

(51) International Patent Classification:
H04L 29/08 (2006.01)(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.(21) International Application Number:
PCT/US2015/068292(22) International Filing Date:
31 December 2015 (31.12.2015)


(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/126,277 27 February 2015 (27.02.2015) US(71) Applicant: **LEVEL 3 COMMUNICATIONS, LLC** [US/US]; 1025 Eldorado Boulevard, Broomfield, Colorado 80021 (US).(72) Inventor: **NEWTON, Christopher**; 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).(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



(57) **Abstract:** A content delivery method including the operations of receiving a uniform resource locator resolution request at an authoritative name server for a domain where the uniform resource resolution request is received based, at least in part, on a host name of the uniform resource resolution request where the host name is uniquely related to a resource associated with the uniform resource resolution request. The method further including the operation of tracking a popularity of the resource based on the host name uniquely related to the resource and providing a location within a network capable of delivering the resource where the provided location is based on the popularity of the resource.

NETWORK ADDRESS RESOLUTION

RELATED APPLICATIONS

[0001] This Patent Cooperation Treaty (PCT) patent application claims priority from U.S. provisional application no. 62/126,277 entitled “NETWORK ADDRESS RESOLUTION,” filed on February 27, 2015, 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 the assignment of an enhanced host name to resources intended for distribution by a content delivery network, and the use of those enhanced host names to track popularity of the content through DNS and to resolve the request to a CDN node based on popularity of the content.

BACKGROUND

[0003] The Internet and the World Wide Web (the “Web”) have become ubiquitous. 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 origin servers or other locations, which are less efficient than edge servers in providing content.

[0004] Numerous forms of content may be served from the CDN. For example, television shows and movies may now be accessed from any number of Web sites, and the shows and movies may actually be served from the CDN. 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, app, or the like 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, and the updates and patches are delivered from a CDN.

[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 comprising the operations of receiving a uniform resource locator resolution request at an authoritative name server for a domain, the uniform resource locator resolution request received based, at least in part, on a host name of the uniform resource locator resolution request, the host name related to a resource associated with the uniform resource resolution request. The method may further include tracking a popularity of the resource based on the host name related to the resource and providing a node within a network capable of delivering the resource, the provided node based on the popularity of the resource. These operations may be performed by a computing device (or devices) including a processing unit, and may be encoded as computer readable instructions stored on a non-transitory computer readable medium.

[0007] Embodiments of the present disclosure may further involve an apparatus comprising at least one processing device in communication with a non-transitory computer readable media including computer executable instructions for receiving a uniform resource locator resolution request at an authoritative name server for a domain, the uniform resource locator resolution request received based, at least in part, on a host name of the uniform resource locator resolution request, the host name related to a resource associated with the uniform resource locator resolution request. The instructions may further be for tracking a popularity of the resource based on the host name uniquely related to the resource.

[0008] These and other embodiments are discussed in more detail below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] 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 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.

[0010] Figure 1 is a network diagram depicting a content delivery network, and other networks, that tracks resource popularity based on a unique host name assigned to a resource, and can resolve requests to a node within the content delivery network based, at least in part, on the popularity of the resource;

[0011] Figure 2 is a flow diagram illustrating a method of popularity tracking and content delivery based on a unique host name assigned to a resource;

[0012] Figure 3 is an example of a computing system that may implement various systems and methods discussed herein.

DETAILED DESCRIPTION

[0013] 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 also involve tracking the popularity of content (resources) based on the use of a unique enhanced host name assigned to the content. In one example, the popularity designation may take the form of globally assigning an enhanced host name to a resource, which may be done at or in association with the origin server for the resource (or a resource containing embedded resources with unique enhanced host names), where the enhanced host name is specific to the resource. Namely and in one example, at the origin server a unique host name for a resource may be assigned or otherwise defined. As discussed herein, the origin server, which may be a collection of servers and may be distributed, is the source for content from any given content provider. Origin servers may be maintained within the CDN or without.

[0014] In many conventional systems, there is not an awareness of the actual resource being requested at DNS time (when the host name is being resolved to an IP address); rather, only the host name is known. By providing a unique host name for the resource, some device within the DNS resolution function can track the popularity of the resource by understanding how many times the unique host name is being resolved. Accordingly, a popularity service is coupled with DNS infrastructure that tracks popularity and considers popularity when resolving a DNS request.

[0015] Examples of enhanced host names for two resources (resource 1 and resource 2) are as follows:

[0016] www.resource1.path.cdn.example.com/path/resource1

[0017] www.resource2.path.cdn.example.com/path/resource2

[0018] In such a situation, the authoritative name servers for example.com are provided with DNS records to delegate the subdomain (in this example, cdn.example.com) to the name servers of the CDN. Thus, when the requesting clients' ISP resolver contacts the authoritative name servers for example.com in order to resolve such an enhanced unique host name, the resolver is directed to the name servers of the CDN instead. The CDN name servers may track popularity of the resources by counting the number of resolution requests for any particular enhanced host name, and use that when determining to which tier of CDN nodes to direct the client. Within the CDN, one or more content server tiers may be arranged to serve both popular and unpopular content. For example, a relatively large number of content servers, positioned geographically proximate to various possible client locations, may be designated for popular content. Popular content requests from a client will be resolved by the CDN to a server within the popular tier, from where it will be served. As the content has been requested some number of times and tracked through DNS and the enhanced host name, the content will initially be accessed from a machine within the popular tier, and then, depending on the architecture of the CDN, may then be cached at that machine.

[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] A resource may also refer to a collection of objects that are combined in some manner to create a resource. For instance, a video or audio stream may be delivered as a series of small chunks (for example of some short duration) of a particular encoding rate. All such chunks, from all such encoding rates, may be considered collectively to a single resource such that they share a common enhanced host name. Similarly, pages that uniquely contain embedded resources may utilise the same enhanced host name for the container (e.g., HTML or XML document) and those embedded resources.

[0021] Figure 1 illustrates one example network configuration for practicing aspects of the present disclosure, while Figure 2 describes a related method. 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 browser 12. The computing device may be some form of conventional device such as personal computer, a thin client, a tablet, a smart phone, a set-top box, a cable box, or the like, that is capable of running a browser and providing a user with the ability to request content over a network connection, or some other app or 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 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.

[0022] To illustrate aspects of the disclosure, consider when a user enters into the browser, the web page address <http://www.example.com>. The browser 12 contacts or otherwise interacts with the resolver 14, which in turn requests the IP address for the www.example.com Web page from DNS 16 (operation 200). 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. The present flow discusses contacting an origin server for the Web page; however, this is done for explanation purposes, and it is possible that DNS may resolve the request to some other device for obtaining the Web page. 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 (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 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 www.example.com. In this example, the embedded resource is a link (e.g., a URL) to some form of content that the user can obtain by selecting the link when the Web page is displayed. In the case of an online news site, for example, the home page or any number of sub-pages (HTML documents) may include any number of links to content (embedded resources) such as articles, short video clips, etc.. In the case of a television or movie portal, the home page or any number of sub-pages, may include any number of links to particular episodes of a television series or specific movie releases.

[0023] The HTML document is provided with URL's 26 of the embedded resources 24 that have unique enhanced host names for the specific resource. Here, the customer of the CDN does not need to understand popularity of the content or

otherwise pre-designate popularity. For example, consider the case where www.example.com is the home page for an on-line video provider and the embedded resource is a link to a recently released movie that did well in theaters. In such an example, when the movie is released by the video provider, it will likely be popular (requested numerous times), and thus the enhanced host name for the movie will be resolved by the CDN name servers 28 numerous times and thereby become designated as “popular” by the CDN.

[0024] In the examples discussed herein, content is either popular or non-popular, and popularity may be designated based on the number of times a specific enhanced host name is resolved by DNS which may specifically be CDN name servers 28 for enhanced host names. However, popularity may be designated in other scales, such as 1 = low, 2 = medium, 3 = high, scale of 1-10 with 1 being unpopular and 10 being the most popular, etc., with thresholds (number of times resolved) pertaining to the scale. Moreover, the scale may be adjusted from an initial designation 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 also be regionalized. For example, a resource in the case of a news story may initially be considered to be popular in the region from which it originates but un-popular in other regions. Say, for example, in the case of a breaking news story in Atlanta, the on-line news story (embedded resource) may be designated popular for Georgia or the Southeast of the United States, but may be designated unpopular for all other geographic regions. Thus, popularity services may be regionalized and popularity tracked based on the number of DNS resolutions in any given geographical region.

[0025] To provide a popularity designation at DNS time, a unique enhanced host name is provided for the embedded resource. With respect to the unique enhanced host names, in one example, if resource 1 (e.g., a recently released block buster movie) is given a unique host name and resource 2 (e.g., a very old esoteric documentary film) is given a different unique host name. As introduced above, two enhanced host names 26 may be: Embedded resource 1 (ER1) = (<http://resource1.path.cdn.example.com/path/resource1>) for the popular content) and Embedded resource 2 (ER2) = <http://resource2.path.cdn.example.com/path/resource2>

(for the unpopular content). The unique URL host names are embedded within an HTML page 22 from which a user may request either of the resources. Two examples enhanced host names are provided herein; however, any number of alternatives host names may be used, and numbers (N) of resources and corresponding enhanced host names, beyond one or two, may also be used.

[0026] Alternatively, the host name need not include an explicit reference to the resource in the path; rather it may include some per-resource identifier, such as an index, a checksum (e.g., MD5) of the path, etc., as discussed in more detail below. The host name could also include the client IP address (or a representation thereof) to count popularity. Alternatively, the relative popularity may be noted at name resolution time, since the same name is likely to be cached for some at the resolver and so may be used by a number of clients without contacting the CDN DNS system.

[0027] Returning to the client request, the HTML document 22 is returned to the client device 10 with the embedded resources as described above (operation 220). When the client then selects either resource, the browser 12 again communicates with the ISP resolver 14 to look-up (through DNS) the enhanced host name associated with the resource. Using the popular example, the ISP resolver would contact DNS to resolve resource1.path.cdn.example.com. In one specific implementation, the authoritative name servers for example.com would delegate the path.example.com subdomain to the authoritative name servers for the CDN 28. When resolving the host name from one of the embedded resources, the ISP resolver associated with the client device would be directed to the authoritative name servers of the CDN, which would then be able to track popularity of each name so requested before determining which popularity tier to assign the request (operation 230). The CDN DNS servers 28 would then select the most appropriate distribution node from within that popularity tier for the client, and return its IP address to the ISP resolver. The CDN name servers may use some form of intelligent traffic management 36 (best distributor algorithms) to direct the client to obtain the content from either a content server in the popular 30 or unpopular 32 domain. Because of the unique enhanced host name, the CDN name servers track the popularity of each unique enhanced host name, which may be achieved by a popularity service 34 associated with the CDN DNS. The CDN name servers also

return an IP address for the best server node within the CDN to deliver the content (operation 240). For popular content, the node will be within the popular tier, for unpopular content, the node will be in the unpopular tier.

[0028] The intelligent traffic management (ITM) function may return an address that is optimally suited to serve the content to the requesting client. So, in a simple example, for a user in New York requesting popular content from a CDN having nodes capable of serving the 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 ITM may instead return the New York node to the client in New York. Note, in one example, the CDN name servers may use the location of the resolver 14 as a proxy for the location of the client device 10.

[0029] The ITM may take into account and use the IP address of the client device, the IP address of the resolver, customer, and the various other forms of information in determining the node from which to serve the content. For example, the ITM may use geographical information to pick the node closest to the client device or resolver 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 that is close but less busy. In some instances, such enhanced node selection may only be performed for popular content with unpopular content served from the origin or tiers besides those at the edge.

[0030] In some instances, for example, the node may not yet have the content stored locally, 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 when operating in a caching CDN, for example.

[0031] The CDN name servers return the determined network address (e.g., URL, URI, etc.) to the requesting client device, and the browser 12 (or other device) then may establish a connection with the CDN node designated to serve the embedded resource. The designated node will either have the embedded resource, or it will obtain the content and serve it to the requesting client. Within the CDN, as client's request

content, the popular tier nodes will quickly have all of the content cached and be able to deliver the content directly upon receiving a request from a client.

[0032] In one possible implementation, the unique hostname may be generated from the path within the base domain. For example, using a function that reverses the order of the path elements and replaces “/” with “.” the path /some/dir/resource.ext might result in resourceext.dir.some.path.example.com.

[0033] In an alternative implementation, a unique hostname may be used for each resource (which would allow the DNS system to track popularity). So, for a resource such as: <http://origin.example.com/path/resource>, the system would use a URL such as: <http://resource.path.cdn.example.com/path/resource>. In this example, the DNS servers see the entire "path," including a representation of the resource and its path. In the example of <http://origin.example.com/path/resource>, the DNS servers would not obtain /path/resource, and hence would be unaware of the path and resource subject to the request. However, in the popularity tracking URL, DNS does see a representation of the resource and path, and hence can track the number of times there is a request for the content. Accordingly, DNS may track resource popularity.

[0034] In an alternative, the system may use a hostname that uniquely represents each /path/resource at the origin in order to provide popularity tracking without having the actual path and resource name provided in the URL and available to DNS. In such an alternative, the system may use an index or a checksum (e.g., MD5) of the /path/resource to provide a unique host name representative of the path/resource, such as: <http://a1242.cdn.example.com/path/resource>. In this example, the name within the cdn.example.com domain is different for each different value of /path/resource. Similar to the enhanced unique host name referenced above, DNS may track the number of requests for a resource, and hence can track popularity.

[0035] Now, if a checksum were used, it could also be used to verify that the path was valid and/or that the hostname is correct for the resource. In some instances, a system may have problems when someone uses the hostname from one URL on a second unrelated URL -- form of a “link hijacking.” To validate a hostname relative to the path, if the hostname were or included a checksum, the receiving CDN node could immediately reject the request if it didn't align with the /path/resource being requested of

it by a client. This check could be partial or complete; that is, the checksum may be over some portion of the /path/resource or it could be of the entire /path/resource. Using a part of /path/resource would mean that the popularity is being tracked for a set of resources rather than per resource, but with the benefit of somewhat reducing the number of names needing to be tracked. This may be of benefit when dealing with groups of closely related objects (e.g., individual chunks of a streaming video asset).

[0036] Figure 3 is an example schematic diagram of a computing system 600 that may implement various methodologies discussed herein. 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.

[0037] 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.

[0038] 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.

[0039] 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.

[0040] As shown, main memory 604 is encoded with an enhanced host name popularity rendezvous engine application 608-1 that supports functionality as discussed above and as discussed further below. For example, in one embodiment, the application 608-1 may include or otherwise implement the various processes and/or instructions described herein. The application 608-1 (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 an EHNPR process 608-2, such as through logic instructions, executing on the processor 602 and based on the application 608-1 stored in main memory or otherwise tangibly stored.

[0041] 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.

[0042] 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., EEPROM and EEPROM); flash memory; or other types of medium suitable for storing electronic instructions.

[0043] 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.

[0044] 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 uniform resource locator resolution request at an authoritative name server for a domain, the uniform resource locator resolution request received based, at least in part, on a host name of the uniform resource locator resolution request, the host name related to a resource associated with the uniform resource resolution request;

tracking a popularity of the resource based on the host name related to the resource; and

providing a node within a network capable of delivering the resource, the provided node based on the popularity of the resource.

2. The content delivery method of claim 1 wherein the node is an edge server within a content delivery network.

3. The content delivery method of claim 1 wherein the host name is constructed from a base domain name and a function of a path to the resource.

4. The content delivery method of claim 3 wherein the host name includes a checksum uniquely related to the path.

5. The content delivery method of claim 4 further comprising:

validating that the host name matches the path and resource based on the checksum.

6. The content delivery method of claim 1 wherein the host name includes an index uniquely related to the resource.

7. The content delivery method of claim 1 further comprising converting a path indicative of the resource to the host name.

8. The content delivery method of claim 1 wherein the host name is unique host name, and is included in a link comprising the form: resource.path.cdn.example.com/path/resource.

9. An apparatus comprising:

at least one processing device in communication with a non-transitory computer readable media including computer executable instructions for:

receiving a uniform resource locator resolution request at an authoritative name server for a domain, the uniform resource locator resolution request received based, at least in part, on a host name of the uniform resource locator resolution request, the host name related to a resource associated with the uniform resource locator resolution request; and

tracking a popularity of the resource based on the host name uniquely related to the resource.

10. The apparatus of claim 9, the computer executable instructions further for providing a node within a network capable of delivering the resource, the provided node based on the popularity of the resource.

11. The apparatus of claim 10 wherein the node is an edge server within a content delivery network.

12. The apparatus of claim 9 wherein the host name is constructed from a base domain name and a function of a path to the resource.

13. The apparatus of claim 12 wherein the host name includes a checksum uniquely related to the path.

14. The apparatus of claim 13, the computer executable instructions further for validating that the host name matches the path and resource based on the checksum.

15. The apparatus of claim 9 wherein the host name includes an index uniquely related to the resource.

16. The apparatus of claim 9 wherein the computer executable instructions are further for converting a path indicative of the resource to the host name.

17. The apparatus of claim 9 wherein the host name is uniquely related to the resource associated with the uniform resource locator resolution request.

18. The apparatus of claim 17 wherein the host name uniquely related to the resource is included in a link comprising the form:

resource.path.cdn.example.com/path/resource




FIG. 1





FIG. 2



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US15/68292

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

B. FIELDS SEARCHED

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

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

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

PatSeer (US, EP, WO, JP, DE, GB, CN, FR, KR, ES, AU, JN, 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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/0289319 A1 (AMAZON TECHNOLOGIES, INC.) 25 September 2014; figures 1, 3A; paragraphs [0034], [0037], [0045], [0047], [0048]; claim 1.	1-3, 7-12, 16-18
---		-----
Y		4-6, 13-15
Y	US 2012/0167227 A1 (HENDERSON, JR., R. et al.) 28 June 2012; paragraphs [0015], [0016], [0032].	4-5, 13-14
Y	US 2014/0053237 A1 (AVENTAIL LLC) 20 February 2014; paragraph [0048].	6, 15

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

18 February 2016 (18.02.2016)

Date of mailing of the international search report

11 MAR 2016

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

Authorized officer

Shane Thomas

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774

摘要

一种内容传送方法，包括操作：在域的授权名称服务器处接收统一资源定位符解析请求，其中，所述统一资源定位符解析请求是至少部分地基于所述统一资源定位符解析请求的主机名而被接收的，其中，所述主机名与关联于所述统一资源定位符解析请求的资源相关。该方法还包括操作：基于与所述资源唯一相关的主机名跟踪所述资源的流行度，并且提供网络内能够传送所述资源的位置，其中，所提供的位置基于所述资源的流行度。