



(19) **United States**

(12) **Patent Application Publication**
Gwozdz

(10) **Pub. No.: US 2011/0087750 A1**

(43) **Pub. Date: Apr. 14, 2011**

(54) **RESOURCE LOCATORS FOR WIDELY DISTRIBUTED SYSTEMS**

(52) **U.S. Cl. 709/206; 709/219**

(75) **Inventor: Daniel Gwozdz, Fall City, WA (US)**

(73) **Assignee: Microsoft Corporation, Redmond, WA (US)**

(21) **Appl. No.: 12/967,591**

(22) **Filed: Dec. 14, 2010**

Related U.S. Application Data

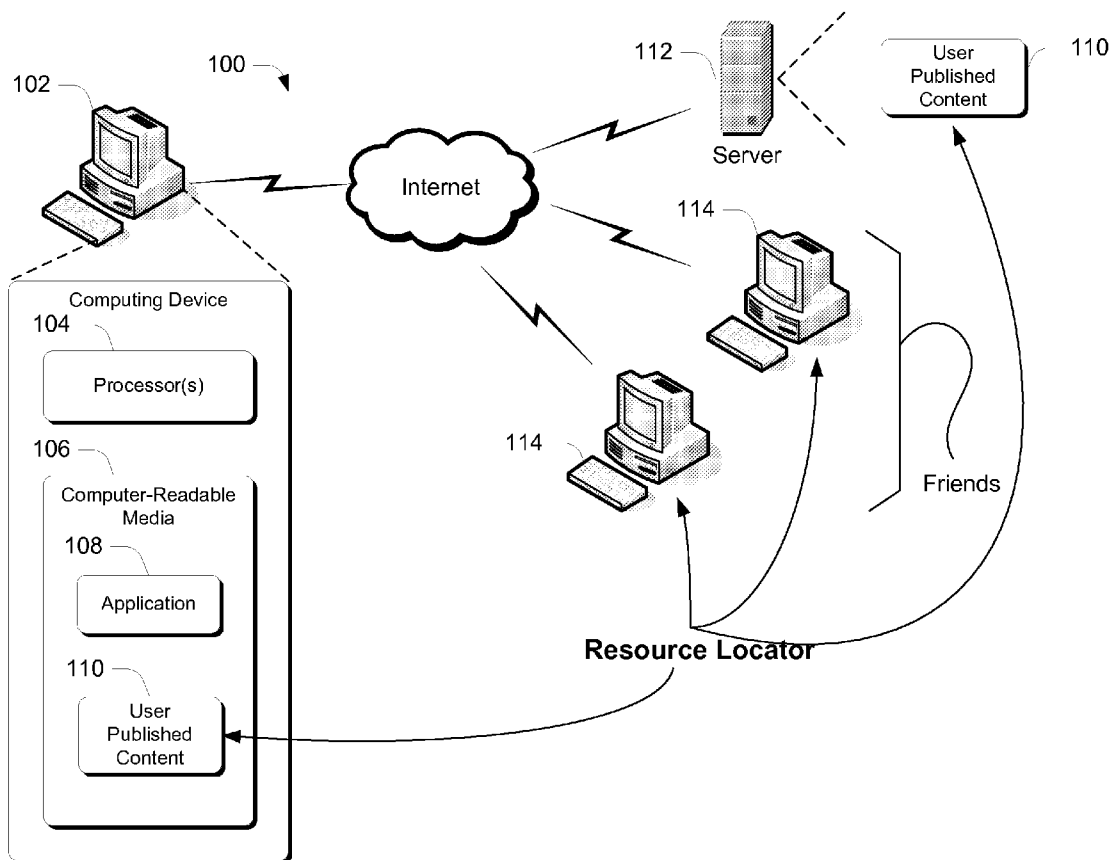
(63) Continuation of application No. 11/420,972, filed on May 30, 2006, now Pat. No. 7,882,095.

Publication Classification

(51) **Int. Cl. G06F 15/16 (2006.01)**

(57) **ABSTRACT**

Implementations provide simple syndication, browsing, searching, and subscribing solutions to networks of friends. In at least some embodiments, the approaches utilize standard HTTP protocols for consumption, subscription, and interaction of data while using the local file system for publication. Some implementations overlay a URL address space that can be handled and routed interchangeably between multiple clients and servers or "locations". In this manner, generated requests for user published content can be satisfied by local caches, servers, or peers through, for example, a peer-to-peer network. Accordingly, some implementations provide a platform to make this possible with a simple publishing model that is fully compatible with existing Internet enabled applications.






Fig. 1




Fig. 3




Fig. 2




Fig. 4

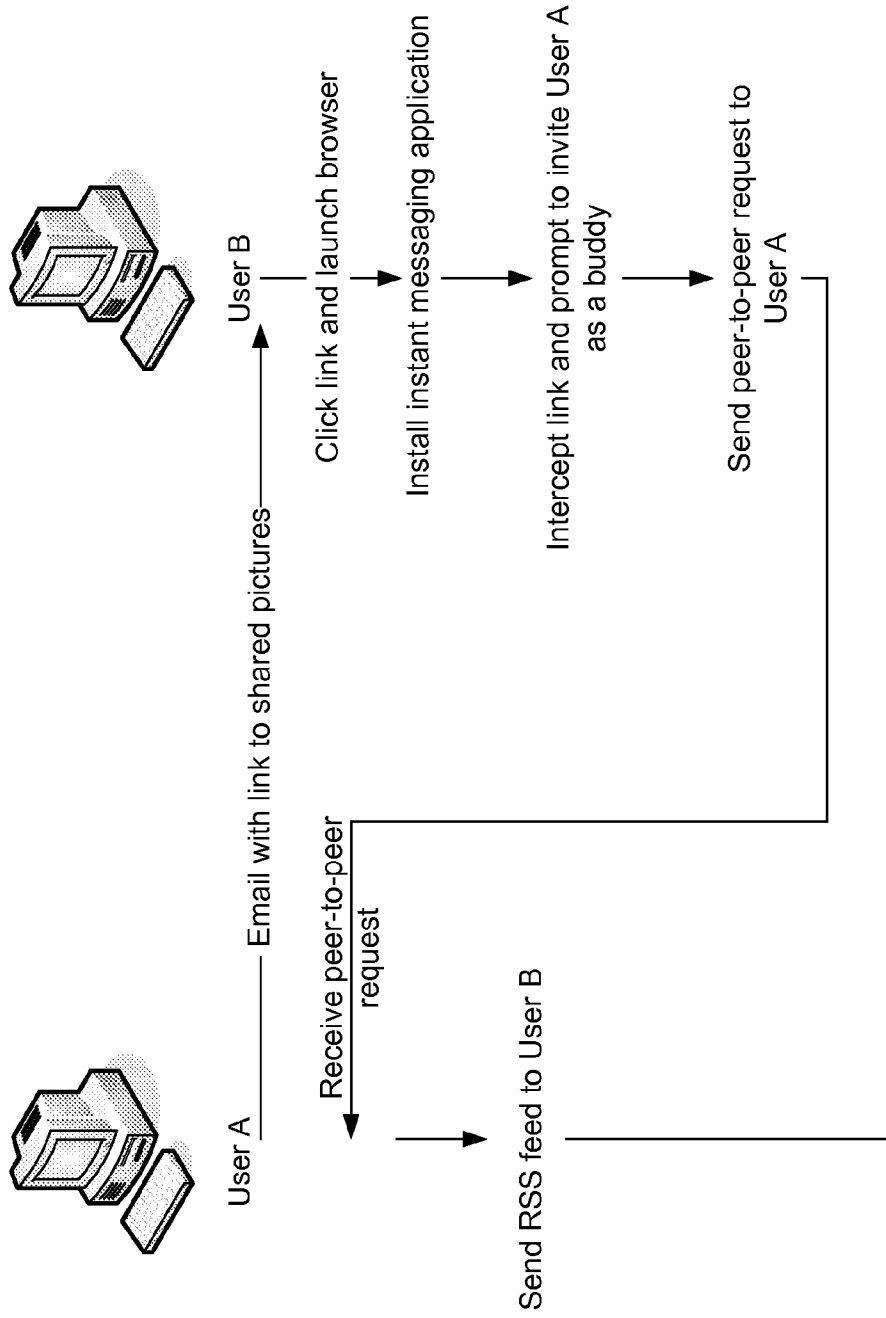


Fig. 5

RESOURCE LOCATORS FOR WIDELY DISTRIBUTED SYSTEMS

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application is a continuation application of, and claims priority to, U.S. patent application Ser. No. 11/420,972, filed on May 30, 2006, to the same inventor as herein, the entire disclosure of which is incorporated herein by reference.

BACKGROUND

[0002] To date, when a user wishes to publish their own content on line for their friends, they face a number of challenges. For example, they may have different kinds of content, such as applications, activities, photos, documents, blogs and the like. Each of these may have their own unique publishing requirements or forums. Many users tend however, to publish their own content to their friends by way of email, as by sending an email with an attachment in the form of the published content. Yet, the email experience in terms of the publishing solution that it provides is not the best technique or solution.

[0003] In terms of content consumption, there are a variety of rich solutions that use Internet standards, such as RSS (Really Simple Syndication), to provide a great user experience. Yet, the gap between easy publishing solutions and rich consumption solutions remains to be meaningfully and richly bridged.

SUMMARY

[0004] The methods and systems described below provide simple syndication, browsing, searching, and subscribing solutions to networks of friends. In at least some embodiments, the approaches utilize standard HTTP protocols for consumption, subscription, and interaction of data while using the local file system for publication. The inventive approach overlays a resource locator or URL address space that can be handled and routed interchangeably between multiple clients and servers or "locations". In this manner, generated requests for user published content can be satisfied by local caches, servers, or peers through, for example, a peer-to-peer network.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 illustrates an exemplary system in which various embodiments can be employed in accordance with one embodiment.

[0006] FIG. 2 is a flow diagram that describes steps in a method in accordance with one embodiment.

[0007] FIG. 3 is a flow diagram that describes steps in a method in accordance with one embodiment.

[0008] FIG. 4 is a flow diagram that describes steps in a method in accordance with one embodiment.

[0009] FIG. 5 is a high level overview of a walk through of one example of how various embodiments can be employed.

DETAILED DESCRIPTION

Overview

[0010] The methods and systems described below provide simple syndication, browsing, searching, and subscribing solutions to networks of friends. In at least some embodi-

ments, the approaches utilize standard HTTP protocols for consumption, subscription, and interaction of data while using the local file system for publication. The inventive approach overlays a resource locator or URL address space that can be handled and resolved interchangeably and routed interchangeably between multiple clients and servers or "locations". In at least some embodiments, the resource locator is resolvable in the traditional manner, but is also capable of being locally intercepted and alternatively resolved. In this manner, generated requests for user published content can be satisfied by local caches, servers, or peers through, for example, a peer-to-peer network. The inventive solutions provide a platform to make this possible with a simple publishing model that is fully compatible with existing Internet enabled applications.

[0011] The discussion that follows is structured in the following way. First, a section entitled "Exemplary Environment" provides a description of one environment in which the various embodiments can be employed. Following this, a section entitled "Exemplary Resource Locator" describes characteristics of an exemplary resource locator and gives an example of one. Next, a section entitled "Using the Resource Locator" describes how an application can use the exemplary resource locator. This section includes three sub-sections: "Publishing Content Using the Resource Locator" describes how a user can build an association between their content and the resource locator; "Consuming Content Using the Resource Locator" describes how applications can consume content using the resource locator; and, finally, "End User Scenarios" provides some exemplary end user scenarios in which the exemplary resource locator is employed.

Exemplary Environment

[0012] FIG. 1 illustrates an exemplary environment, in accordance with one embodiment, in which the inventive embodiments can be employed generally at 100.

[0013] Here, system 100 includes a client computing device 102 which includes one or more processors 104 and one or more computer-readable media 106 on which executable, computer-readable instructions reside. In this example, computer-readable media 106 includes code that implements one or more applications 108. The various embodiments described below can be implemented by any suitable type of application. In but one example, the inventive embodiments are implemented in connection with an instant messaging application. Microsoft's Live Messenger constitutes but one example of such an application. It is to be understood that other applications can be utilized without departing from the spirit and scope of the claimed subject matter.

[0014] Generally, instant messaging applications allow users to instantly communicate and collaborate with what are termed their "contacts" or "buddies"—also referred to as "friends" in this application. A user can build a list of contacts or buddies by providing a friendly name such as "Paul" and an address. Instant messaging applications utilize the notion of "presence", which makes instant communication and collaboration with contacts even more effective. Presence gives users the ability to find each other and stay constantly updated with each other's online status. The powerful capabilities of presence are usually provided through a presence-based network that utilizes one or more servers or services, as will be appreciated by the skilled artisan.

[0015] The protocol used for initialization and communication on an instant messaging session depends upon the

server or service being used. For example, text can be carried over a TCP connection, a TCP UDP or Secure Sockets Layer (SSL), to name just a few. Typically, an initiating client sends a request to start a conversation with the contact to the server, which is then forwarded to the other client. The instant messaging communication can then proceed. Specifically, message text can now be sent to the server and forwarded to the other client. How the message text is delimited depends upon the server and protocols being used. In some systems, text is sent using a Hypertext Transfer Protocol (HTTP) or Extensible Markup Language (XML) in HTTP message encapsulated in the TCP connection with the server. Other features that instant messaging applications can provide include, by way of example and not limitation, voice and video calls, application sharing and file transfer, all of which will be appreciated by the skilled artisan.

[0016] Continuing, computing device **102** also includes user-published content **110** which can include any type of content. For example, such content can include, by way of example and not limitation, a user's pictures or images, favorites (e.g. hyperlinks), lists, ratings, comments, annotations, video, audio, blogs, applications created by the user, daily thoughts and the like. With respect to applications that are published by a user, consider the following.

[0017] One of the things that HTTP allows is the sending of verbs such as PUT and GET. In this context, user published content can reside in the form of an application that can execute actions on a remote device using GET or PUT parameters. The inventive embodiments described above and below can allow for remote publishing of content using a PUT verb. Hence, applications running on a remote device (such as WEBDAV applications or a media center control panel) or an application that runs on the requestor's device (such as a web page or an AJAX application) or both (such as an AJAX interactive conversation activity) can be enabled through the use of the inventive embodiments described herein. At least one example of such an application is provided below.

[0018] As such and in addition, system **100** includes other locations where the user's published content can reside. For example, such user-published content can reside on a server **112** and/or in various local caches on the computing devices of the user's friends, such computing devices shown generally at **114**.

[0019] It is to be appreciated and understood that while the computing devices in the FIG. **1** example are illustrated as desktop devices, such devices can take other forms such as laptop or notebook computers, handheld computers, or any other type of computing device that can be used in connection with instant messaging functionality.

Exemplary Resource Locator

[0020] In the illustrated and described embodiment, the user can utilize a unique resource locator or URL, and can associate the unique resource locator with their user-published content. Accordingly, the URL is redefined, in a sense, to reference an individual's or user's content and what that content is. This resource locator is resolvable to multiple different locations at which their published content can reside, as indicated in the illustration. By using a unique, multiple location-resolvable resource locator, the user can publish their content to their friends in an easy and straightforward manner, as will become apparent below.

[0021] In one embodiment, the resource locator is an Internet-compatible, universal naming scheme for addressing and

enumerating data or information based on people or individuals. In the context of this document, an "individual" can include, by way of example and not limitation, a person, a group of persons, and/or an entity such as a company or corporation. In at least one embodiment, the resource locator has multiple parts. As but one non-limiting example of such a resource locator, consider the following.

[0022] In one embodiment, a first part of the resource locator constitutes a universal prefix. This prefix is used as an identifier to identify the scheme of which the resource locator is a part. In addition, the prefix allows the resource locator to be resolved by either a client or a server. An example of a first part of the resource locator might be "friendcast.msn.com".

[0023] In one embodiment, a second part of the resource locator is associated with or defines an endpoint. This part of the resource locator identifies or specifies whose data the resource locator represents. An example of this part is "dangw@msn.com". Hence, the data or information with which the resource locator is associated belongs to or is associated with an individual whose alias is "dangw". Optionally, this part of the resource locator can also be associated with a specific machine or computing device. In this case, this part of the identifier might look like: "dangw@msn.com@home", where the "home" portion identifies a particular machine or computing device.

[0024] In one embodiment, a third part of the resource locator is associated with what is referred to as an application namespace. This part of the identifier represents or identifies an application with which the requested data or information is associated. There can be several different types of applications. For example, applications such as MyPictures or MyCalendar can be associated with a user's pictures or calendar information respectively. In this case, this part of the locator might look like "My %20Pictures".

[0025] In one embodiment, a fourth part of the resource locator is associated with application-specific data or information that is associated with the application referenced by the third part of the resource locator. In this particular example, if the third part of the resource locator is "MyPictures", then an exemplary fourth part might look like: "/Hawaii/img01.jpg". Hence, in this example, the fourth part would reference a JPG image (img01) presumably taken in Hawaii.

[0026] Putting all of the parts of the exemplary resource locator together, one would have the following:

[0027] "http://friendcast.msn.com/dangw@msn.com@home/MyPictures/Hawaii/img01.jpg".

[0028] In the illustrated and described embodiment, and as will become apparent below, this resource locator can be placed into a browser, an RSS reader or any other Internet-enabled application and subsequently routed and resolved as described below.

[0029] In one embodiment, the endpoint part of the resource locator can be encoded so that it does not identify an individual or a specific machine. For example, in the resource locator above, the endpoint "dangw@msn.com@home" may be encoded and take the form "CID@GUID", where CID (which maps to "dangw@msn.com") and GUID (which maps to "@home") are unique strings that are mappable to the specific endpoint reference that names an individual and, optionally, a specific computing device.

[0030] Further, in another embodiment, the resource locator may include a session GUID that uniquely identifies a particular session which is referenced by the resource locator.

In the above example, the session identifier might be prepended to the resource locator. This would be used, for example, in situations where two individual users are participating in a session-associated activity, such as playing a game that is referenced by the session identifier. Such can be the case, for example, where one user publishes a new chess game to his friends. Specifically, by sending a link in the form of a resource locator that identifies the chess game to a friend, the user can initiate a game of chess that is identified by the session identifier. This way, the user can participate in multiple different games of chess without the risk of a collision between the games. One way in which this type of interaction can take place is described below in more detail.

Using the Resource Locator

[0031] A resource locator, such as the one described above, can be used to associate resources or information, such as published content, with different resolvable locations. This can enable individuals to publish their information to their friends and others in a very powerful, yet simple way. The discussion below starts first with an explanation of how one can utilize the resource locator to make this association, and then follows with an explanation of how one can consume user-published content using the resource locator.

Publishing Content Using the Resource Locator

[0032] FIG. 2 is a flow diagram that describes steps in a method in accordance with one embodiment. The method can be implemented in connection with any suitable hardware, software, firmware or combination thereof. In but one embodiment, the method can be implemented using an application that executes on a user's computing device.

[0033] Step 200 defines a resource locator that is resolvable to multiple different locations. This step can be implemented by, for example, a software application that presents the user with a user interface through which they can define the resource locator. The multiple different locations can correspond to locations at which information, such as user-published or publishable content resides or can reside. For example, a user can keep their published content on their local computing device in a particular directory. Additionally, the user may maintain some of their published content on a remote server that is accessible through the Internet. Alternately or additionally, the user may have some of their published content already present on one or more of their friends' computing devices. In at least some embodiments, the resource locator is resolvable using a presence-based network, such as one that supports presence-based applications like instant messaging applications.

[0034] Step 202 associates the resource locator with information, such as the user's published content, which can be located at the multiple different locations. This step can be performed as part of the definition step mentioned above. For example, in the example above, the resource locator includes a first part that defines a universal prefix which is resolvable by a client or a server, and a second part that defines a user and, optionally, a computing device. The resource locator also refers to an application namespace (the third part) and application-specific data or information (which, in this example, is the user-published content). In this case, this step is performed by providing the last two parts of the resource locator described above.

[0035] Accordingly, once the user has created their content, defining a resource locator and associating it with their content makes this content easily available to their friends and others. In addition, as will become apparent below, in at least some embodiments, any updates to a user's published content can be quickly and simply conveyed to their friends without any intervention by the user who created the content.

[0036] It is to be appreciated and understood that the user's published content can be published on their local machine and/or optionally pushed to a server cache.

Consuming Content Using the Resource Locator

[0037] FIG. 3 is a flow diagram that describes steps in a content consumption method in accordance with one embodiment. The method can be implemented in connection with any suitable hardware, software, firmware or combination thereof. In but one embodiment, the method can be implemented using an application that executes on a user's computing device.

[0038] Step 300 generates a request for information, such as user-published content, represented by a resource locator. This step can be performed on a user's local machine or device by, for example, providing a resource locator, such as the one described above, to a suitably configured application, such as a browser or RSS reader. In the illustrated and described embodiment, the request can be an HTTP request.

[0039] Step 302 intercepts the request. This step can be performed by any suitably configured component or application. For example, this step can be performed by a suitably configured handler component which may or may not comprise part of an instant messaging application. In the illustrated and described embodiments, interception takes place locally.

[0040] Step 304 attempts to resolve the request, on the local machine or device, using a presence-based network. In the illustrated example, the presence-based network can comprise one that is utilized by an instant messaging application. This step can be performed in any suitable way. For example, as part of implementing this step, the intercepting component can check and ascertain whether a friend or user with which the request is associated is on-line and whether there is a peer-to-peer connection established with this friend and the friend's remote device. That is, this friend may appear on the user's buddy list and if so, the presence-based network can easily ascertain if this friend is on-line. If there is a peer-to-peer connection established with this friend, then the request can be resolved using the peer-to-peer connection.

[0041] If, on the other hand, this request cannot be resolved using the presence-based network, step 306 can optionally attempt to resolve the request using a local cache. For example, the information that is the subject of the request may have been previously cached on the user's local device. If this is the case, then the request might be resolved by resort to the local cache. Of course, it is to be appreciated and understood that this step can occur before step 304.

[0042] Alternately or additionally, if the request cannot be resolved using the presence-based network, then step 308 can attempt to resolve the request through a remote server. That is, the information or content with which the request is associated (or a form thereof) may also be maintained on a remote server or accessible through a remote server. In this case, the method can attempt to resolve the request using the remote server. For example, the information or content may be content that belongs to a friend of a friend. For example, if users

A, B and C are friends, user A might be able to get user B's content from user C's cache. In this example, the remote server can facilitate acquisition of the content by, for example, ascertaining whether the requested content resides in its cache and/or attempting to establish a peer-to-peer connection between user A and user C.

[0043] FIG. 4 is a flow diagram that describes steps in another content consumption method in accordance with one embodiment. The method can be implemented in connection with any suitable hardware, software, firmware or combination thereof. In but one embodiment, the method can be implemented using an application that executes on a user's computing device.

[0044] Step 400 generates, on a local device, a request for information represented by a resource locator having an endpoint. Examples of resource locators are given above. In this particular embodiment, the resource locator is resolvable by an appropriately configured application via a presence-based network using a peer-to-peer connection, or through a remote server that does not use a peer-to-peer connection. In this example, the application can comprise any suitably-configured application, such as an instant messaging application.

[0045] If the local device does not have the application, or if the peer-to-peer connection does not exist with a computing device associated with the endpoint, step 402 attempts to resolve the request using a remote server. Alternately, if the local device has the application or if the peer-to-peer connection does exist, the request can be attempted to be resolved using the application or the peer-to-peer connection as described above.

[0046] In an event that the local device does not have the suitably configured application, the remote server can invite the user to install the application. That is, responsive to attempting to resolve the request, the server may prompt the user to install the application. The user can then opt to install the application and proceed to attempt to have the request resolved.

[0047] If the user's request is associated with a user that does not appear on their buddy list, then the server can send the user originating the request a communication asking whether the request-originating user wishes to invite the user associated with the endpoint to join the request-originating user's buddy list. If the request-originating user sends an invitation and the invitation is accepted, then the request can be attempted to be resolved using a peer-to-peer connection between these users.

[0048] In this instance, an original request which may not have had the appropriate application to complete the request, or which did not have the appropriate permissions (as by being associated with a buddy that appears on the request-originating user's buddy list) is nonetheless capable of being resolved by performing a couple of different actions: installing the application and/or adding the endpoint-associated user to the appropriate buddy list.

End User Scenarios

[0049] The above-described embodiments can be used in various routing, resolution, syndication and consumption scenarios. The following examples are provided to give the reader a tangible idea of how the various embodiments can be employed. The examples are not intended to limit application of the claimed subject matter.

[0050] With regard to routing and resolution, consider the following. When a resource locator such as that describe

above is requested, a component, such as an http server, installed on the client can intercept and resolve the request. This client component can route the request as appropriate. For example, it can redirect the request to a server, satisfy the request with a locally cached response, or tunnel the request through a peer-to-peer transport directly to the owner of the data as described above. If the local http server is not installed, the local system will naturally direct the request to the appropriate server which can either satisfy the request or install the client component as described above.

[0051] With regard to syndication, consider the following. When the request is routed to the publisher of the content, a registered handler installed on the publisher's computing device can satisfy the request. The handler can map the resource locator to a hard drive location and return either a file as is, or a folder as an RSS feed. In the case where the resource locator includes a search parameter, a desktop search application, such as Windows® Desktop Search can run on the publisher's machine to generate an RSS search response. The RSS feed can also contain an XSLT reference to format the data to HTML for browsers. If the request was a WebDAV request, the response can be formatted as XML that is used by the client to give a file system experience.

[0052] With regard to consumption of content, consider the following. Since the address is a resource locator or URL and the protocol being used is HTTP, any web-enabled application can consume the data. For example, the data could be HTML, XML, RSS, or binary files like JPEG. In some cases a richer set of verbs is possible, such as those provided by WebDAV, as will be appreciated by the skilled artisan. In a browser the HTML can be a web application just like conventional servers provide. Most things that can be done with HTTP can now be done in a friend-to-friend manner, without the complications of setting up a server. Widgets (or Gadgets) could be simple RSS readers subscribed to a friend's URLs.

[0053] In addition, when a resource locator is resolved, at least some embodiments can aggregate lists of information from different locations and present those aggregated lists to a user. For example, a user may have some of their published content at a remote server and other of their published content on their local device. If this is the case, and a request for this data is received, the list that is returned to the requester can contain an aggregation of the content at both locations.

[0054] With regard to security, consider the following. When an HTTP request is intercepted locally, the request can be routed and transformed through the appropriate permissions and security. Although the local client is issuing plain text HTTP requests, the data that leaves and enters the machine can be restricted via ABCH and encrypted via a passport.

[0055] FIG. 5 illustrates a walkthrough that shows an end to end scenario in an instant messaging application in the form of Windows® Live Messenger. Here, permissions are enforced through the buddy list and encryption is through the secure peer-to-peer transport. In this example, since the Messenger code base is invoked, it can provide security prompts and enforcement without the client application being aware. Security is applied at both ends of the communication ensuring that only mutually verified criteria will allow a transaction to complete.

[0056] In FIG. 5, there are two users—User A and User B. User A wishes to share his pictures with User B. Initially, User A, using an instant messaging application, sends an e-mail to

user B, as indicated in FIG. 5, containing a link to his shared pictures. This link might assume the form:

[0057] `click here`

[0058] User B, who is not an instant messaging application user, clicks the link which launches her browser and takes her to the friendcast.msn.com server which up sells her to install MSN Messenger.

[0059] After User B installs Messenger, the link is intercepted by the Messenger client and User B is prompted to invite User A as a buddy. That is, in this example, the response is a page asking her if she'd like to invite User A as a buddy.

[0060] After User B has User A as a buddy, the link will cause Messenger to send a peer-to-peer request to User A. User A's Messenger client receives the request from User B, validates her permissions, and generates an RSS file where each item is either a folder with a FriendCast link to the folder, or a file with an enclosure link to the file. The feed is sent back to User B's machine.

[0061] User B's browser can now render the RSS file with an applied XSLT to present a pleasing HTML view. User B can look at the files and folders and click on them to browse or download. User B can type, for example, 'vacation' into the search field of the rendered page. This, in turn, navigates the browser to the link

[0062] `"http://friendcast.msn.com/a@msn.com/Shared%20Pictures?search=vacation"`,

[0063] which is routed to User A.

[0064] User A's client handles the request by creating an RSS response from the files found using a search program, such as Desktop Search. User B sees the search results and clicks on an 'RSS' button on the page. Her registered RSS reader is launched and the search results are shown as items in a feed.

[0065] Now, assume that User A downloads new photos from his camera and tags them with the 'vacation' category using his favorite photo management tool. User B's RSS reader can poll the URL and show the new photos within moments of their availability. User B is now richly subscribed to User A's photos and User A needs to take no explicit action for new photos to be available to User B after copying the photos off his camera.

[0066] If, in the example above, the initial link from User A was presented by an RSS reader instead of a browser, then the up sell and add the buddy prompts would have shown up as RSS items.

More Examples

[0067] As one can surely appreciate, there are simply enumerable user scenarios that are made possible using the above-described embodiments. As for some further non-limiting examples, consider the following.

[0068] Using the above-described approach, a link on an MSN Space photo, for example, could fetch a full resolution photo using a peer-to-peer connection, or a medium resolution photo from a server. If both buddies are online at the same time, the Spaces server would be spared the throughput cost. Also, if the photo is no longer on the server and if not peer-to-peer system is in place, the server can fetch the photo, store it and then satisfy the request.

[0069] A browser application, such as Internet Explorer® 7, can automatically download calendar appointments in an RSS feed and place them in and overlaid calendar in Outlook®. A calendar handler can automatically publish some-

one's Outlook® calendar to their friends using the techniques described above. Accordingly, within minutes of creating an appointment, a user's friends would be able to see the appointment.

[0070] Additionally, a user can record a daily thought to an audio file, save it to their hard drive, and within moments, that file is placed on their friends' media devices by their media casting RSS reader.

[0071] As another example, a user can save their comments (blog) in a text file to their hard drive. The system can then automatically syndicate from the local drive content in an RSS format.

[0072] As further examples, consider the following.

[0073] One content type that can be published is an HTML page with JavaScript that is an AJAX application. The code runs locally and it can make requests that causes the code to run remotely. As an example, the AJAX application might implement a game—such as a game of chess that could operate as follows. User A receives a link from user B to an instance of a game of chess in the form of an HTML application. User A's browser runs the application that makes the requests to User B to get the state of the game then to render it and enforce rules. When user A makes a move it generates a request to user B. User B's application executes and store the move, and then updates both users to the new state. When user B makes a move, the request is handled locally to execute and store the move and then update both users to the new state. All requests in this scenario reside in the form of an URL, such as those described above.

[0074] Further, the content can be an application that resides at the location. One example of this is content that is an HTML page that invokes remove verbs by use of a URL with GET or PUT parameters. For example, a media center personal computer can publish an HTML page with links that get routed to the media center personal computer and cause it to take actions such as setup a recording or delete a recording. Alternately or additionally, the media center personal computer can expose these verbs as a SOAP API thus allowing two media center personal computers to discover and communicate with each other exchanging verbs and content.

[0075] As but another example, a request on a down-level device can get routed to a server which then routes the request to another location or locations. This request could go to the server since a local handler might not exist. The server can then proxy the request to one or more local machines through a presence network. For example, a user can navigate to a URL on their web-enabled phone to search their home computer and other computers for a document, and then tell the computer(s) to e-mail the document to their work e-mail account. The request goes to the server then the server routes the request to the user's computer(s). The computer(s) return the search results to the server which can then aggregate results and transform them to a phone-consumable format. Links then can invoke verbs, routed similarly, that tell the computer(s) to send the e-mail.

CONCLUSION

[0076] The methods and systems described above provide simple syndication, browsing, searching, and subscribing solutions to networks of friends. In at least some embodiments, the approaches utilize standard HTTP protocols for consumption, subscription, and interaction of data while using the local file system for publication. The inventive approach overlays a URL address space that can be handled

and routed interchangeably between multiple clients and servers or “locations”. In this manner, generated requests for user published content can be satisfied by local caches, servers, or peers through, for example, a peer-to-peer network. Accordingly, the inventive solutions provide a platform to make this possible with a simple publishing model that is fully compatible with existing Internet enabled applications. [0077] Although the invention has been described in language specific to structural features and/or methodological steps, it is to be understood that the invention defined in the appended claims is not necessarily limited to the specific features or steps described. Rather, the specific features and steps are disclosed as preferred forms of implementing the claimed invention.

1. A method comprising:
 - generating, by a processor of a computing device, a request for content, wherein the request is represented by a resource locator comprising an endpoint associated with the content; and
 - attempting, by an application, executed on the computing device, that accesses a presence-based network, to resolve the request using the presence-based network by:
 - determining whether a participant of the presence-based network, referred to by the resource locator, has a current presence on the presence-based network; and
 - when the participant of the presence-based network, referred to by the resource locator, has a current presence on the presence-based network, obtaining the content from the participant referred to by the resource locator.
2. The method according to claim 1, further comprising, when attempting to resolve the request using the presence-based network fails, resolving the request by resolving the resource locator through a remote server having access to a copy of the content.
3. The method according to claim 2, wherein the copy of the content accessed through the remote server is of a different resolution or quality than the content available from the participant through the presence-based network.
4. The method according to claim 1, further comprising, prior to attempting to resolve the request using the presence-based network, attempting to resolve the request by resolving the resource locator to obtain a copy of the content stored on a local cache of the computing device.
5. The method according to claim 1, wherein the application is an instant messaging application that establishes a presence of a user of the computing device as participating in the presence-based network when the instant messaging application is initiated to access the presence-based network.
6. The method according to claim 1, wherein:
 - the participant of the presence-based network is one of a plurality of participants, each having a presence in the presence-based network; and
 - the presence-based network is monitored by one or more servers coordinating the presence-based network for determining participants currently having a presence in the presence-based network.
7. The method according to claim 1, wherein:
 - the participant of the presence-based network is one of a plurality of participants of the presence-based network that are affiliated based on a buddy list; and
 - members of the buddy list are able to establish a presence in the presence-based network with respect to other

members of the buddy list, but not with respect to participants of the presence-based network that are not members of the buddy list.

8. The method according to claim 7, further comprising receiving, as a result of attempting to resolve the request using the presence-based network, a communication asking whether a first user of the computing device wishes to invite a second user associated with the endpoint to join a buddy list of the first user.
9. The method according to claim 1, wherein the resource locator comprises:
 - a first part including a universal prefix;
 - a second part defining the endpoint;
 - a third part associated with an application namespace; and
 - a fourth part which pertains to application-specific information.
10. The method according to claim 1, wherein the endpoint is further associated with at least one of:
 - an individual corresponding to the participant; or
 - a specific computing device corresponding to the participant.
11. The method according to claim 1, wherein:
 - the resource locator further comprises a session GUID that identifies and is unique to a particular session; and
 - the session GUID is used to identify a particular session between a user of the computing device and the participant to avoid collisions with other sessions between the user and other participants of the presence-based network.
12. A computing device comprising:
 - a processor coupled for communication with computer-readable media; and
 - an application maintained on the computer-readable media and executed on the processor to perform operations including:
 - receiving a request for content, wherein:
 - the content is published by a user of a remote computing device;
 - the request is represented as a resource locator comprising an endpoint associated with the content; and
 - the resource locator is resolvable to obtain the content through a presence-based network using a peer-to-peer connection or through a remote server that does not utilize a peer-to-peer connection; and
 - resolving the resource locator to obtain the content corresponding to the request.
13. The computing device according to claim 12, further comprising a local cache on the computing device, wherein the resource locator is also resolvable to obtain the content from the local cache on the computing device.
14. The computing device according to claim 12, wherein:
 - the application is configured to attempt to resolve the resource locator to obtain the content from a peer over the presence-based network; and
 - when the attempt to resolve the resource locator to obtain the content from the peer fails, the application is configured to resolve the resource locator to obtain the content from the remote server.
15. The computing device according to claim 14, wherein when the attempt to resolve the resource locator to obtain the content from a peer is successful, the application satisfies the request the peer-to-peer connection on the presence-based network by receiving the content from a remote device,

wherein the remote device is associated with an individual who appears on a buddy list of a user of the computing device.

16. The computing device according to claim 12, wherein: the application establishes a user of the computing device as currently having a presence in the presence-based network; and

the user is one of a plurality of participants in the presence-based network, each capable of having a presence in the presence-based network monitored by one or more servers coordinating the presence-based network for determining participants currently having a presence in the presence-based network.

17. The computing device according to claim 16, wherein when a first participant of the presence-based network associated with the endpoint does not have a current presence in the presence-based network, the application attempts to resolve the resource locator through at least one of:

the presence-based network at a device of a second participant that is a member of a group of which the first participant and the user of the computing device are also members; or

the remote server without peer-to-peer interaction.

18. One or more computer-readable media maintaining computer-readable instructions which, when executed, implement operations comprising:

generating, on a local device, a request for information, wherein the request is represented as a resource locator comprising an endpoint associated with user-published content;

attempting to resolve the resource locator through a presence-based network by determining whether a participant of the presence-based network referred to by the resource locator has a current presence on the presence-based network, wherein the presence-based network provides presence information regarding which participants currently have a presence; and

when the participant of the presence-based network referred to by the resource locator has a current presence on the presence-based network, obtaining the content from the participant through the presence-based network.

19. The one or more computer-readable media of claim 18, the operations further comprising receiving, as a result of attempting to resolve the resource locator through the presence-based network, an invitation to install an application for joining the presence-based network.

20. The one or more computer-readable media of claim 18, the operations further comprising, when the participant of the presence-based network referred to by the resource locator does not have a current presence on the presence-based network, obtaining the content from a remote server that does not use a peer-to-peer connection.

* * * * *