A META-SEARCH WEB SERVICE-BASED ARCHITECTURE FOR PEER-TO-PEER COLLABORATION AND VOICE-OVER-IP

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to the field of networking, and, more particularly, to discovering end-user devices and establishing a collaboration application session between/among the end-user devices.

2. Description of the Related Art

P2P is a communications model in which two or more parties (or "peers") have the same communications capabilities and any party can initiate a communications session. This differs from other communications models, such as the "client/server" model and the "master/slave" model, in which different nodes may have different communications capabilities. P2P communications may be implemented by giving each communications node both server and client capabilities. In recent years
, the term "P2P" has generally come to describe applications in which a group of users can use a network (*e.g.*, the Internet) to directly exchange files through at least one mediating server. In most P2P models, the mediating server is effectively "hidden" from the end-user; thus, the end-user is led to believe that a "direct" connection with another end-user is made.

A P2P network is generally a type of transient network that allows a group of computer users to connect with each other and directly discover and/or transfer files stored in each other's computers (*e.g.*, stored in a hard drive). The P2P network is generally created through the use of P2P networking software such as Kazaa (now commonly called FastTrack), Napster and Gnutella.

The initial growth of P2P networks was primarily due to its ability to directly share multimedia (*e.g.*, music, movies). However, unresolved legal issues associated with the sharing of protected works may reduce the use of P2P networks for such purpose. Therefore, it would be advantageous to use

P2P networks in other ways, for example, in gaming and business applications.

SUMMARY OF THE INVENTION

In one aspect of the present invention, a system for discovering potential devices on a peer-to peer (P2P) network is provided. The system includes a seeker device and a plurality of potential devices operatively connected to the P2P network. Each of the plurality of potential devices is associated with one or more identity files. Each of the identity files comprising a plurality of searchable elements. One or more of the plurality of potential end-user devices post their one or more identity files on the P2P network. The seeker device searches the P2P network to discover one or more of the plurality of potential devices based on the one or more identity files of the plurality of the potential devices. The seeker device initiates a collaboration session with the one or more potential devices.

In another aspect of the present invention, a method for a seeker device discovering potential collaborators on a peer-to peer (P2P) network is provided. The method includes the steps of discovering one or more entry point nodes to the P2P network, registering a seeker device on the P2P network, performing identity provisioning on a P2P network, performing one or more searches on the P2P network, obtaining one or more search results for potential collaborators on the P2P network, and initiating at least one of an application and a service to form a collaboration session with one or more potential collaborators from the search results.

In yet another aspect of the present invention, a method for a seeker device discovering potential collaborators on a peer-to peer (P2P) network is provided. The method includes the steps of registering with a P2P network, initiating a Web service to a Web service provider, requesting an available P2P server on the P2P network from the Web service provider using the Web service, registering the available P2P server in a Web service cluster using the Web service, performing identity self-provisioning on the P2P network, obtaining one or more search results searching for a potential collaborator on the P2P network, obtaining service and identity availability for each search

result, narrowing the number of search results to generate a narrowed result list, and initiating a collaboration session with one or more potential collaborators on the narrowed result list.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention may be understood by reference to the following description taken in conjunction with the accompanying drawings, in which like reference numerals identify like elements, and in which:

Figure 1 depicts a flow diagram of a method for a seeker device discovering potential collaborators on a peer-to peer (P2P) network, in accordance with one embodiment of the present invention;

Figures 2(a) and 2(b) depicts XML code for an exemplary registration process, in accordance with one illustrative embodiment of the present invention;

Figure 3 depicts XML code for an exemplary response to a request action, in accordance with one illustrative embodiment of the present invention;

Figure 4 depicts a process of self-provisioning, in accordance with one illustrative embodiment of the present invention;

Figure 5(a) depicts an exemplary process for a Web service based remote query of P2P devices, in accordance with one illustrative embodiment of the present invention;

Figure 5(b) depicts an alternate exemplary process for a Web service based remote query of P2P devices, in accordance with one illustrative embodiment of the present invention;

Figure 5(c) depicts an exemplary process for a Web service based remote query of call center agents, in accordance with one illustrative embodiment of the present invention;

Figure 6(a) depicts XML code of an exemplary search form, in accordance with one illustrative embodiment of the present invention;

Figure 6(b) depicts XML code of another exemplary search form, in accordance with one illustrative embodiment of the present invention;

Figure 6(c) depicts XML code of yet another exemplary search form, in accordance with one illustrative embodiment of the present invention;

Figure 7 depicts an exemplary identity and an exemplary search form with search criteria, in accordance with one illustrative embodiment of the present invention; and

Figure 8 depicts a diagram illustrating connectivity between a super peer P2P network and/or a P2P entry point, in accordance with one illustrative embodiment of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Illustrative embodiments of the invention are described below. In the interest of clarity, not all features of an actual implementation are described in this specification. It will of course be appreciated that in the development of any such actual embodiment, numerous implementation-specific decisions must be made to achieve the developers' specific goals, such as compliance with system-related and business-related constraints, which will vary from one implementation to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking for those of ordinary skill in the art having the benefit of this disclosure.

While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof have been shown by way of example in the drawings and are herein described in detail. It should be understood, however, that the description herein of specific embodiments is not intended to limit the invention to the particular forms disclosed, but on the contrary, the intention is to cover all modifications, equivalents , and alternatives falling within the scope of the invention as defined by the appended claims.

It is to be understood that the systems and method described herein may be implemented in various forms of hardware, software, firmware, special purpose processors, or a combination thereof. In particular, at least a portion of the present invention is preferably implemented as an application comprising program instructions that are tangibly embodied on one or more

program storage devices (*e.g.*, hard disk, magnetic floppy disk, RAM, ROM, CD ROM, etc.) and executable by any device or machine comprising suitable architecture, such as a general purpose digital computer having a processor, memory, and input/output interfaces. It is to be further understood that, because some of the constituent system components and process steps depicted in the accompanying Figures are preferably implemented in software, the connections between system modules (or the logic flow of method steps) may differ depending upon the manner in which the present invention is programmed. Given the teachers herein, one of ordinary skill in the related art will be able to contemplate these and similar implementations of the present invention.

The present invention allows a collaboration seeker to locate and query potential collaborators in a public (*e.g.*, in an Airport-lounge, coffee shop) or in a private/fixed wireline or wireless home environment (*e.g.*, using ADSL broadband fixed or wireless LAN DSL access routers).

The present invention leverages peer-to-peer ("P2P") networks created by one of a variety of known P2P networking applications/software (*e.g.*, Kazaa, OpenNAP, Gnutella, FastTrack, LimeWire, etc.) such that a collaboration seeker can search for potential collaborators from a collaboration pool. The collaboration seeker may search for the potential collaborators using user-defined search criteria. The term "collaboration pool," as used herein, refers to a group of people, machines, or devices participating in the P2P network. The term "potential collaborators," as used herein, refers to a collaboration pool that has been refined using the user-defined search criteria. Thus, the number of potential collaborators can be no greater than that of the collaboration pool. Once one or more potential collaborators have been identified, the collaboration seeker may establish a collaboration session with the one or more potential collaborators. Collaboration sessions include, but are not limited to, applications and/services, such as voice-over-IP (VoIP) telephony, multi-player gaming, multi-person business applications, and machine-to-machine applications. Collaboration sessions are usually initiated using a collaboration application. Collaboration sessions generally involve interaction between two or more

instances of the same application, or equivalent applications (*e.g.*, VoIP software phones).

Each collaborator in the collaboration pool may be connected to the P2P network using one or more end-user devices. End-user devices may include, but are not limited to, personal digital assistants ("PDAs"), laptop computers with wireless networking capabilities, and mobile phones (*e.g.*, a Smartphone). It is understood that the primary purpose of the P2P network in the present invention is to provide a means by which to identify end-user devices, and, thus, potential collaborators. Once the user has identified one or more potential collaborators from the collaboration pool, the collaboration seeker can obtain the IP address of the one or more potential collaborators and commence a collaboration session.

It should be understood that knowledge of the IP address allows the user to directly contact the potential collaborator with or without use of the P2P networking software. For example, a collaboration session may be made between two or more collaborators over the Internet or a local area network ("LAN").

The present invention can refine and reduce the number of collaborators returned from the collaboration pool by querying the potential collaborators' devices. The collaboration seeker initiates a search based on a search criteria. This may be implemented on, for example, a P2P collaboration application in which the search criteria may be present on a customized search form. If not previously cached, the collaboration seeker can download the search form from, for example, a Web Service Provider ("WSP"). The WSP may be fee-based and separate from an Internet service provider ("ISP"). The P2P collaboration application may automatically search the P2P network and return a result set of potential collaborators.

The majority of the popular existing P2P networks have demonstrated their ability to scale globally and to support significant user populations. In addition, some of these P2P networks are created using widely available open-source software, such as OpenNAP. By utilizing existing P2P network technologies, a number of advantages are automatically inherited. Among these are that the network will be "self-provisioning," which removes a

significant burden from the ISP. The term "self-provisioning," as used herein, refers to the ability of an end-user device to be a client-server (*i.e.*, having "servant" capability).

It is understood that the present invention is not limited to any particular type of P2P network or P2P networking software. In one embodiment, a private P2P network (*e.g.*, using an open-source implementation) may be used. In an alternate embodiment, one or more existing public P2P networks may be "piggybacked." In yet another embodiment, a hybrid approach may be used utilizing both a private and public P2P network. Further, there have been various models proposed for P2P networks, such as serverless, super-node server and hierarchical server architectures, each having its respective benefits.

The present invention builds upon the benefits of a P2P network by offering an application/service-independent framework in strong relation with any of a variety of known Web services and Web service architectures ("WSAs"), such as the one defined by the World Wide Web Consortium®. Web services provide a means of interoperating between different software applications, running on a variety of platforms and/or frameworks. This generic approach enables the framework to be used as a vehicle for offering and deploying multiple applications and services -- not just a single, dedicated service as with the vast majority of known P2P applications.

The present invention may use a meta-search web service-based architecture. In a meta-search engine, for example, a user submits keywords in a search box. The engine transmits the search to several individual search engines. Results are acquired from all the search engines queried. Meta-search engines generally do not own a database of Web pages. Instead, the engines send search terms to databases maintained by search engine companies.

Referring now to Figure 1, a flow diagram 100 illustrating a method for discovering potential collaborators for a collaboration session on a P2P network is shown, in accordance with one embodiment of the present invention. An end-user device (of the collaboration seeker) enters (at 105) a

hotspot area (*i.e.*, a geographic area covered by a wireless network, for example, the Internet), or connects to any IP network using any other means. The end-user device may be any of a variety of portable electronic devices capable of connecting to the hotspot, such as a PDA, a laptop, or a mobile phone.

The end-user device registers (at 110) with a P2P network, which allows the end-user device to utilize the P2P network. This may be done manually by the collaboration seeker or automatically once the end-user device enters (at 105) the hotspot, or connects to any IP network any other means. This registration process may be performed using a Web service and SOAP (Simple Object Access Protocol) protocols over HTTP (HyperText Transfer Protocol), FTP (File Transfer Protocol), etc. over IP (Internet Protocol) networks. For the sake of simplicity, our illustrative embodiments primarily utilize SOAP protocols used over HTTP. However, it should be understood that any of a variety of protocols may be utilized as is contemplated by those skilled in the art. An exemplary registration process is shown in Figure 2(a) and 2(b). Figure 2(a) shows a Web service request by an end-user device to use the P2P network. Figure 2(b) shows a response by the P2P acknowledging receipt of the Web service request.

Referring again to Figure 1, the end-user device initiates (at 115) a Web service using HTTP/XML (eXtensible Markup Language)/SOAP protocols to a known Web service provider ("WSP"). This initiation step may include discovering the WSP. The WSP may be discovered using a UDDI (Universal Description, Discovery and Integration) Web service registry and business entries. It is understood that the discovery of the WSP may be optional (*i.e.*, if knowledge of a particular WSP is hard-coded in the end-user device).

The end-user device requests (at 120) an available P2P server. In one embodiment, the end-user device sends a "request action" using the Web service to the WSP. The Web service may be capable of performing any number of request actions, or receiving any number of responses, actions, or methods. Request actions and responses may be defined in a WSDL service descriptor file. The WSDL service descriptor file is typically published by the

WSP into a public UDDI business service registration server. Examples of request actions include "requestServerReport" and "getServerReportResult," as shown in Figures 2 and 3, respectively. Figure 3 illustrates a response by the WSP showing a list a of SOAP service protocol elements reporting an available P2P server, the port to be used, and the number of users connected on the port. This information may be sent to the end-user device using XML/SOAP protocol elements. As shown in Figure 3, the command "getServerReportResult" returns a list of one or more available P2P network servers. In an alternate embodiment, the command "getServerReportResult" may also return a particular P2P network type (*e.g.*, OpenNap, Kazaa, Gnutella, eMule, Kademlia, or Limewire), including an "entry server" and a P2P access port to be used. The "entry server" is primarily used for registering for the first time on the particular P2P network. As shown in Figure 3, the end-user device supports only one P2P network (*i.e.*, OpenNap using port 8888). It is understood, however, that the end-user device may support any number of P2P network protocols. It is also understood that the request actions, "requestServerReport" and "getServerReport," are shown only for illustrative purposes, and other command method and functions may be implemented, as is contemplated by those skilled in the art.

Referring again to Figure 1, the end-user device registers (at 125) the P2P server into the Web service cluster (*i.e.*, a group of two or more computers implementing Web services). The registration process connecting and disconnecting one or more P2P servers into the Web service cluster may be performed using standard middleware architecture platform required to offer Web services, as is known to those skilled in the art. By using any professional J2EE (Java 2 Platform, Enterprise Edition) Web service architecture (*e.g.*, WebSphere® from IBM, WebLogic® from BEA, JBOSS), P2P protocols can be used to send statistics and other P2P protocol elements to a J2EE Web service adapter ("J2EE Adapter"). The J2EE Adapter is a composition of a "Connection" and a "Resource Adapter" according to the J2EE Connector Architecture Specifications, JSR112 V1.5. The J2EE Adapter may automatically register a P2P "entry" server as a legacy Enterprise Information System ("EIS") (*e.g.*, IBM® CICS transaction system,

SAP/R3 application) on the Web service cluster. This architecture allows key service features, such as (a) load balancing between P2P servers, (b) redundancy, (c) database synchronization, (d) proprietary EIS messaging translation in XML Web services messages, and (e) multiple EIS adapters for each additional supported P2P protocol.

The end-user device performs (at 130) identity self-provisioning. As previously stated, the term "self-provisioning" refers to the ability of the end-user device to be a client-server. The process of "identity self-provisioning" provides an identity of the end-user via the end-user device. The identity may be as simple or enhanced (*i.e.*, complex) as the user or WSP desires. A WSP may provide a simple or enhanced template/form for entering identity data. In one embodiment, the simple form may be embedded in or automatically downloaded to the end-user device as a default form, and the enhanced forms may be queried and downloaded by the end-user, as he so desires. Figure 4(a) illustrates the case where an identity form is embedded as part of the end-user device. The collaboration seeker enters data into the identity form. The identity is posted to the P2P network and a post reply is returned. Figure 4(b) illustrates the alternate case where the end-user device queries the WSP for the identity form. A Web service request is sent to the WSP and the WSP returns the identity form to the end-user device. The end-user enters data into the identity form. The end-user device posts the identity to the P2P network and a post reply is returned.

For security reason, forms completed by the end-user may be checked against a reference template hosted by, for example, the WSP. This avoids incorrectly formatted files and provides a more simplified and secure search process, which is described in greater detail below. Further, it is understood the type and content of information published in the P2P network is entirely up to the end-user. Thus, the end-user may choose to not populate all possible fields in an identity form.

Referring now to Figure 5(a), a first exemplary search process is shown, in accordance with one embodiment of the present invention. A service provider 505 is responsible for delivering "service descriptors" on-demand to devices which require a particular search. A user device 510

desires from the service provider 505 a way to get access to search information and the associated search process. Therefore, the user device 510 requests a so-called "Web service descriptor file" (hereinafter "WSDL file").

If the service provider 505 has the WSDL file 515 which complies with the request of the user device 510, the WSDL file 515 is transmitted to the user device 510. Based on the transmitted WSDL file 515, a first SOAP client 520 can initiate a search process by transmitting a protocol element <GetCity> in a query 525. The query 525 may not contain any city name. The protocol element <GetCity> is received by a first SOAP server 530 on a queried device 535A. A second SOAP client 540 responds with a protocol element <CityResponse>, which includes the name of the city that was registered on the queried device 535A. The aforementioned query 525 may be performed on other devices, such as 535B and 535C.

Referring now to Figure 5(b), a second exemplary search process is shown, in accordance with one embodiment of the present invention. The process of Figure 5(b) is similar to that of Figure 5(a), except the requested information in the query 525 is not a city, but a real-time information, such as "business availability" or "leisure availability" of a particular device.

According to a non-predictable/non-deterministic method, the queried device 535A can change the protocol element sent in a real-time fashion. The determination of the "device availability" may be determined by any means, such as local user manipulation or real-time and automatic determination.

Referring now to Figure 5(c), a third exemplary search process is shown, in accordance with one embodiment of the present invention. The process of Figure 5(c) is similar to that of Figure 5(b), except that Figure 5(c) is applied to an unknown end-user (*e.g.*, here a call center agent) 535A. In the case, the transmitted protocol element refers, in real-time, to whether a call agent 535A who is "on duty" is able to respond to a call request. In this example, the call agent 535A may respond that it is "on duty" and has certain "skills" registered on it. However, the transmitted protocol element may be repeated by other systems or processes that may not be sent by the end-user itself.

Figure 6(a) illustrates an exemplary simple personal identity form using XML. The form stores (a) a first name, (b) a last name, (c) an email, (d) a country, and (e) a gender. The file is saved on either (a) a P2P shared directory on the end-user device, or (b) on a distributed Hash Table "stored" on the P2P network itself, such that other users can search for the end user "John Smith" based on the criteria set forth in the identity form.

Figure 6(b) illustrates an exemplary business identity form for the same user, John Smith, as in Figure 6(b). The business identity form contains entries for the business type (*i.e.*, a car repair shop), the location of the business, and a registration or license number. The business identity form may be stored in the same P2P directory as the personal identity form.

Figure 6(c) illustrates an exemplary enhanced identity form. As is illustrated, the enhanced identity provides more information than the simple identity form of Figure 5(a). It is understood that the identity forms of Figures 5(a)-5(c) are shown only for illustrative purposes, and other form types and elements may be used, as is contemplated by those skilled in the art.

Referring again to Figure 1, the end-user device searches (at 135) for a potential collaborator from the collaboration pool. A simple search process will now be described. It is appreciated that the simple search process is shown only for illustrative purposes, and other search processes may be contemplated by those skilled in the art.

The simple "search process" may comprise at least one of the following two basic phases: (1) a P2P filename search; and (2) P2P file parse. In one embodiment, the search process may be performed using open protocols and international IT standards:

a) by getting a Web service (WSDL file/service descriptor) from a UDDI business service provider;
b) by implementing the Web service on the device according to the WSDL descriptor; and
c) by parsing (using SOAP protocols and XML grammar) each of the returned files using Web service descriptors available from step b above.

The P2P file search phase involves narrowing the identity files using the filename. The identity files may also be referred to as "user profiles." As shown in Figure 5(a), the filename of the simple personality identity form includes the first and last name of the end-user. Thus, a search for a particular end-user named "John Smith" may be performed by searching for the file <John.Smith.xml>. Other information may be searched in the filename as is contemplated by those skilled in the art (*e.g.*, <Car.Mechanics.Repair.xml>).

Referring now to Figure 7, assume that a collaboration seeker, John, wants to search for a person named "Marie Joe" on the P2P network. The device may initiate a search on the appropriate P2P network for all files published and defined as <Marie.Joe.xml>.

Each queried P2P server may respond with one or more registered <Marie.Joe.xml> files, including each file's respective IP address. It is understood that a response from an end-user device with a file indicates that the end-user device is online in the P2P network.

If more than one file is returned, the end-user would have to check each of the potential partners manually, which can become a tedious and time-consuming process. Thus, it may be advantageous to further perform a P2P file parse. In the example above, the collaboration seeker may input additional search criteria such that the number of <Marie.Joe.xml> files can be narrowed to down to one. The types of additional search criteria (*i.e.*, "masks") may be limited to a search form provided by the WSP or a third party to the collaboration seeker. Any number of masks can be used, including, but not limited to, e-mail address, city, country, gender, street, business, specialty, birthday, business address, etc.

The list above is meant to be only illustrative and not exhaustive. The number and types of masks depends largely on the WSP. The input fields offered on an identity form will generally fit a specific WSP search operation. For example, if the search field "Business" is offered in an identity form, then a querying SOAP action <GetBusiness> may be offered as a possible parsing criteria. Input masks may be standard XML structured files and can be

downloaded (or distributed) to end-users devices using Web services protocols and SOAP attachments standards.

Referring again to Figure 1, the end-user device obtains (at 140) the service or identity availability for each search result, which, in the example above, are "Marie.Joe.xml" files, using SOAP/XML standard protocols on top of HTTP transport protocols. This generally allows the end-user device to query files across network address translators ("NATs") and corporate firewalls, as HTTP is mostly used on Internet port 80 and cannot be blocked unless Internet access is blocked completely. It is understood that many other protocols (or mechanisms using other protocols in combination with HTTP, such as UDP) may be used as is known to those skilled in the art. Many of those are publicly available, and some allow the definition of a preferred way to exchange information across firewalls. To avoid downloading every <Marie.Joe.xml> file on the P2P network and performing an off-line parsing of the files, the end-user device may use an online remote querying using SOAP protocols.

The end-user device narrows (at 145) the number of search results using one or more search criteria. In the example above, the search can be refined by querying/searching, for example, the <city> listed for the potential collaborator.

A search agent (called, *e.g.*, NapGear) uses the <city> name to perform a SOAP action, <GetCity>, to all end-devices that published the <Marie.Joe.xml> file in the P2P network. Assuming they are turned on, each peer device will respond using a SOAP Response, for example, <CityResponse> (*e.g.*, "New York"). If a peer device is turned off, a timeout will occur after an allotted time has passed.

As a result of a SOAP response (or timeout in case of no response) devices and/or files with non-appropriate responses are discarded from each device, for example, because "New York" doesn't fit to the requested location provided in the search form (*i.e.*, "Los Angeles"). If additional search refinement is necessary, the search results may be further narrowed with additional Web service requests using additional search criteria.

Once an end-user has found one or more potential collaborators, it is appreciated, that the end-user may add the one or more potential collaborators to a directory on the end-user device. The directory may be implemented in any of a variety of ways, such as a buddy list, a chat list, a user group, etc. This step, although optional, allows the end-user to seek a known potential collaborator without performing the search step, as described in greater detail above, again. The directory may be informed by the P2P network of a potential collaborator's availability and may automatically update the potential collaborator's IP address.

The end-user device initiates (at 150) a collaboration session using any of a variety of collaboration applications (*e.g.*, VoIP, online gaming application). The end-user device generally selects and initiates the particular collaboration application and the one or more potential collaborators follow suit. In one embodiment, the collaboration application is initiated on a direct P2P connection using a known IP address. For example, the end-user device may initiate VoIP software (*e.g.*, SJPhone®) using an IP-address (for H.323) or an email address (for SIP). The IP address of the potential collaborator may be provided directly by the P2P network or the XML parsing process. This IP address is used as a command line parameter when starting the collaboration application on the end user's device. In an alternate embodiment, the collaboration application may be initiated remotely using Web services between the end-user device and the potential collaborator.

In the case of gaming applications, a similar procedure applies. Search criteria may be different and may not require knowledge of the potential collaborators. Search criteria for gaming applications may include, but is not limited to, "preferred gaming application" and "highest score." Such criteria may be sufficient to engage an online game or virtual competition between two or more known or unknown users.

The end-user device leaves (at 155) the hotspot area or any other IP access point or network access, and removes the posted files from the P2P network. If the end-user device cannot remove its identity from the P2P network, the P2P nodes can discard the user and his identity files after a determined period of time (*i.e.*, a timeout).

Mobile devices can become part of the P2P server architecture themselves. For example, simple P2P servers may report to a "super peer" P2P server. This P2P architecture improves network robustness and enables faster searches. Super peers are well known to those skilled in the art and are extensively used in popular P2P networks.

Figure 8 illustrates how simple mobile devices connect to mobile P2P servers, which, in turn, can connect to a super peer P2P server and/or a P2P network entry point, in accordance with one embodiment of the present invention. In this example, even if two simple end-user devices are now connected together, they can discover each other using a simple request to the super peer. By letting simple end-user devices connect to two or more P2P embedded servers, we greatly improve the network redundancy. As shown in Figure 7, if server A fails (eg. loses its wireless connection) all devices connected can be reached by server B which is still connected to the super peer server.

It should be understood that the forgoing description of the present invention should not be limited to "end-user" or personal devices. For example, we can also utilize any of a variety machines connected to P2P network, such that the machines are capable of performing automatic searches. Machines can perform searches by (a) manually gathering information from an operator, or (b) automatically gathering information in real-time (eg. a captured measurement unit like temperature or pressure).

The particular embodiments disclosed above are illustrative only, as the invention may be modified and practiced in different but equivalent manners apparent to those skilled in the art having the benefit of the teachings herein. Furthermore, no limitations are intended to the details of construction or design herein shown, other than as described in the claims below. It is therefore evident that the particular embodiments disclosed above may be altered or modified and all such variations are considered within the scope of the invention. Accordingly, the protection sought herein is as set forth in the claims below.

WHAT IS CLAIMED IS:

1. A system which enables discovery of potential collaborators on a peer-to-peer (P2P) network, comprising:

a seeker device; and

a plurality of potential collaborators operatively connected to a P2P network;

wherein the seeker device comprises:

means for registering with a P2P network;

means for initiating a Web service to a Web service provider;

means for requesting an available P2P server on the P2P network from the Web service provider using the Web service;

means for registering the available P2P server in a Web service cluster using the Web service;

means for performing identity self-provisioning on the P2P network;

means for obtaining one or more search results searching for a potential collaborator on the P2P network;

means for obtaining service and identity availability for each search result;

means for narrowing the number of search results to generate a narrowed result list; and

means for initiating a collaboration session with one or more potential collaborators on the narrowed result list.

2. The system of claim 1, wherein the seeker device is a seeker end-user device and the plurality of potential devices are a plurality of potential end-user devices.

3. The system of claim 2, wherein each of the seeker end-user device and the plurality of end-user devices comprises at least one of a personal digital assistant, a laptop, and a cellular phone.

4. The system of claim 1, wherein each of the one or more identity filed of the plurality of the potential devices is downloaded from a Web service provider in response to the seeker device sending a Web service request to the Web service provider.

5. The system of claim 1, wherein the seeker end-user device logs on a Web service provider to gain access to the P2P network using Web services and simple-object access protocols (SOAP) over hypertext transfer protocol (HTTP) and internet protocol (IP) networks.

6. The system of claim 1, wherein the seeker device is a machine connected to an IP network.

7. The system of claim 1, wherein the P2P network comprises at least one of Kazaam OpenNAP, Gnutella, FastTrack, LimeWire, eMule/Kademila, and Napster.

8. The system of claim 1, wherein each of the one or more identity files of the plurality of the potential devices comprises an extensible markup language (XML) file.

9. The system of claim 1, wherein the collaboration session is independent of the P2P network.

10. A method for a seeker device discovering potential collaborators on a peer-to-peer (P2P) network, comprising:

registering with a P2P network;

initiating a Web service to a Web service provider;

requesting an available P2P server on the P2P network from the Web service provider using the Web service;

registering the available P2P server in a Web service cluster using the Web service;

performing identity self-provisioning on the P2P network;

obtaining one or more search results searching for a potential collaborator on the P2P network;

obtaining service and identity availability for each search result;

narrowing the number of search results to generate a narrowed result list; and

initiating a collaboration session with one or more potential collaborators on the narrowed result list.

11. The method of claim 10, wherein registering with a P2P network comprises registering automatically with a P2P network when the seeker end-user device connects to an IP network.

12. The method of claim 10, wherein initiating a Web service to a Web service provider comprises initiating a Web service to a Web service provider using HTTP/CML/SOAP protocols.

13. The method of claim 10, further comprising discovering the Web service provider using a UDDI Web service registry and business entities.

14. The method of claim 10, wherein requesting an available P2P server on the P2P network from the Web service provider using the Web service comprises sending a Web service request using a Web service to the Web service provider, the Web service request requesting a list of available P2P servers.

15. The method of claim 14, wherein sending a Web service request using a Web service to the Web service provider comprises sending a Web service request defined in a WSDL service descriptor file using a Web service to the Web service provider.

16. The method of claim 10, wherein performing identity self-provisioning on the P2P network comprises:

receiving an identity form from the Web service provider in response to a Web service request, the identity form comprises a plurality of information fields;

populating one or more of the plurality of information fields; and

posting the identity form on the P2P network;

17. The method of claim 10, wherein obtaining one or more search results searching for a potential collaborator on a P2P network comprises performing a P2P filename search.

18. The method of claim 10, wherein narrowing the number of search results to generate a narrowed result list comprises:

downloading a search form from the Web service provider in response to a Web service request, the search form comprising a plurality of search fields;

populating one or more of the plurality of search fields;

narrowing the one or more search results based on the one or more of the plurality of search fields; and

storing the results from the step of narrow in the narrowed result list.

19. The method of claim 10, wherein initiating a collaboration session with one or more potential collaborators on the narrowed result list comprises with one or more potential collaborators on the narrowed result list independent of the P2P network.

20. A machine-readable medium having instructions stored thereon for execution by a processor to perform a method for a seeker device discovering potential collaborators on a peer-to-peer (P2P) network, comprising:

registering with a P2P network;

initiating a Web service to a Web service provider;

requesting an available P2P server on the P2P network from the Web service provider using the Web service;

registering the available P2P server in a Web service cluster using the Web service;

performing identity self-provisioning on the P2P network;

obtaining one or more search results searching for a potential collaborator on the P2P network;

obtaining the service and identity availability for each search result;

narrowing the number of search results to generate a narrowed result list; and

initiating a collaboration session with one or more potential collaborators on the narrowed result list.