




US 20100057894A1

(19) **United States**(12) **Patent Application Publication**
Glasser(10) **Pub. No.: US 2010/0057894 A1**(43) **Pub. Date: Mar. 4, 2010**(54) **TARGETED CACHING TO REDUCE
BANDWIDTH CONSUMPTION****Publication Classification**(75) Inventor: **Alan Glasser**, Manalapan, NJ (US)(51) **Int. Cl.**
G06F 15/177 (2006.01)(52) **U.S. Cl.** **709/222**Correspondence Address:
AT&T Legal Department - LNA
Attn: Patent Docketing
Room 2A- 207, One AT & T Way
Bedminster, NJ 07921 (US)(57) **ABSTRACT**

A system includes a name server, an edge cache server, and a local cache server. The name server is configured to provide an anycast IP address in response to a request for an IP address of an origin hostname from a client system. The edge cache server is configured to respond to the anycast IP address and a unicast IP address and to retrieve content from an origin. The local cache server includes a storage and is configured to respond to the anycast IP address, to retrieve content from the edge cache server, and provide the content to a client system.

(73) Assignee: **AT&T CORP.**, New York, NY (US)(21) Appl. No.: **12/199,200**(22) Filed: **Aug. 27, 2008**




FIG. 2




FIG. 3




FIG. 4

TARGETED CACHING TO REDUCE BANDWIDTH CONSUMPTION

FIELD OF THE DISCLOSURE

[0001] The present disclosure generally relates to communications networks, and more particularly relates to targeted caching to reduce bandwidth consumption.

BACKGROUND

[0002] Packet-switched networks, such as networks based on the TCP/IP protocol suite, can distribute a rich array of digital content to a variety of client applications. One popular application is a personal computer browser for retrieving documents written in the Hypertext Markup Language (HTML) over the Internet via the Hypertext Transfer Protocol (HTTP). Frequently, these documents include embedded content. Where once the digital content consisted primarily of text and static images, digital content has grown to include audio and video content as well as dynamic content customized for an individual user. Similarly, software updates can be distributed through the network via HTTP.

[0003] It is often advantageous when distributing digital content across a packet-switched network to divide the duty of answering content requests among a plurality of geographically dispersed servers. For example, popular Web sites on the Internet often provide links to “mirror” sites that replicate original content at a number of geographically dispersed locations. A more recent alternative to mirroring is content distribution networks (CDNs) that dynamically redirect content requests to a cache server situated closer to the client issuing the request. CDNs either co-locate cache servers within Internet Service Providers or deploy them within their own separate networks.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] It will be appreciated that for simplicity and clarity of illustration, elements illustrated in the Figures have not necessarily been drawn to scale. For example, the dimensions of some of the elements are exaggerated relative to other elements. Embodiments incorporating teachings of the present disclosure are shown and described with respect to the drawings presented herein, in which:

[0005] FIG. 1 is a diagram illustrating a content distribution network (CDN) in accordance with one embodiment of the present disclosure;

[0006] FIG. 2 is block diagram illustrating a system for distributing content to client systems in accordance with one embodiment of the present disclosure;

[0007] FIG. 3 is a flow diagram illustrating a method of distributing content to client systems in accordance with one embodiment of the present disclosure; and

[0008] FIG. 4 is an illustrative embodiment of a general computer system.

[0009] The use of the same reference symbols in different drawings indicates similar or identical items.

DETAILED DESCRIPTION OF THE DRAWINGS

[0010] The numerous innovative teachings of the present application will be described with particular reference to the presently preferred exemplary embodiments. However, it should be understood that this class of embodiments provides only a few examples of the many advantageous uses of the innovative teachings herein. In general, statements made in

the specification of the present application do not necessarily limit any of the various claimed inventions. Moreover, some statements may apply to some inventive features but not to others.

[0011] FIG. 1 shows a geographically dispersed network **100**, such as the Internet. Network **100** can include routers **102**, **104**, and **106** that communicate with each other and form an autonomous system (AS) **108**. AS **108** can connect to other ASs that form network **100** through peering points at routers **102** and **104**. Additionally, AS **108** can include client systems **110**, **112**, and **114** connected to respective routers **102** and **106** to access the network **100**. Router **102** can provide ingress and egress for client system **110**. Router **106** can provide ingress and egress for both of client systems **112** and **114**.

[0012] Client systems **110**, **112**, and **114** can retrieve information from an origin server **116**. For example, client system **112** can retrieve content, such as videos, music, portable document format files, and software updates or patches, provided by origin server **116**. Alternatively, the content may be a data file for a distributed computing project. Client systems **110**, **112**, and **114** can retrieve the data file from the origin server **116**, process a portion of the data file, and return a result file to either origin server **116** or another server. The results file for the client systems **110**, **112**, and **114** can be combined with result files from additional client systems to obtain a result for the data file.

[0013] The time normally required for client system **112** to retrieve the information from the origin server **116** is related to the size of the file, the distance the information travels, and available bandwidth along the route. Additionally, the load on the origin server **116** is related to the number of client systems **110**, **112**, and **114** that are actively retrieving information from the origin server **116**. As such, the resources such as processor, memory, and bandwidth available to the origin server **116** limit the number of client systems **110**, **112**, and **114** that can simultaneously retrieve information from the origin server **116**.

[0014] Additionally, the network can include an edge cache server **118** and a local cache server **120** replicating content from the origin server **116** that can be located more closely within the network to the client systems **110**, **112**, and **114**. Edge cache server **118** can connect to router **102**, and local cache server **120** can connect to router **106**. Client systems **110**, **112**, and **114** can retrieve content from edge cache server **118** or local cache server **120** to decrease the time required for delivery of the content. When the content is a data file for a distributed computing project, the client systems **110**, **112**, and **114** may return the result files to the origin server **116**, the edge cache server **118**, the local cache server **120**, or any combination thereof. In an embodiment, the network **100** may contain a plurality of local cache servers and/or a plurality of edge cache servers geographically distributed to reduce latency to additional client systems.

[0015] In an embodiment, edge cache server **118** and local cache server **120** can both respond to an anycast address. Generally, an anycast address is an Internet Protocol (IP) address that identifies two or more receivers, but only one of them is chosen at any given time to receive information from any given sender. Routing protocols utilized by network **100** can direct traffic along a shortest path between the sender and the receiver. In the case of an anycast address, each router in the network can send packets addressed to the anycast address towards the closest system announcing the anycast address.

For example, router **102** can direct packets addressed to the anycast address to edge cache server **118**, router **106** can direct packets addressed to the anycast address to the local cache server **120**. From the router **104**, packets addressed to the anycast address must pass to either router **102** or router **106**. In an embodiment where the network distance from router **104** to router **102** is substantially similar to the network distance from router **104** to router **106**, network policies can be used to determine a preferred route. Packets addressed to the anycast address from client system **110** can first reach router **102** and be directed to edge cache server **118**. Packets addressed to the anycast address from client systems **112** and **114** can first reach router **106** and can be directed to local cache server **120**.

[0016] In an embodiment, the network can utilize a routing protocol, such as border gateway protocol (BGP), to determine routes through the network. A router can receive network layer reachability information (NLRI) from peer devices. Based on the NLRI, the router can determine a best next hop for network traffic destined for an IP address. The best next hop can be the peer device with a shortest network distance to the destination IP address. The local cache server **120** and the edge cache server **118** can participate in a BGP group, providing NLRI information indicating that the local cache server is a preferred destination for packets addressed to the anycast address. Peer routers can determine a preferred route for packets addressed to the anycast address based on network policies and their distance to the local cache server and to the edge cache server.

[0017] Client system **112** may send a request for content to local cache server **120**. If local cache server **120** has the information stored in a cache, it can provide the content to client system **112**. This can decrease the distance the information travels and reduce the time to retrieve the information. Alternatively, when local cache server **120** does not have the content, it can retrieve the content from edge cache server **118** prior to providing the content to the client system **112**. In an alternate embodiment, the local cache server **120** may retrieve the content from another local cache server **120** within the network.

[0018] Edge cache server **118** may either provide cached content to local cache server **120** or retrieve the content from the origin server **116** prior to providing the content to local cache server **120**. In another alternate embodiment, the edge cache server **118** may retrieve the content from another server replicating the content stored on the origin server **116**, such as another edge cache server **118**.

[0019] Edge cache server **118** and local cache server **120** can cache the content in order to respond to additional requests for the content without first retrieving the content from another server. In this way, the content may be retrieved from the origin server **116** only once, reducing the load on origin server **116** and network **100** such as, for example, when client systems **114** and **110** request the same content.

[0020] FIG. 2 shows a block diagram illustrating an exemplary system **200** for distributing content from an origin server **202** to client systems **204** and **206**. The system **200** can include a Domain Name System (DNS) server **208**, a local cache server **210**, an edge cache server **212**, and a map server **214**. The local cache server **210** can be a diskless cache server. For example, the local cache server **210** may have a non-volatile memory, such as a flash memory for storing an operating system and settings, and a substantial amount of volatile memory such as RAM for caching content files. In an

embodiment, the local cache server **210** may lose the content files from the volatile memory during a loss of power or a restart. Alternatively, the local cache server **210** may utilize the non-volatile memory for storing content files during a powered down state.

[0021] In an embodiment, client system **204** can send a request **216** to the DNS server **208**. The request **216** can include a hostname of the origin server **202**. The DNS server **208** can send a response **218** including an anycast address. The anycast address can be an IP address shared by the local cache server **210** and the edge cache server **212**. Client system **204** can send a request **220** for the content using the anycast address. In an embodiment, the content can be a relatively popular content, such as a Microsoft Windows update, that can be requested by a significant number of clients within a relatively short period of time. When the network distance from client system **204** to local cache server **210** is less than the network distance from client system **204** to edge cache server **212**, the local cache server **210** can receive the request **220**. In an embodiment, the local cache server **210** may periodically receive an instruction **222** from the map server **214**. The instruction **222** can include a map server table. The map server table can include the anycast IP address and a unicast IP address for the edge cache server **212** associated with the anycast IP address. For example, the map server **214** can periodically push the map server table to the local cache server **210**. Alternatively, the local cache server **210** can periodically request the map server table from the map server **214**. In a further embodiment, the local cache server **210** may not announce a route for the anycast address until it receives the map server table.

[0022] When the local cache server **210** does not have the content, the local cache server can send a request **224** to edge cache server **212** for the content. When edge cache server **212** does not have the content, edge cache server **212** may send a request **226** to the origin server **202**. The origin server **202** can send a response **228** including the content to the edge cache server **212**. The edge cache server **212** can store the content and send a response **230** including the content to the local cache server **210**. Similarly, the local cache server **210** can store the content and send a response **232** including the content to the client system **204**.

[0023] In another embodiment, client system **206** can send a request **234** to the DNS server **208**. The request can include the hostname of the origin server **202**. The DNS server **208** can send a response **234** including the anycast address. Client system **206** can send a request **238** for the content using the anycast address. When the network distance from client system **204** to local cache server **210** is greater than the network distance from client system **204** to the edge cache server **212**, the edge cache server **212** can receive the request **238**. When the edge cache server **212** has the content, the edge cache server **212** can send a response **240** including the content to the client system **206**.

[0024] FIG. 3 shows a flow diagram illustrating an exemplary method for providing client systems with content. At **302**, the client system can request an IP address of an origin server. The request can include a hostname of an origin server. At **304**, the DNS server can provide an anycast address. The anycast address can be an IP address shared by a local cache server and an edge cache server. At **306**, the client system can send a request for the content to the anycast address. At **308**, the network can determine if the request is routed to the local cache server. When network distance from the client system

to the edge cache server is less than the network distance from the client system to the local cache server, the request can be routed to the edge cache server rather than to the local cache server. At 310, the edge cache server can determine if the content is cached or if the content must be retrieved from the origin server. When the content is cached at the edge cache server, the edge cache server can provide the content to the client, as illustrated at 312. Alternatively, when the content is not cached at the edge cache server, the edge cache server can retrieve the content from the origin server, as illustrated at 314, prior to providing the content to the client, as illustrated at 312.

[0025] Alternatively, when the network distance from the client system to the edge cache server is greater than the network distance from the client system to the local cache server, the request can be routed to the local cache server rather than the edge cache server. At 316, the local cache server can determine if the content is cached. When the content is cached at the local cache server, the local cache server can provide the content to the client, as illustrated at 318.

[0026] Alternatively, when the content is not cached at the local cache server, the local cache server can request the content from the edge cache server, as illustrated at 320. At 322, the edge cache server can determine if the content is cached or if the content must be retrieved from the origin server. When the content is cached at the edge cache server, the edge cache server can provide the content to the local cache server, as illustrated at 324, and the local cache server can provide the content to the client, as illustrated at 318. Alternatively, when the content is not cached at the edge cache server, the edge cache server can retrieve the content from the origin server, as illustrated at 326, and can provide the content to the local cache server, as illustrated at 324.

[0027] FIG. 4 shows an illustrative embodiment of a general computer system 400. The computer system 400 can include a set of instructions that can be executed to cause the computer system to perform any one or more of the methods or computer based functions disclosed herein. The computer system 400 may operate as a standalone device or may be connected, such as by using a network, to other computer systems or peripheral devices.

[0028] In a networked deployment, the computer system may operate in the capacity of a server or as a client user computer in a server-client user network environment, or as a peer computer system in a peer-to-peer (or distributed) network environment. The computer system 400 can also be implemented as or incorporated into various devices, such as a personal computer (PC), a tablet PC, a set top box, a personal digital assistant (PDA), a mobile device, a palmtop computer, a laptop computer, a desktop computer, a communications device, a wireless telephone, a land-line telephone, a control system, a camera, a scanner, a facsimile machine, a printer, a pager, a personal trusted device, a web appliance, a network router, switch or bridge, or any other machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. In a particular embodiment, the computer system 400 can be implemented using electronic devices that provide voice, video or data communication. Further, while a single computer system 400 is illustrated, the term "system" shall also be taken to include any collection of systems or sub-systems that individually or jointly execute a set, or multiple sets, of instructions to perform one or more computer functions.

[0029] The computer system 400 may include a processor 402, such as a central processing unit (CPU), a graphics processing unit (GPU), or both. Moreover, the computer system 400 can include a main memory 404 and a static memory 406 that can communicate with each other via a bus 408. As shown, the computer system 400 may further include a video display unit 410 such as a liquid crystal display (LCD), an organic light emitting diode (OLED), a flat panel display, a solid-state display, or a cathode ray tube (CRT). Additionally, the computer system 400 may include an input device 412 such as a keyboard, and a cursor control device 414 such as a mouse. Alternatively, input device 412 and cursor control device 414 can be combined in a touchpad or touch sensitive screen. The computer system 400 can also include a disk drive unit 416, a signal generation device 418 such as a speaker or remote control, and a network interface device 420 to communicate with a network 426. In a particular embodiment, the disk drive unit 416 may include a computer-readable medium 422 in which one or more sets of instructions 424, such as software, can be embedded. Further, the instructions 424 may embody one or more of the methods or logic as described herein. In a particular embodiment, the instructions 424 may reside completely, or at least partially, within the main memory 404, the static memory 406, and/or within the processor 402 during execution by the computer system 400. The main memory 404 and the processor 402 also may include computer-readable media.

[0030] The illustrations of the embodiments described herein are intended to provide a general understanding of the structure of the various embodiments. The illustrations are not intended to serve as a complete description of all of the elements and features of apparatus and systems that utilize the structures or methods described herein. Many other embodiments may be apparent to those of skill in the art upon reviewing the disclosure. Other embodiments may be utilized and derived from the disclosure, such that structural and logical substitutions and changes may be made without departing from the scope of the disclosure. Additionally, the illustrations are merely representational and may not be drawn to scale. Certain proportions within the illustrations may be exaggerated, while other proportions may be minimized. Accordingly, the disclosure and the FIGs. are to be regarded as illustrative rather than restrictive.

[0031] The Abstract of the Disclosure is provided to comply with 37 C.F.R. §1.72(b) and is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing Detailed Description of the Drawings, various features may be grouped together or described in a single embodiment for the purpose of streamlining the disclosure. This disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter may be directed to less than all of the features of any of the disclosed embodiments. Thus, the following claims are incorporated into the Detailed Description of the Drawings, with each claim standing on its own as defining separately claimed subject matter.

[0032] The above disclosed subject matter is to be considered illustrative, and not restrictive, and the appended claims are intended to cover all such modifications, enhancements, and other embodiments which fall within the true spirit and scope of the present disclosed subject matter. Thus, to the maximum extent allowed by law, the scope of the present

disclosed subject matter is to be determined by the broadest permissible interpretation of the following claims and their equivalents, and shall not be restricted or limited by the foregoing detailed description.

What is claimed is:

1. A system comprising:
 - a name server configured to provide an anycast Internet Protocol (IP) address in response to a request for an IP address of an origin hostname from a client system;
 - an edge cache server configured to respond to the anycast IP address and to a unicast IP address, to retrieve content from an origin, and to provide content to the requesting system; and
 - a local cache server including a storage, the local cache server configured to respond to the anycast IP address, to retrieve the content from the edge cache server, and to provide the content to the client system.
2. The system of claim 1 wherein storage is a diskless storage.
3. The system of claim 1 wherein the local cache server is configured to announce a route for the anycast IP address, the route directing requests to the local cache server.
4. The system of claim 3 wherein the local cache server is configured to announce the route using border gateway protocol.
5. The system of claim 3 wherein the local cache server is configured to announce the route after receiving a table to convert the anycast IP address to the unicast IP address.
6. The system of claim 1 wherein the edge cache server is configured to announce a route for the anycast IP address, the route directing requests to the edge cache server.
7. The system of claim 1 wherein the edge cache server is configured to announce the route using border gateway protocol.
8. The system of claim 1 further comprising a map server, the map server configured to provide the local cache server with a table to convert the anycast IP address to the unicast IP address.
9. The system of claim 8 wherein the map server pushes the table to the local cache server.
10. The system of claim 8 wherein the local cache server periodically requests the table from the map server.
11. A method comprising:
 - receiving a request for content from a client system, the request addressed to an anycast Internet Protocol (IP) address;

- determining that content is not available in a local storage;
- determining a unicast IP address of an edge cache server based on the anycast IP address;
- retrieving the content from the edge cache server; and
- providing the content to the client system.

12. The method of claim 11 wherein the local storage is a diskless storage.

13. The method of claim 11 further comprising storing the content in the local storage and providing the content to a second client system.

14. The method of claim 11 further comprising announcing a route for the anycast address.

15. The method of claim 14 wherein the announcing is accomplished using border gateway protocol.

16. The method of claim 11 wherein determining the unicast IP address includes accessing a table.

17. The method of claim 16 further comprising receiving the table from a mapping server.

18. A device comprising:

- a local storage configured to store content; and
- a processor in communication with the local storage, the processor configured to:

- receive a request for content from a client system, the request addressed to an anycast Internet Protocol (IP) address;

- determine that content is not available in a local storage;
- determine a unicast IP address of an edge cache server based on the anycast IP address;

- retrieve the content from the edge cache server; and
- provide the content to the client system.

19. The device of claim 18 wherein the local storage is a diskless storage.

20. The device of claim 18 wherein the processor is further configured to providing the content to a second client system.

21. The device of claim 18 wherein the processor is further configured to announce a route for the anycast address.

22. The device of claim 21 wherein the process is further configured to announce the route using border gateway protocol.

23. The device of claim 18 wherein the processor is further configured to receive a table from a mapping server and to determine the unicast IP address based on the table.

* * * * *