, Url, Title, Author and the like. In addition,
12 there is a Parent property that points back to the associated actual feed, and an Id
13 property so that an application can manipulate the Id versus having to iterate over all
14 items. In addition, there is an Enclosures property which is the collection of the item's
15 enclosures of the type IEnclosure. Further, an IsRead property enables an application to
16 indicate whether a particular item has been read.

17 Moving to the Enclosure object, consider the following. This object has
18 properties that include a Type property (e.g. mp3) and Length property that describes
19 the length of a particular enclosure. There is also the LocalAbsolutePath to a particular
20 enclosure. The Download() method allows individual enclosures to be downloaded and
21 used by applications.

22 Conclusion

23 The web content syndication platform described above can be utilized to manage,
24 organize and make available for consumption content that is acquired from the Internet.
25 The platform can acquire and organize web content, and make such content available for
consumption by many different types of applications. These applications may or may

1 not necessarily understand the particular syndication format. An application program
2 interface (API) exposes an object model which allows applications and users to easily
3 accomplish many different tasks such as creating, reading, updating, deleting feeds and
4 the like. In addition, the platform can abstract away a particular feed format to provide a
5 common format which promotes the useability of feed data that comes into the platform.
6 Further, the platform processes and manages enclosures that might be received via a
7 web feed in a manner that can make the enclosures available for consumption to both
8 syndication-aware applications and applications that are not syndication-aware.

9 Although the invention has been described in language specific to structural
10 features and/or methodological steps, it is to be understood that the invention defined in
11 the appended claims is not necessarily limited to the specific features or steps described.
12 Rather, the specific features and steps are disclosed as preferred forms of implementing
13 the claimed invention.

CLAIMS

1 **1.** A system comprising:
2
3 one or more computer-readable media;
4
5 computer-readable instructions on the one or more computer-readable media
6 which, when executed, implement:

7 an RSS platform that is configured to receive and process RSS data in one
8 or more formats; and

9 code means configured to enable different types of applications to access
10 RSS data that has been received and processed by the RSS platform.

11
12 **2.** The system of claim 1, wherein the RSS platform is configured to receive
13 and process RSS data in multiple different formats.

14
15 **3.** The system of claim 1, wherein the RSS platform is configured to receive
16 and process RSS data in multiple different formats, and wherein the RSS platform is
17 configured to convert the multiple different formats into a common format.

18
19 **4.** The system of claim 1, wherein said different types of applications include
20 applications other than RSS readers.

21
22 **5.** The system of claim 1, wherein said different types of applications include
23 applications that do not understand formats in which the RSS data is received by the
24 platform.

1 **6.** The system of claim 1, wherein said different types of applications include
2 applications that do not understand formats in which the RSS data is received by the
3 platform, and wherein the different types of applications comprise one or more of a web
4 browser application, email application or a media player application.

5
6 **7.** The system of claim 1, wherein the different types of applications comprise
7 one or more of a web browser application, email application or a media player
8 application.

9
10 **8.** The system of claim 1, wherein the code means exposes an object model in
11 which feed subscriptions are modeled as a hierarchy of folders, and wherein the object
12 model provides access to a shared list of feed subscriptions.

13
14 **9.** The system of claim 1, wherein the code means is configured to enable one
15 or more applications that are not subscribed to a feed to access associated RSS data that
16 is received and processed by the RSS platform.

17
18 **10.** A system comprising:
19 one or more computer-readable media;
20 a set of APIs embodied on the computer-readable media, the set of APIs
21 comprising one or more methods that enable at least one application to access RSS data
22 that has been processed and stored in a feed store; and
23 wherein said at least one application does not understand an RSS format in which
24 the RSS data was originally embodied.
25

1 **11.** The system of claim 10, wherein said one or more methods comprise a
2 method that provides access to a data store in which one or more enclosures are stored,
3 and a method that can be used to discover a relationship between an enclosure and its
4 associated feed item.

5
6 **12.** The system of claim 10, wherein said one or more methods comprise
7 methods that provide access for multiple different types of applications.

8
9 **13.** The system of claim 10, wherein said one or more methods comprise
10 methods that provide access for multiple different types of applications, and wherein at
11 least one of said multiple different types of applications understands an RSS format in
12 which the RSS data was originally embodied.

13
14 **14.** The system of claim 10, wherein said one or more methods comprise
15 methods that provide access for multiple different types of applications, and wherein at
16 least one of said multiple different types of applications understands an RSS format in
17 which the RSS data was originally embodied and comprises an RSS reader.

18
19 **15.** The system of claim 10, wherein said one or more methods comprise
20 methods that provide access for multiple different types of applications, and wherein at
21 least one of said multiple different types of applications understands an RSS format in
22 which the RSS data was originally embodied and comprises a web browser application.

23
24 **16.** The system of claim 10, wherein said one or more methods comprise a
25 method that enables an application to access data associated with a web feed to which
said application is not subscribed.

1 **17.** The system of claim 10, wherein said one or more methods comprise
2 methods that model feed subscriptions as a hierarchy of folders.

3
4 **18.** The system of claim 10, wherein said at least one application comprises
5 an email application.

6
7 **19.** The system of claim 10, wherein said at least one application comprises a
8 web browser application.

9
10 **20.** The system of claim 10, wherein said at least one application comprises a
11 media player application.

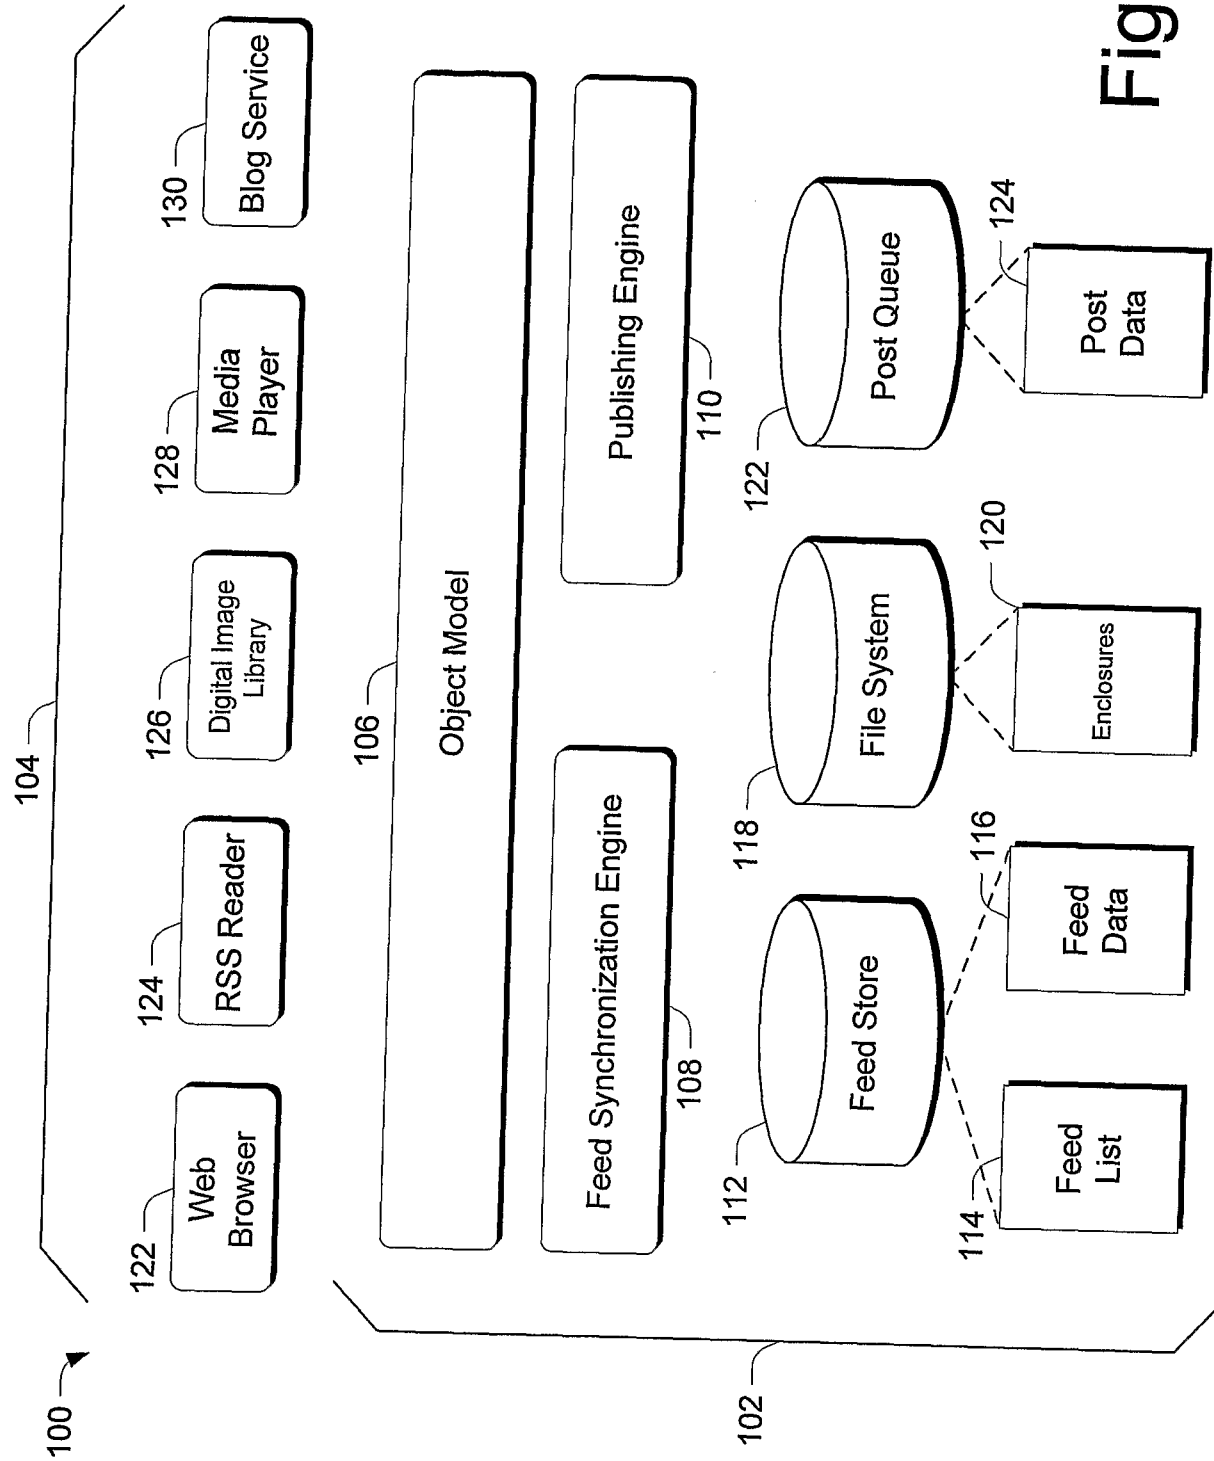


Fig. 1

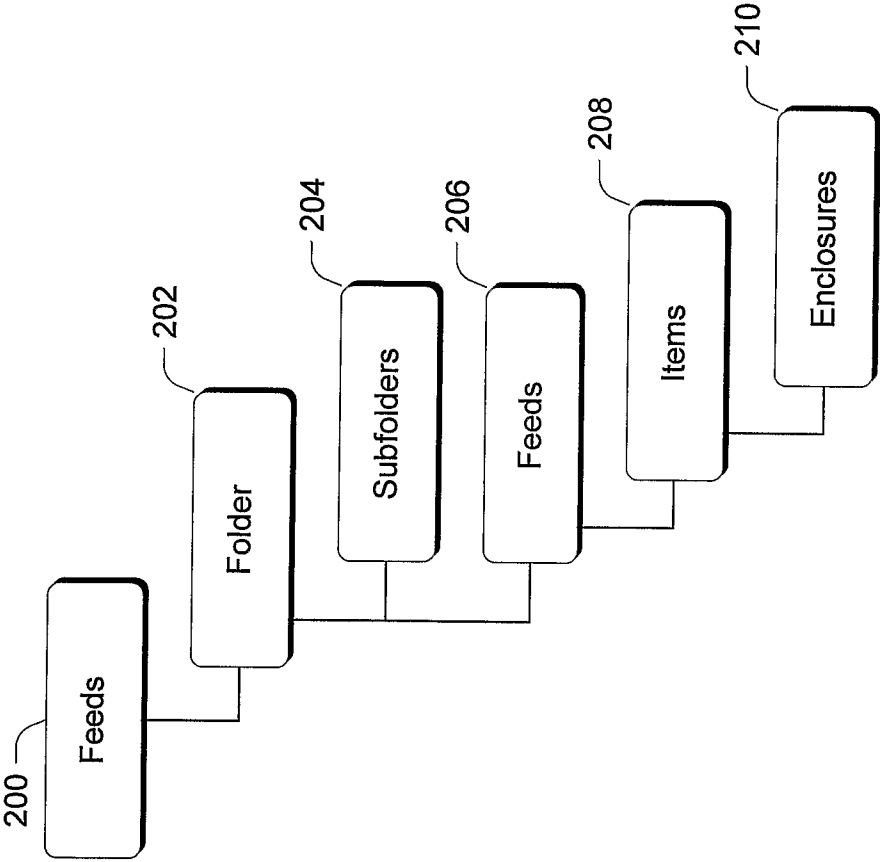


Fig. 2

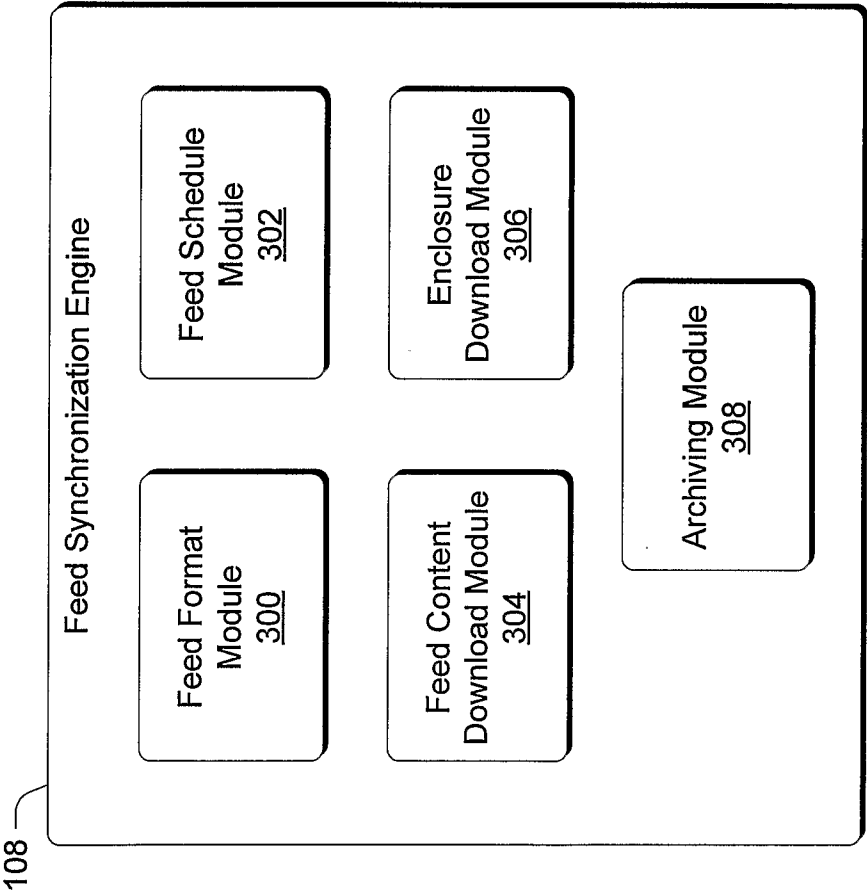


Fig. 3




Fig. 4

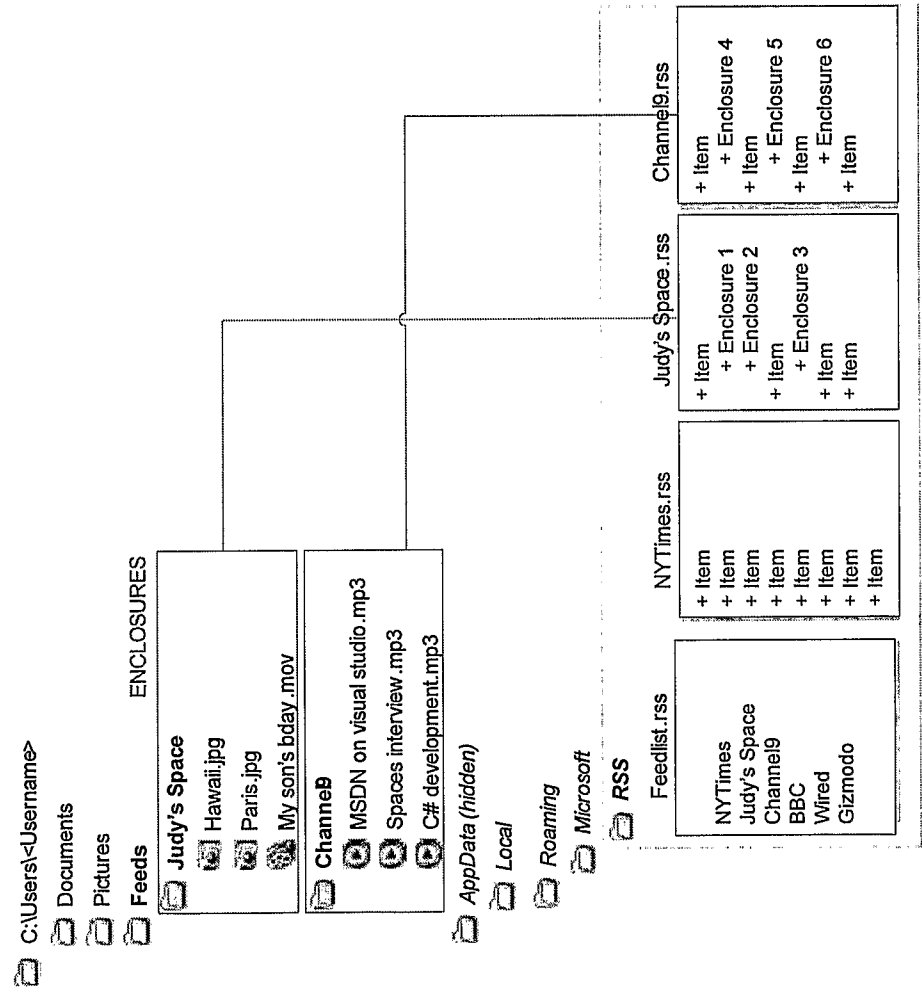


Fig. 5

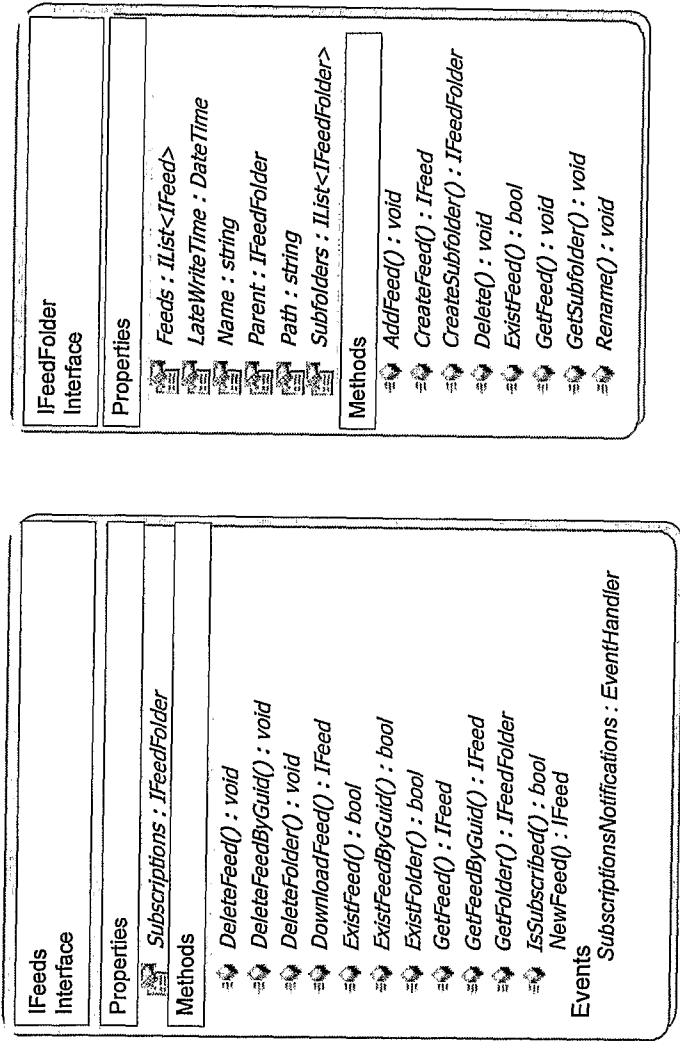


Fig. 6

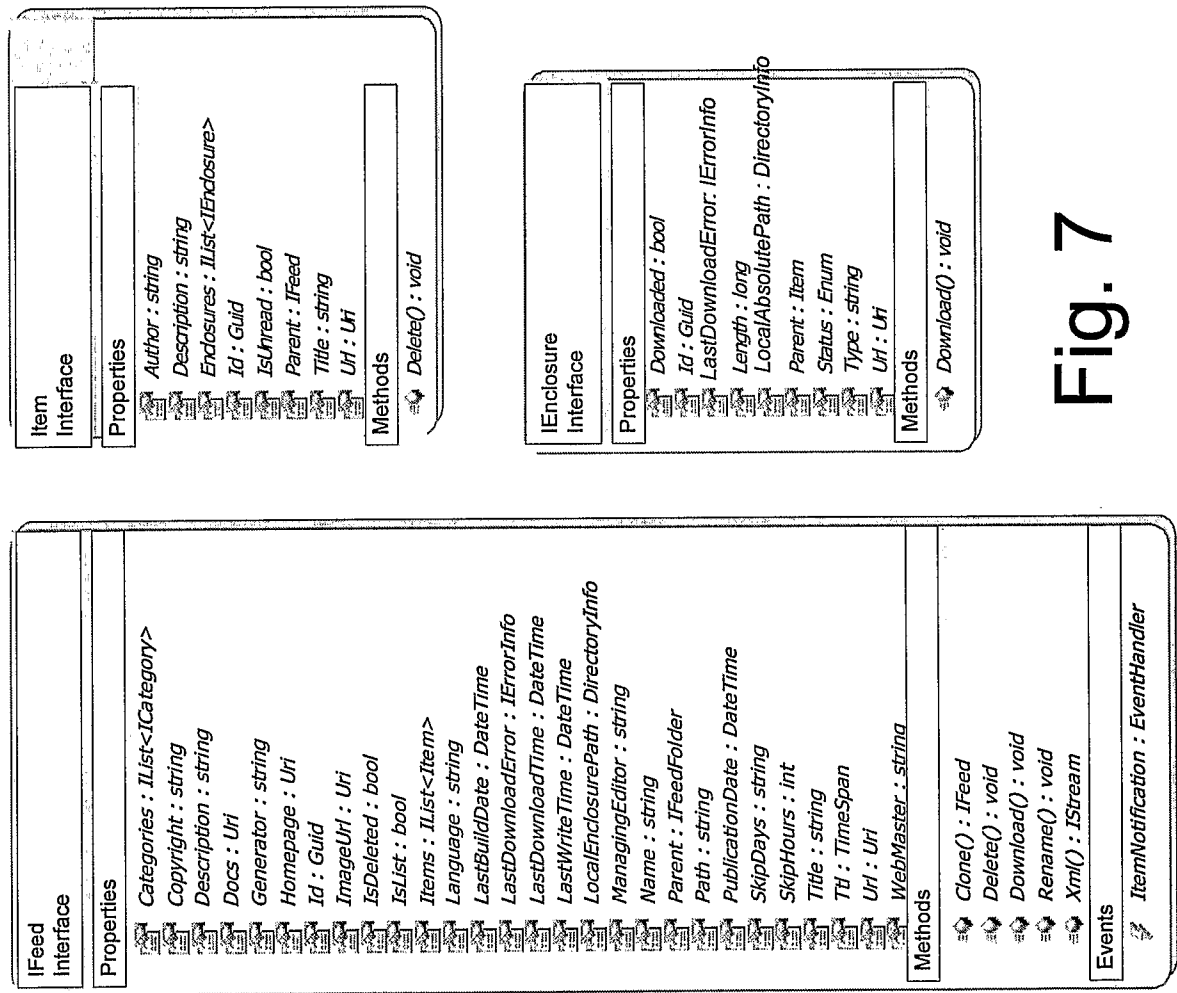


Fig. 7

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US06/23184

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 9/44 (2006.01)

USPC - 717/114

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8) G06F 9/44; G06F 7/00; G06F 15/16; G06F 15/00; G06F 3/00; G06F 9/00; G06Q 30/00; G06Q 20/00 (2006.01)

USPC - 717/114; 707/010; 709/201; 715/513, 760, 749; 705/026, 016

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

MicroPatent (USG USA EPA EPB WO JP DEG DEA DET DEU GBA FRA), IP.com, DialogPro

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2005/0033657 A1 (HERRINGTON et al) 10 February 2005 (10.02.2005) entire document	1-20
A	US 2002/0156905 A1 (WEISSMAN) 24 October 2002 (24.10.2002) entire document	1-20
A	US 2003/0014406 A1 (FAIETA et al) 16 January 2003 (16.01.2003) entire document	1-20
A	US 2004/0230676 A1 (SPIVACK et al) 18 November 2004 (18.11.2004) entire document	1-20
A	US 6,732,102 B1 (KHANDEKAR) 4 May 2004 (04.05.2004) entire document	1-20
A	US 6,708,172 B1 (WONG et al) 16 March 2004 (16.03.2004) entire document	1-20
A	US 2005/0119910 A1 (SCHNEIDER) 2 June 2005 (02.06.2005) entire document	1-20
A	US 5,926,796 A (WALKER et al) 20 July 1999 (20.07.1999) entire document	1-20

☐ Further documents are listed in the continuation of Box C.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

19 October 2006

Date of mailing of the international search report

30 NOV 2006

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774