



- (51) **International Patent Classification:**
H04L 29/08 (2006.01)
- (21) **International Application Number:**
PCT/US2015/068284
- (22) **International Filing Date:**
31 December 2015 (31.12.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/098,930 31 December 2014 (31.12.2014) US
- (71) **Applicant:** LEVEL 3 COMMUNICATIONS, LLC [US/US]; 1025 Eldorado Boulevard, Broomfield, Colorado 80021 (US).
- (72) **Inventors:** NEWTON, Christopher; 1025 Eldorado Boulevard, Broomfield, Colorado 80021 (US). POWER, William R.; 1025 Eldorado Boulevard, Broomfield, Colorado 80021 (US).
- (74) **Agents:** DURBIN, Gregory P. et al.; Polsinelli PC, 900 W. 48th Place, Suite 900, Kansas City, Missouri 64112 (US).

- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:
— with international search report (Art. 21(3))

(54) **Title:** NETWORK ADDRESS RESOLUTION




FIG. 1

(57) **Abstract:** A content delivery method, and related apparatus, that involves the operations of receiving a request for a content resource including at least one embedded resource with a tag. Upon identification of the tag, using at least one delivery parameter to generate a modified embedded resource. Embedding the modified embedded resource in a content resource, such as an HTML document. Delivering, such as in response to request from a browser, the content resource with the modified embedded resource.

WO 2016/109797 A1

NETWORK ADDRESS RESOLUTION

RELATED APPLICATIONS

[0001] This Patent Cooperation Treaty (PCT) application claims priority from U.S. provisional application no.62/098,930 entitled "NETWORK ADDRESS RESOLUTION," filed on December 31, 2014, the entire contents of which are fully incorporated by reference herein for all purposes.

TECHNICAL FIELD

[0002] Aspects of the present disclosure relate to network address resolution, and particularly involve introduction of an application programming interface into the resolution stream to embed link resolution information within an HTML document or other content returned to a requesting client.

BACKGROUND

[0003] The Internet and the World Wide Web (the "Web") have become ubiquitous. Thousands or even tens of thousands of content providers (publishers) now use the Internet (and, particularly, the Web) to provide all sorts of content to numerous clients all over the world. In order to offload the job of serving some or all of their content, many content providers now subscribe to content delivery networks (CDNs). Using a CDN, content can be served to clients from the CDN (e.g., from one or more servers in the CDN) instead of from the content provider's server(s). In a caching CDN, content may also be cached on some or all of the CDN servers, either before being served or in response to specific requests for that content. Having content cached within edge servers of the CDN enhances the performance of the CDN because the content does not have to be retrieved from mid-tier servers origin servers or other locations, which are less efficient than edge servers in providing content.

[0004] Content is constantly being accessed from networks. For example, television shows and movies may now be accessed from any number of Web sites. Print

newspapers have migrated to the Web and provide portals through which clients operating some form of computing device (e.g., PC, smart phone, or tablet), with a browser may access numerous forms of content, such as short video clips, articles, images, and audio tracks. Software updates and patches, once only provided on disc and mailed to recipients, are now routinely distributed to devices using only network connections. All of these forms of content may be distributed using CDNs.

[0005] It is with these observations in mind, among others, that various aspects of the present disclosure were conceived and developed.

SUMMARY

[0006] Embodiments of the present disclosure involve a content delivery method including the operations of receiving a request for a content resource including at least one embedded resource with a tag. Upon identification of the tag, using at least one delivery parameter to generate a modified embedded resource. And, delivering the content resource with the modified embedded resource.

[0007] Embodiments may also include a content delivery device comprising at least one processor in communication with a non-transitory computer readable medium including computer readable instructions for generating an application programming interface (API). At the API, receiving an embedded resource and at least one parameter associated with the embedded resource. The instructions may further include generating a modified embedded resource providing information as to a location within a content delivery network for obtaining the embedded resource.

[0008] Embodiments may also include a non-transitory computer readable media comprising an object document, such as an HTML document, comprising at least one embedded resource, the at least one embedded resource including a tag that causes the embedded resource to be analyzed to construct a modified embedded resource to replace the embedded resource, the modified embedded resource causing a browser receiving the document with the modified embedded resource to initiate communication

with a content delivery network to obtain content associated with the embedded resource.

[0009] Finally, embodiments may also include a content delivery method comprising initiating an HTTP connection with a server, the HTTP connection including a client IP address that is available to the server. Receiving an HTML document including at least one embedded resource, the embedded resource including a reference to a content delivery network whereby the reference to the content delivery network is influenced by the client IP address. And, using the reference, initiating a request directly to the content delivery network to obtain content associated with the at least one embedded resource.

[0010] These and other embodiments are discussed in more detail below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The foregoing and other objects, features, and advantages of the present disclosure set forth herein should be apparent from the following description of particular embodiments of those inventive concepts, as illustrated in the accompanying drawings. Also, in the drawings the like reference characters may refer to the same parts throughout the different views. The drawings depict only typical embodiments of the present disclosure and, therefore, are not to be considered limiting in scope.

[0012] Figure 1 is a network diagram and message flow related to generating an embedded resource for inclusion in a network asset, such as an HTML document, the embedded resource providing a connection to a node of a CDN bypassing DNS, in one example;

[0013] Figure 2 is a flow diagram illustrating a method of identifying a tagged resource and thereafter providing a modified embedded resource by passing various parameters to an API for identification of a node of a CDN able to serve content associated with the embedded resource;

[0014] Figure 3 is a flow diagram illustrating a method of generating a modified resource for embedding in a requested asset having a tagged resource;

[0015] Figure 4 is a network diagram and message flow related to generating an embedded resource from within a CDN for inclusion in a network asset, such as an HTML document, the embedded resource providing a connection to a node of the CDN bypassing DNS, in one example;

[0016] Figure 5 is a flow diagram illustrating a method of identifying a tagged resource and thereafter providing a modified embedded resource by passing various parameters to an API within a CDN, the API for identification of a node of a CDN able to serve content associated with the embedded resource;

[0017] Figure 6 is an example of a computing system that may implement various systems and methods discussed herein.

DETAILED DESCRIPTION

[0018] Aspects of the present disclosure involve systems and methods for selecting an edge server, rack, cluster or other device, which may be in a content delivery network (CDN), to serve content to a client. Aspects of the present disclosure involve an alternative to using conventional domain name system (DNS) servers for edge selection by instead calling an API or otherwise querying the CDN when constructing a uniform resource locator (URL) for the resource to be delivered from the CDN. The API call may include a name of the resource, a location of the client requesting the resource and/or an IP address for the client which may be used to determine the location of the client, and/or the popularity of the resource being requested (or embedded within a resource). The API or other computing engine associated therewith generates a target or list of targets (nodes capable of serving the content from the CDN), and provides one or more locations from which the content (or embedded content) may be served. In one particular form, the API may return one or more virtual IP addresses (VIPs) for one or more clustered edge servers. The API may also return the protocol to use for requesting the content. The API may also return a list of absolute URL's from which the

content may be retrieved, and/or a list of hostnames and/or IP addresses that could be used to construct the URLs at the calling application.

[0019] The term “content” as used herein means any kind of data, in any form, regardless of its representation and regardless of what it represents. The term “content” may include, without limitation, static and/or dynamic images, text, audio content, including streamed audio, video content, including streamed video, web pages, computer programs, documents, files, and the like. Some content may be embedded in other content, e.g., using markup languages such as HTML and XML. Content includes content which is created or formed or composed specifically in response to a particular request. The term “resource” is sometimes used herein to refer to content.

[0020] Figure 1 illustrates one example network configuration for practicing aspects of the present disclosure, while Figures 2 and 3 describe methods. The illustrated network would include many other components involved in servicing a request and providing content; however, the diagram here focuses on only some components involved in the system and one of ordinary skill will readily understand other conventional components within the network. With respect to content delivery, in this example a request originates from a computing device 10 operating a user agent 12, such as a browser. The computing device may be some form of computing device such as personal computer, a thin client, a tablet, a smart phone, a set-top box, a cable box, or the like, which is capable of running a browser and providing a user with the ability to request content over a network connection, or some other application suitable for interacting with a network and requesting content therefrom. The user device may be a consumer device deployed in a public network, may be an enterprise device deployed in a private network, or may be some other type of device. The network connection, in many instances, will be provided by an Internet service provider (ISP) operating one or more resolvers 14. A resolver (also referred to as a DNS resolver) is a client-side device used to resolve a network address (e.g., translate a domain name into an IP address) for a requested resource. Generally speaking, the resolver contacts (e.g., queries) a DNS infrastructure 16, which may involve many different components, to resolve a network address for a resource (operation 200).

[0021] To illustrate aspects of the disclosure, consider when a user enters, into the browser Web page, an address `http://www.example.com`. The browser 12 contacts the resolver 14, which in turn requests the internet protocol (IP) address for the Web page from DNS 16. DNS returns the IP address of an origin server 18 capable of delivering a hypertext markup language (HTML) document 20 that will render the Web page in the browser. In order to obtain the HTML document, the client device establishes a TCP connection, which may encapsulate an HTTP protocol communication, with the origin server 18 (operation 210). An origin server is a device (or collection of devices) maintaining original copies of content for distribution to client's and/or distribution of copies of content to other nodes. While the present disclosure is illustrated in association with a connection to an origin server, the connection may instead be made with some other device designated in DNS to serve the requested resource. In this example, the origin server is connected with a database 22, which organizes and manages stored resources in some form of storage device (or devices), such as a storage appliance, that includes the HTML document 20 with an embedded resource 24. The HTML document is what the browser uses to display the Web page `http://www.example.com/`. In this example, the embedded resource 24A is a link to some form of content that the user can obtain by selecting the link when the Web page is displayed.

[0022] The embedded resource may include a tag indicating that a resource address compute engine 30 should be contacted via a call to an API 26 (operation 220). For instance, the HTML document may include an embedded resource of the form:

```
http://${CALLAPI:ADDRONLY}/path/resource.txt, or  
${CALLAPI:FULLURL,path=$ENV{PATH},pop=$POP}.
```

[0023] In both instances, the tag is "\$" which causes the origin server 18 to call the API 26 to retrieve the address (only) or the complete URL. Hence, the presence of the tagged resource 24A in the HTML document 20 at the origin servers causes the origin server to call the API, the result of which causes the HTML document returned to the

requesting client to include a modified embedded resource 24B with an address or URL to the CDN whereby the content may be retrieved.

[0024] When the origin server receives the request (or at some time prior to receiving the request), the origin server calls an API in order to communicate with the compute engine 30 to format the address for the embedded resource. The presence of the tagged embedded resource prompts the API call. In this example, the address for the embedded resource is not static. Instead, an address that is optimally suited to serve the content to the requesting client is returned. So, in a simple example, the HTML document is initially published with a tagged embedded resource 24A that would identify the resource, and may also include other attributes such as popularity. If the user were in New York, which may be determined by including the requesting client's IP address or information, such as geographical information, based on the client IP address in the API call, and the CDN is capable of serving the requested content from both California and New York, it may not be optimal to serve the content from California due to any number of reasons including cost, latency, quality, optimal CDN use, and the like. In the present system, the API/compute engine may instead cause the creation of an embedded resource 24B (also referred to as a modified embedded resource) that will direct the client to the New York node. While described herein as an API call (or calls) in conjunction with an associated compute engine, other applications may be used to compute and return the modified embedded resource. In some instances, only an API is referenced but it should be recognized that such reference includes not only the interface but also the associated compute routine, function, and the like for processing inputs to the API and generating a modified embedded resource. In any event, a modified embedded resource is generated and substituted for the resource, where the modified resource identified a node able to serve the content (operation 230).

[0025] The API may obtain information used in determining the address from various sources. To begin, the request from the origin server may include information. Since the client has contacted the origin server over a TCP connection, the client IP address may be included in the request from the client. Accordingly, the origin server may transmit the client IP address to the API and/or the origin server may use the client IP to

obtain or generate information 28 to share with the API. For example, the origin server may obtain geographic information for the client, based on the client IP address, or may transmit the IP address to the API, which may obtain geographic information for the client. The location of the client may be used in selecting an optimal node to serve a request for the embedded resource. Note, by embedding an IP address to serve the content directly, the client device may bypass DNS and instead connect directly to the source of the embedded resource. Regardless, the HTML document 20 with the modified imbedded resource 24B is returned to the client device (operation 240).

[0026] The origin server, or other devices associated with the customer managing the origin server, may also obtain popularity information for the content. For example, consider the case where <http://www.example.com/> is the home page for an on-line video provider and the embedded resource is for a recently released movie that did well in theaters. In such an example, when the movie is released by the video provider, it may initially be quite popular, and this popularity may be tracked and shared with the API (and its compute engine). In the same example, the customer may preset the popularity to “high,” for example, based on its understanding that the initial release of the video will be highly popular among its subscribers (user base). The scale of popularity may be in any number of possible forms – 1 = low, 2 = medium, 3 = high, scale of 1-10 with 1 being unpopular and 10 being the most popular, etc. The customer defined popularity may be included in the tagged embedded resource of the HTML page as initially defined at the origin or may be provided as a separate input to the API. The scale may be based on the number of requests for the content, which may be based on discrete times (e.g., number of requests per day, per minute, per second, etc.). Finally, popularity may be predicted, and may also be regionalized. For example, a resource in the case of a news story may initially become popular in some particular region or at some particular time. Say, for example, an early morning national breaking news story may become actually popular on the East Coast of the United States. The customer may then predict its popularity on the West Coast and the system may automatically set it to a “high” value for all requests.

[0027] Besides client information, including geography, and popularity, the origin server may also pass other information such as policy information, political information, restrictions, etc. For example, the distribution of some content may be geographically limited – e.g., obtaining real-time sporting content may be limited in the geographic area of one or both teams participating in the event. Some content may have geopolitical distribution restrictions. This information may also be stored separately at the CDN or otherwise separately accessible by the CDN and passed to or otherwise used in generating the modified embedded resource. Regardless, referring to Fig. 3, the compute engine receives CDN information, as well as other information, it may use to identify a node, collection of nodes, VIP, or the like, capable of receiving the content (operation 300).

[0028] Similar information may be provided to the API or accessible by the API from other sources. The CDN, for example, may maintain its own popularity information and may track the same, which may account for requests originating with other customers of the CDN. It is often the case that a CDN will have numerous customers, and handle some or all content delivery for those customers. Further, in many instances, customers will have Web properties with the same content. Accordingly, clients may contact any of those Web properties to obtain the content. Hence, the CDN has a larger view of the popularity of some content.

[0029] In any event, the API 26 provides the interface to a computing engine 30 that uses the information, passed to the engine through the API, to determine one or more nodes 32 in the CDN to serve the content upon receiving a request from the origin server (or from some other request source) related to a tagged resource, or resource otherwise identified for modification (operation 310). The computing engine may take into account and use client IP, customer (through, for example a customer ID), popularity in some or all forms, the delivery protocol of the resource, and the various other forms of information in determining the node from which to serve the content (operation 320). For example, the computing engine may use geographical information to pick the node (CS1) closest to the client to serve the content. The computing engine may further take into account network traffic, and if the closest node is too busy, then

select another node (e.g., CS2) that is close but less busy. The computing engine may further embed an instruction, along with the IP address, for the node to cache a copy of the content should the content meet some popularity threshold, and likely be requested again. In some instances, for example, the node may not yet have the content, and upon receiving the request, have to request the content from an origin server, another node, or elsewhere. In such instances, the instruction would cause the node to cache and retain the content and be able to directly serve the next request.

[0030] The API returns an absolute URL from which the content may be retrieved, and/or a list of hostnames and/or IP addresses that could be used to construct the URLs at the calling application (operation 330). Similarly, the origin server may pass the document to the API, and the API causes the modified embedded resource to be directly embedded and return the modified document to the origin server. Thus, when the client browser displays the Web page, for example, the modified embedded resource (e.g., link) may include an absolute URL for a node (e.g., edge server(s)) within the CDN that can deliver the content. When the user selects the link, the browser will directly contact the node for the content. In the case of a hostname, the embedded resource may include specific reference to the CDN such that name servers 34 of the CDN are contacted directly, bypassing any other DNS processing such that the CDN may perform any form of intelligent traffic management (ITM) or other forms of edge selection and return an IP address of a node within the CDN able to process the request. It may be the case that the hostname will also include a reference to a node capable of serving the request as determined by the API. In such a case, the hostname references a CDN location as is closely tailored to a location for serving the request but allows for some level of fail over, load balancing, or the like, relative to an absolute URL. In such a case, the hostname may be mapped to a VIP associated with various machines capable of serving the request.

[0031] In an alternative implementation, which may also involve an API and calling the same in order to resolve an optimal node to serve a request, the process of contacting the API and constructing the optimal network address is instead handled by the CDN and does not necessarily involve a customer origin server or other devices

outside the CDN. While embodiments are discussed herein as involving an API, it should be noted that an API is a convenient and efficient way for disparate systems to uniformly and predictably pass information to the computing engine and for the computing engine to return information. However, other forms of communicating with and passing information to the computing engine are also possible. Moreover, the API may include the functionality of the engine and not necessarily be considered a discrete mechanism from the computing engine functions.

[0032] Referring now to Figs. 4 and 5, in the same manner as the system and process of Figs. 1-3, a client device 40 makes an initial request for a network resource using a browser 42, which request is received (operation 510). For example, the client requests `http://www.example.com`. The client's resolver 44 queries DNS 46, which returns a CNAME from an origin server – name service (OS-NS) 48 associated with the requested content (operation 520). A CNAME is a mechanism used in DNS to specify the name or address of another domain. In this case, the CNAME points to the domain of the CDN 50 – e.g., `www.example.com.cdn.net`. In such a case, the client device (or other accessing device) will seamlessly be directed to contact the name server 52 or servers for the `cdn.net` domain. The name server will resolve the request and return an entry point to the CDN, e.g., a node 54, from which to obtain the page 56. Regardless, however, of how the client is directed to the CDN, like the example above, the request to the CDN from the client device may involve an HTTP connection and may include the IP address of the requesting device (client IP) and information as to the requested page (operation 530). The page (e.g., `www.example.com`) has one or more embedded resources, which may be tagged resources 58A as discussed above. Hence, the CDN node (e.g., an edge server or the CDN origin server) contacted for the page will recognize the presence of a tagged embedded resource 58A and contact the API/compute node 60 to return an embedded resource (referred to as the “modified embedded resource”) 58B to include in the page returned to the requesting client device 40 (operations 540 and 550).

[0033] The API/compute engine will have access to various types of information from various possible sources. For example, as discussed above, the API may be passed

information and may receive information from other sources, which information can be used by the compute engine to generate an address for an embedded link within the page. The page and embedded link (e.g., link to resource) are returned to the requesting client. In one possible example, the embedded resource may include a VIP or a direct link to an edge server to deliver the resource. The compute engine may take into account the location of the client resolver, the location of the client (based on client IP), policy, popularity as tracked by the CDN or included in the original embedded resource, geography restrictions, protocol, and other parameters, when generating the modified embedded resource 58B, which information may be stored, computed or otherwise accessed from a CDN information store 62. Accordingly, when the page 56 (e.g., the HTML document) is returned to the client device 40, and the modified embedded resource 58B is selected, the client device (e.g., browser) may directly contact the node as defined by the API/compute engine or otherwise contact the CDN to obtain the resource.

[0034] Figure 6 is example schematic diagram of a computing system 600 implementing a CDN node compute engine (and API) 608 that may collect the various data discussed herein, receive requests from origin server, a CDN-OS, or otherwise and process the data and request to determine an optimal CDN node. The computing system for the application 608 includes a bus 601 (i.e., interconnect), at least one processor 602 or other compute element, at least one communication port 603, a main memory 604, a removable storage media 605, a read-only memory 606, and a mass storage device 607. Processor(s) 602 can be any known processor, such as, but not limited to, an Intel® Itanium® or Itanium 2® processor(s), AMD® Opteron® or Athlon MP® processor(s), or Motorola® lines of processors. Communication port 603 can be any of an RS-232 port for use with a modem based dial-up connection, a 10/100 Ethernet port, a Gigabit port using copper or fiber, or a USB port. Communication port(s) 603 may be chosen depending on a network such as a Local Area Network (LAN), a Wide Area Network (WAN), or any network to which the computer system 600 connects. The rendezvous application may be in communication with peripheral devices (e.g., display screen 630, input device 616 via Input/Output (I/O) port 609).

[0035] Main memory 604 can be Random Access Memory (RAM) or any other dynamic storage device(s) commonly known in the art. Read-only memory 606 can be any static storage device(s) such as Programmable Read-Only Memory (PROM) chips for storing static information such as instructions for processor 602. Mass storage device 607 can be used to store information and instructions. For example, hard disks such as the Adaptec® family of Small Computer Serial Interface (SCSI) drives, an optical disc, an array of disks such as Redundant Array of Independent Disks (RAID), such as the Adaptec® family of RAID drives, or any other mass storage devices, may be used.

[0036] Bus 601 communicatively couples processor(s) 602 with the other memory, storage and communications blocks. Bus 601 can be a PCI / PCI-X, SCSI, or Universal Serial Bus (USB) based system bus (or other) depending on the storage devices used. Removable storage media 605 can be any kind of external hard drives, thumb drives, Compact Disc – Read Only Memory (CD-ROM), Compact Disc – Re-Writable (CD-RW), Digital Video Disk – Read Only Memory (DVD-ROM), etc.

[0037] Embodiments herein may be provided as a computer program product, which may include a machine-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform a process. The machine-readable medium may include, but is not limited to, optical discs, CD-ROMs, magneto-optical disks, ROMs, RAMs, erasable programmable read-only memories (EPROMs), electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, flash memory, or other type of media/machine-readable medium suitable for storing electronic instructions

[0038] As shown, main memory 604 is encoded with an embedded resource application 650-1, which may include functionality described relative the API/compute engine discussed herein and that supports functionality as discussed above and as discussed further below. For example, in one embodiment, the application 650-1 may include or otherwise implement the various processes and/or instructions described herein. The application 650-1 (and/or other resources as described herein) can be

embodied as software code such as data and/or logic instructions (e.g., code stored in the memory or on another computer readable medium such as a disk) that supports processing functionality according to different embodiments described herein. During operation of one embodiment, processor(s) 602 accesses main memory 604 via the use of bus 601 in order to launch, run, execute, interpret or otherwise perform the logic instructions of the 650-1.

[0039] The description above includes example systems, methods, techniques, instruction sequences, and/or computer program products that embody techniques of the present disclosure. However, it is understood that the described disclosure may be practiced without these specific details. In the present disclosure, the methods disclosed may be implemented as sets of instructions or software readable by a device. Further, it is understood that the specific order or hierarchy of steps in the methods disclosed are instances of example approaches. Based upon design preferences, it is understood that the specific order or hierarchy of steps in the method can be rearranged while remaining within the disclosed subject matter. The accompanying method claims present elements of the various steps in a sample order, and are not necessarily meant to be limited to the specific order or hierarchy presented.

[0040] The described disclosure may be provided as a computer program product, or software, that may include a machine-readable medium having stored thereon instructions, which may be used to program a computer system (or other electronic devices) to perform a process according to the present disclosure. A machine-readable medium includes any mechanism for storing information in a form (e.g., software, processing application) readable by a machine (e.g., a computer). The machine-readable medium may include, but is not limited to, magnetic storage medium, optical storage medium (e.g., CD-ROM); magneto-optical storage medium, read only memory (ROM); random access memory (RAM); erasable programmable memory (e.g., EPROM and EEPROM); flash memory; or other types of medium suitable for storing electronic instructions.

[0041] It is believed that the present disclosure and many of its attendant advantages should be understood by the foregoing description, and it should be apparent that various changes may be made in the form, construction and arrangement of the components without departing from the disclosed subject matter or without sacrificing all of its material advantages. The form described is merely explanatory, and it is the intention of the following claims to encompass and include such changes.

[0042] While the present disclosure has been described with reference to various embodiments, it should be understood that these embodiments are illustrative and that the scope of the disclosure is not limited to them. Many variations, modifications, additions, and improvements are possible. More generally, embodiments in accordance with the present disclosure have been described in the context of particular implementations. Functionality may be separated or combined in blocks differently in various embodiments of the disclosure or described with different terminology. These and other variations, modifications, additions, and improvements may fall within the scope of the disclosure as defined in the claims that follow.

CLAIMS

WHAT IS CLAIMED IS:

1. A content delivery method comprising:
 - receiving a request for a content resource including at least one embedded resource with a tag;
 - upon identification of the tag, using at least one delivery parameter to generate a modified embedded resource; and
 - delivering the content resource with the modified embedded resource.
2. The content delivery method of claim 1 wherein the request is received over an http connection and the at least one delivery parameter includes a requesting client IP address, and further comprising:
 - passing information based on the client IP address to an application programming interface (API);
 - in association with the API, obtaining a parameter associated with the embedded resource; and
 - identifying, based on the requesting client IP address and the parameter, location information for at least one device in a content delivery network that can serve the embedded resource; and
 - generating the modified embedded resource to include reference to the at least one device.
3. The content delivery method of claim 1 wherein the tag is a unique tag.

4. The content delivery method of claim 2 wherein the parameter is a popularity parameter.

5. The content delivery method of claim 2 wherein the parameter is a delivery protocol parameter.

6. The content delivery method of claim 2 wherein the information is the requesting client IP address.

7. The content delivery method of claim 2 wherein the information is geographic information.

8. The content delivery method of claim 2 wherein the location information includes at least one of an IP address, a virtual IP address, one or more absolute uniform resource locators, and one or more hostnames.

9. The content delivery method of claim 1 wherein the request is over an http connection and the at least one delivery parameter is a client location identifier, the method further comprising:

passing the client location identifier to an application programming interface (API) to generate the modified embedded resource relating to a content delivery node to service the embedded resource based on the client location identifier.

10. The content delivery method of claim 1 wherein the content resource is a hypertext markup language (HTML) document and the embedded resource is a link.

11. The content delivery method of claim 1 wherein the request is received at a content delivery network, the request based on a CNAME identifying a domain of the content delivery network.

12. A content delivery device comprising:

at least one processor in communication with a non-transitory computer readable medium including computer readable instructions for:

generating an application programming interface (API);

at the API, receiving an embedded resource and at least one parameter associated with the embedded resource;

generating a modified embedded resource providing information as to a location within a content delivery network for obtaining the embedded resource.

13. The content delivery device of claim 11 wherein the computer readable instructions are further for replacing the embedded resource with the modified embedded resource in a resource.

14. The content delivery device of claim 12 wherein the resource is an HTML document and the modified embedded resource is a link.

15. The content delivery device of claim 12, the computer readable instructions further for identifying a unique tag associated with the embedded resource, and upon identifying the tag, passing the embedded resource to the API.

16. The content delivery device of claim 12 wherein the modified embedded resource includes at least one of an IP address to facilitate communication with a content delivery network to obtain content associated with the embedded resource, a virtual IP address to facilitate communication with a content delivery network to obtain content associated with the embedded resource, one or more absolute uniform resource locators to facilitate communication with a content delivery network to obtain content associated with the embedded resource, and one or more hostnames to facilitate communication with a content delivery network to obtain content associated with the embedded resource.

17. A non-transitory computer readable media comprising:

an HTML document comprising at least one embedded resource, the at least one embedded resource including a tag that causes the embedded resource to be analyzed to construct a modified embedded resource to replace the embedded resource, the

modified embedded resource causing a browser receiving the HTML document with the modified embedded resource to initiate communication with a content delivery network to obtain content associated with the embedded resource.

18. The non-transitory computer readable media of claim 17 wherein the tag causes the HTML document and the tag to be passed to an API to cause the embedded resource to be analyzed to construct the modified embedded resource.

19. The non-transitory computer readable medium of claim 17 wherein the modified embedded resource includes at least one of an IP address to facilitate communication with a content delivery network to obtain content associated with the embedded resource, a virtual IP address to facilitate communication with a content delivery network to obtain content associated with the embedded resource, one or more absolute uniform resource locators to facilitate communication with a content delivery network to obtain content associated with the embedded resource, and one or more hostnames to facilitate communication with a content delivery network to obtain content associated with the embedded resource.

20. A content delivery method comprising:


initiating an HTTP connection with a server, the HTTP connection including a client IP address that is available to the server;

receiving an HTML document including at least one embedded resource, the embedded resource including a reference to a content delivery network whereby the reference to the content delivery network is influenced by the client IP address; and

using the reference, initiating a request directly to the content delivery network to obtain content associated with the at least one embedded resource.

21. The content delivery method of claim 20 wherein the reference includes at least one of an IP address to facilitate communication with the content delivery network to obtain the content associated with the embedded resource, a virtual IP address to facilitate communication with the content delivery network to obtain the content

associated with the embedded resource, one or more absolute uniform resource locators to facilitate communication with the content delivery network to obtain the content associated with the embedded resource, and one or more hostnames to facilitate communication with the content delivery network to obtain the content associated with the embedded resource.



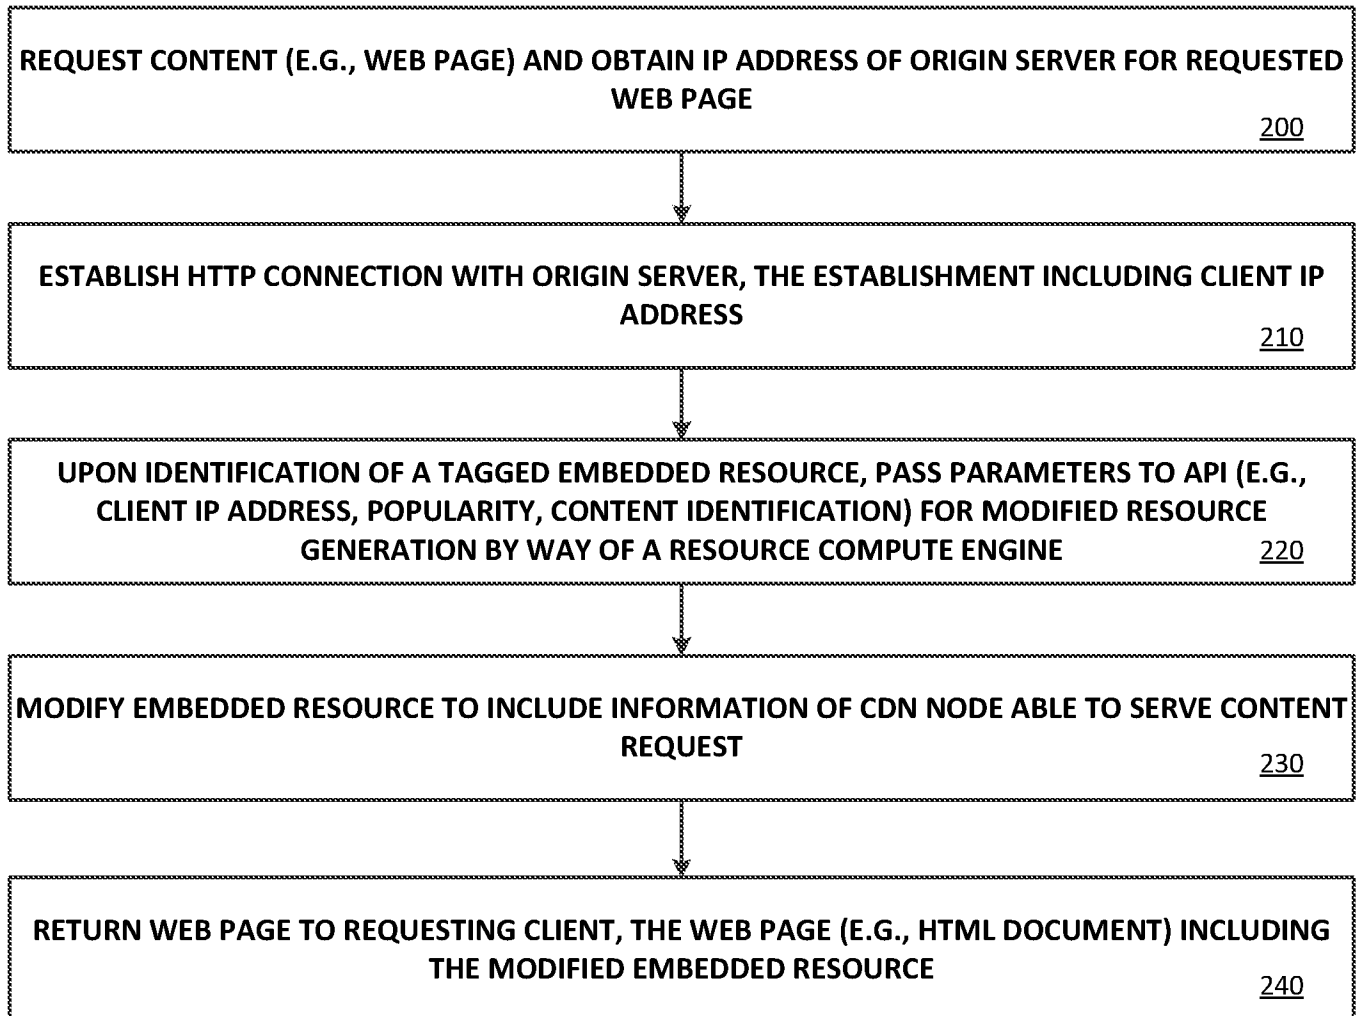


FIG. 2

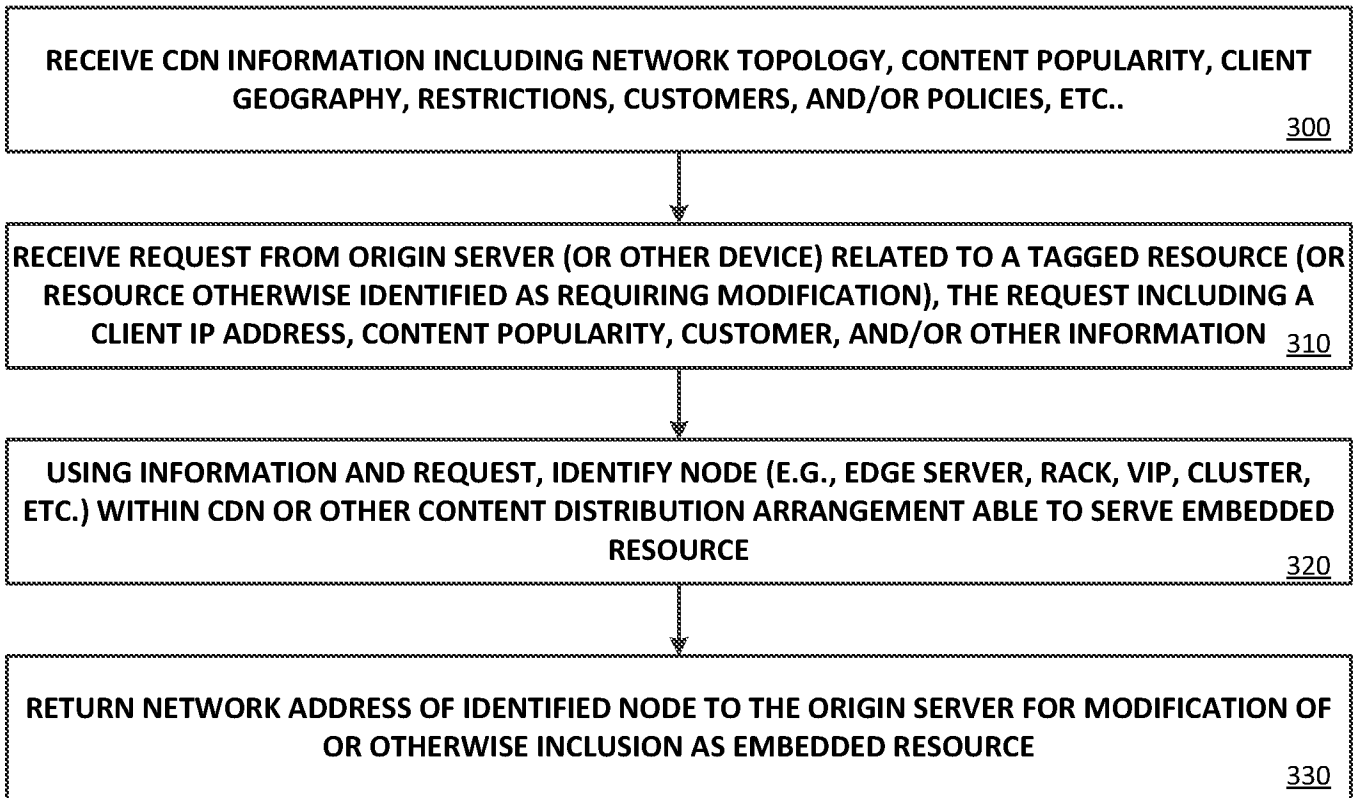


FIG. 3




FIG. 4

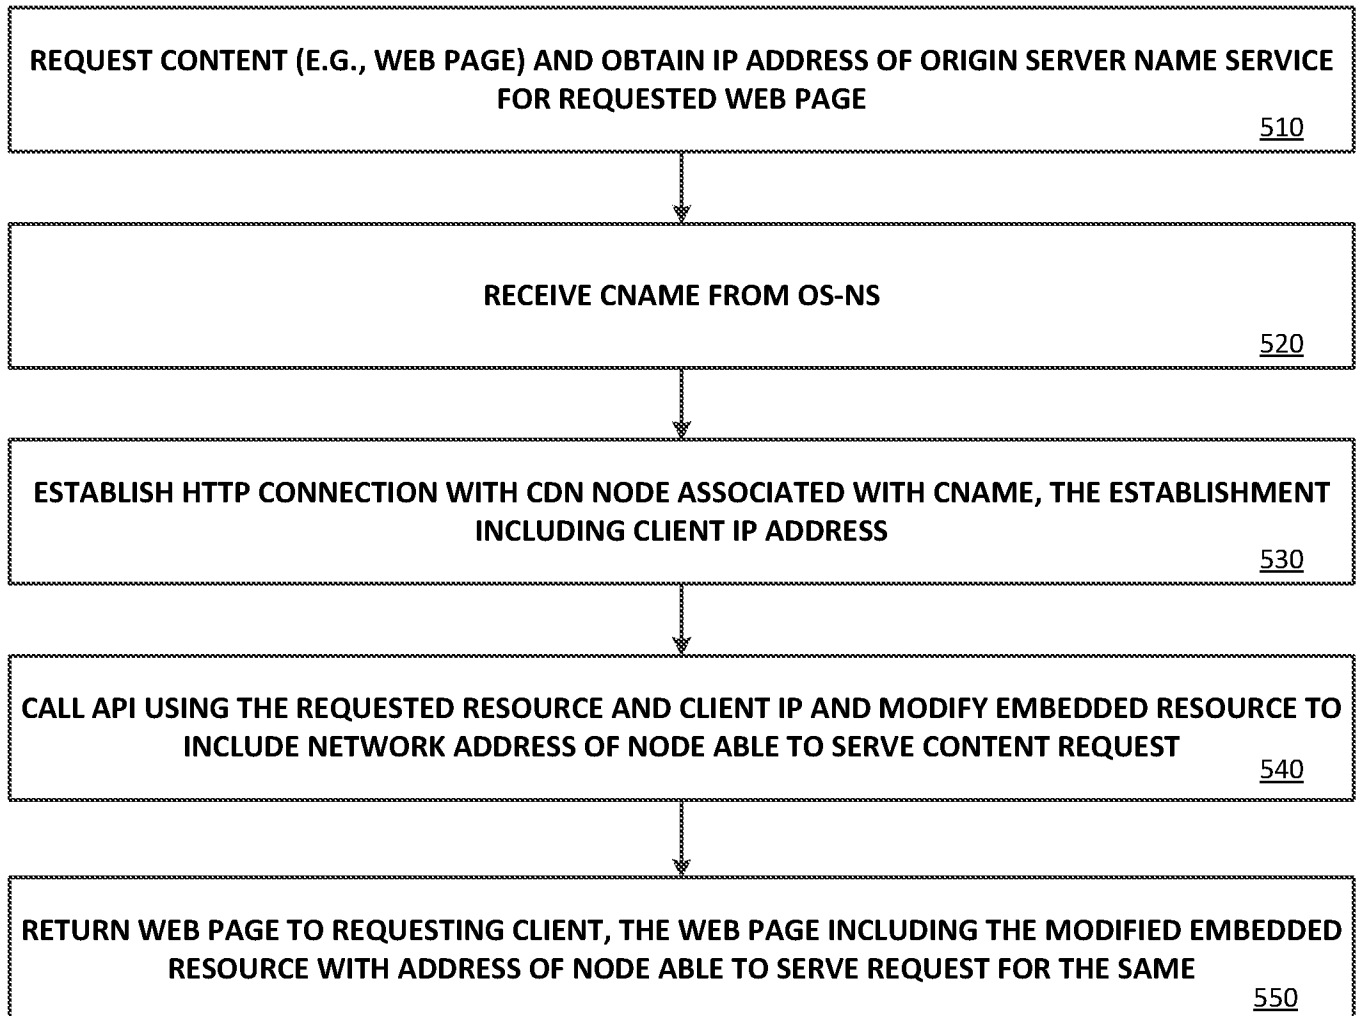


FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US15/68284

<p>A. CLASSIFICATION OF SUBJECT MATTER IPC(8) - H04L 29/08 (2016.01) CPC - H04L 61/30, 61/1511; G06F 17/30876 According to International Patent Classification (IPC) or to both national classification and IPC</p>														
<p>B. FIELDS SEARCHED</p> <p>Minimum documentation searched (classification system followed by classification symbols) IPC(8): H04L 29/06, 29/08, 29/12, 29/00, 12/00; G06F 17/30 (2016.01) CPC: H04L 61/30, 61/15, 61/1511, 61/00; G06F 17/30876, 17/30</p> <p>Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched</p> <p>Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) PatSeer (US, EP, WO, JP, DE, GB, CN, FR, KR, ES, AU, IN, CA, INPADOC Data); ProQuest; Google Patents; Google Scholar; IEEE Keywords: track, popularity, information, url, host, content, media, resource, threshold, edge, node, server, allocate, instantiate, provision, domain, dns, cdn, cname, api, application programmer interface</p>														
<p>C. DOCUMENTS CONSIDERED TO BE RELEVANT</p> <table border="1" style="width:100%; border-collapse: collapse;"> <thead> <tr> <th style="width:10%;">Category*</th> <th style="width:70%;">Citation of document, with indication, where appropriate, of the relevant passages</th> <th style="width:20%;">Relevant to claim No.</th> </tr> </thead> <tbody> <tr> <td>X --- Y</td> <td>US 2014/0289319 A1 (AMAZON TECHNOLOGIES, INC.) 25 September 2014; figures 1, 3A; paragraphs [0034], [0037], [0045], [0047], [0048]; claim 1.</td> <td>1-4, 6-21 ----- 5</td> </tr> <tr> <td>Y</td> <td>US 2013/0086358 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 04 April 2013; paragraphs [0013], [0043], [0045].</td> <td>5</td> </tr> <tr> <td>A</td> <td>US 8,250,211 B2 (SWILDENS, E. et al.) 21 August 2012; entire document.</td> <td>1-21</td> </tr> </tbody> </table>			Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.	X --- Y	US 2014/0289319 A1 (AMAZON TECHNOLOGIES, INC.) 25 September 2014; figures 1, 3A; paragraphs [0034], [0037], [0045], [0047], [0048]; claim 1.	1-4, 6-21 ----- 5	Y	US 2013/0086358 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 04 April 2013; paragraphs [0013], [0043], [0045].	5	A	US 8,250,211 B2 (SWILDENS, E. et al.) 21 August 2012; entire document.	1-21
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.												
X --- Y	US 2014/0289319 A1 (AMAZON TECHNOLOGIES, INC.) 25 September 2014; figures 1, 3A; paragraphs [0034], [0037], [0045], [0047], [0048]; claim 1.	1-4, 6-21 ----- 5												
Y	US 2013/0086358 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 04 April 2013; paragraphs [0013], [0043], [0045].	5												
A	US 8,250,211 B2 (SWILDENS, E. et al.) 21 August 2012; entire document.	1-21												
<p><input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.</p>														
<p>* Special categories of cited documents:</p> <table style="width:100%;"> <tr> <td style="width:50%;"> <p>“A” document defining the general state of the art which is not considered to be of particular relevance</p> <p>“E” earlier application or patent but published on or after the international filing date</p> <p>“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)</p> <p>“O” document referring to an oral disclosure, use, exhibition or other means</p> <p>“P” document published prior to the international filing date but later than the priority date claimed</p> </td> <td style="width:50%;"> <p>“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</p> <p>“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</p> <p>“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</p> <p>“&” document member of the same patent family</p> </td> </tr> </table>			<p>“A” document defining the general state of the art which is not considered to be of particular relevance</p> <p>“E” earlier application or patent but published on or after the international filing date</p> <p>“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)</p> <p>“O” document referring to an oral disclosure, use, exhibition or other means</p> <p>“P” document published prior to the international filing date but later than the priority date claimed</p>	<p>“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</p> <p>“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</p> <p>“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</p> <p>“&” document member of the same patent family</p>										
<p>“A” document defining the general state of the art which is not considered to be of particular relevance</p> <p>“E” earlier application or patent but published on or after the international filing date</p> <p>“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)</p> <p>“O” document referring to an oral disclosure, use, exhibition or other means</p> <p>“P” document published prior to the international filing date but later than the priority date claimed</p>	<p>“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</p> <p>“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</p> <p>“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</p> <p>“&” document member of the same patent family</p>													
<p>Date of the actual completion of the international search</p> <p>18 February 2016 (18.02.2016)</p>		<p>Date of mailing of the international search report</p> <p align="center">14 MAR 2016</p>												
<p>Name and mailing address of the ISA/ Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-8300</p>		<p>Authorized officer</p> <p align="center">Shane Thomas</p> <p>PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774</p>												