

(51) International Patent Classification:
G06F 15/16 (2006.01)(21) International Application Number:
PCT/US2009/047832(22) International Filing Date:
18 June 2009 (18.06.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/165,343 30 June 2008 (30.06.2008) US(71) Applicant (for all designated States except US): **AMAZON TECHNOLOGIES, INC.** [US/US]; P.O. Box 8102, Reno, NV 89507 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **SIVASUBRAMANIAN, Swaminathan** [IN/US]; 1200 12th Avenue South, Suite 1200, Seattle, WA 98144-2734 (US). **RICHARDSON, David, R.** [US/US]; 1200 12th Avenue South, Suite 1200, Seattle, WA 98144-2734 (US). **SCOFIELD, Christopher, L.** [US/US]; 1200 12th Avenue South, Suite 1200, Seattle, WA 98144-2734 (US). **MARSHALL, Bradley, E.** [US/US]; 1200 12th Avenue South, Suite 1200, Seattle, WA 98144-2734 (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, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (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, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: REQUEST ROUTING USING NETWORK COMPUTING COMPONENTS

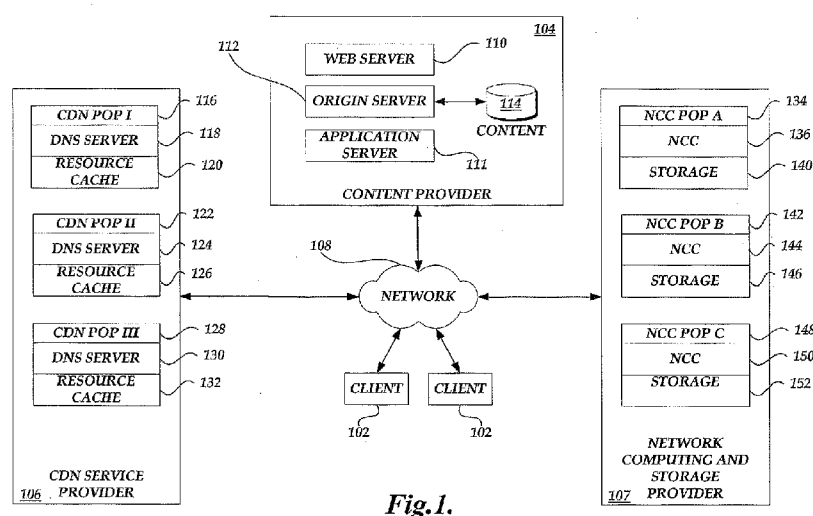


Fig. 1.

(57) Abstract: A system, method and computer-readable medium for request routing. A DNS server at a content delivery network service provider obtains a DNS query corresponding to a resource requested from a client computing device and associated with a first resource identifier. The first resource identifier includes a first portion with DNS information and a second portion with path information. The DNS server selects a network computing component for processing the requested resource based on the DNS portion of the resource identifier and transmits information identifying the selected network computing component to the client computing device.

REQUEST ROUTING USING NETWORK COMPUTING COMPONENTS

BACKGROUND

[0001] Generally described, computing devices and communication networks can be utilized to exchange information. In a common application, a computing device can request content from another computing device via the communication network. For example, a user at a personal computing device can utilize a software browser application to request a Web page from a server computing device via the Internet. In such embodiments, the user computing device can be referred to as a client computing device and the server computing device can be referred to as a content provider.

[0002] Content providers are generally motivated to provide requested content to client computing devices often with consideration of efficient transmission of the requested content to the client computing device and/or consideration of a cost associated with the transmission of the content. For larger scale implementations, a content provider may receive content requests from a high volume of client computing devices which can place a strain on the content provider's computing resources. Additionally, the content requested by the client computing devices may have a number of components, which can further place additional strain on the content provider's computing resources.

[0003] With reference to an illustrative example, a requested Web page, or original content, may be associated with a number of additional resources, such as images or videos, that are to be displayed with the Web page. In one specific embodiment, the additional resources of the Web page are identified by a number of embedded resource identifiers, such as uniform resource locators ("URLs"). In turn, software on the client computing devices typically processes embedded resource identifiers to generate requests for the content. Often, the resource identifiers associated with the embedded resources reference a computing device associated with the content provider such that the client computing device would transmit the request for the additional resources to the referenced content provider computing device. Accordingly, in order to satisfy a content request, the content provider would provide client computing devices data associated with the Web page as well as the data associated with the embedded resources.

[0004] Some content providers attempt to facilitate the delivery of requested content, such as Web pages and/or resources identified in Web pages, through the utilization of a content delivery network (“CDN”) service provider. A CDN server provider typically maintains a number of computing devices in a communication network that can maintain content from various content providers. In turn, content providers can instruct, or otherwise suggest to, client computing devices to request some, or all, of the content provider’s content from the CDN service provider’s computing devices.

[0005] With reference to previous illustrative example, the content provider can leverage a CDN service provider with the modification or substitution of resource identifiers associated with the embedded resources. Specifically, the resource identifiers can reference a computing device associated with the CDN service provider such that the client computing device would transmit the request for the additional resources to the referenced CDN service provider computing device. Typically, the content provider facilitates the utilization of a CDN provider by including CDN-provider specific resources identifiers in requested content (e.g., Web pages). This approach generally corresponds to an “offline” process implemented by the content provider in advance of receipt of a request for the original content from the client computing devices. Accordingly, modifications to resource identifiers, such as to provide alternative resources identifiers for the same CDN service provider, to provide additional information utilized by CDN service providers in processing the request for content and/or to identify alternative CDN service providers, can become inefficient as they typically require implementation of the offline process the content provider.

[0006] As with content providers, CDN providers are also generally motivated to provide requested content to client computing devices often with consideration of efficient transmission of the requested content to the client computing device and/or consideration of a cost associated with the transmission of the content. Accordingly, CDN service providers often consider factors such as latency of delivery of requested content in order to meet service level agreements or to generally improve the quality of delivery service.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The foregoing aspects and many of the attendant advantages of this invention will become more readily appreciated as the same become better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein:

[0008] FIGURE 1 is a block diagram illustrative of content delivery environment including a number of client computing devices, content provider, a content delivery network service provider, and a network computing and storage provider;

[0009] FIGURE 2 is a block diagram of the content delivery environment of FIGURE 1 illustrating the registration of a content provider with a content delivery service provider;

[0010] FIGURE 3 is a block diagram of the content delivery environment of FIGURE 1 illustrating the generation and processing of a content request from a client computing device to a content provider;

[0011] FIGURE 4 is a block diagram of the content delivery environment of FIGURE 1 illustrating one embodiment of the generation and processing of a DNS query corresponding to an embedded resource from a client computing device to a content delivery network service provider;

[0012] FIGURES 5A-5C are block diagrams of the content delivery environment of FIGURE 1 illustrating another embodiment of the generation and processing of a DNS query corresponding to an embedded resource from a client computing device to a content delivery network service provider and the subsequent generation and processing of DNS queries corresponding to a first and a second alternative resource identifier from a client computing device to a content delivery network;

[0013] FIGURES 6A-6B are block diagrams of the content delivery environment of FIGURE 1 illustrating the generation and processing of embedded resource requests from a client computing device to a content delivery network service provider;

[0014] FIGURE 7 is a flow diagram illustrative of a request routing routine implemented by a content delivery network service provider for selecting a cache server or network computing component; and

[0015] FIGURE 8 is a flow diagram illustrative of another embodiment of a request routing routine implemented by a content delivery network service provider for selecting a cache server or network computing component.

DETAILED DESCRIPTION

[0016] Generally described, the present disclosure is directed to routing of a DNS query from a client computing device to a network computing component via a content delivery network ("CDN") for processing requested content associated with the DNS query. Specifically, aspects of the disclosure will be described with regard to the routing of a client computing device DNS query within a CDN service provider domain utilizing one or more resource identifiers which include an application identifier in a DNS portion of the resource identifier. Although various aspects of the disclosure will be described with regard to illustrative examples and embodiments, one skilled in the art will appreciate that the disclosed embodiments and examples should not be construed as limiting.

[0017] FIGURE 1 is a block diagram illustrative of content delivery environment 100 for the management and processing of content requests. As illustrated in FIGURE 1, the content delivery environment 100 includes a number of client computing devices 102 (generally referred to as clients) for requesting content from a content provider and/or a CDN service provider. In an illustrative embodiment, the client computing devices 102 can correspond to a wide variety of computing devices including personal computing devices, laptop computing devices, hand-held computing devices, terminal computing devices, mobile devices, wireless devices, various electronic devices and appliances and the like. In an illustrative embodiment, the client computing devices 102 include necessary hardware and software components for establishing communications over a communication network 108, such as a wide area network or local area network. For example, the client computing devices 102 may be equipped with networking equipment and browser software applications that facilitate communications via the Internet or an intranet.

[0018] Although not illustrated in FIGURE 1, each client computing device 102 utilizes some type of local DNS resolver component, such as a DNS Name server, that generates the DNS queries attributed to the client computing device. In one embodiment, the local DNS resolver component may be provided by an enterprise network to which the client computing device 102 belongs. In another embodiment, the local DNS resolver component may be provided by an Internet Service Provider (ISP) that provides the communication network connection to the client computing device 102.

[0019] The content delivery environment 100 can also include a content provider 104 in communication with the one or more client computing devices 102 via the communication network 108. The content provider 104 illustrated in FIGURE 1 corresponds to a logical association of one or more computing devices associated with a content provider. Specifically, the content provider 104 can include a web server component 110 corresponding to one or more server computing devices for obtaining and processing requests for content (such as Web pages) from the client computing devices 102. The content provider 104 can further include an origin server component 112 and associated storage component 114 corresponding to one or more computing devices for obtaining and processing requests for network resources from the CDN service provider. The content provider 104 can still further include an application server computing device 111, such as a data streaming server, for processing streaming content requests. One skilled in the relevant art will appreciate that the content provider 104 can be associated with various additional computing resources, such as additional computing devices for administration of content and resources, DNS name servers, and the like. For example, although not illustrated in FIGURE 1, the content provider 104 can be associated with one or more DNS name server components that would be authoritative to resolve client computing device DNS queries corresponding to a domain of the content provider.

[0020] With continued reference to FIGURE 1, the content delivery environment 100 can further include a CDN service provider 106 in communication with the one or more client computing devices 102 and the content provider 104 via the communication network 108. The CDN service provider 106 illustrated in FIGURE 1 corresponds to a logical association of one or more computing devices associated with a

CDN service provider. Specifically, the CDN service provider 106 can include a number of Point of Presence (“POP”) locations 116, 122, 128 that correspond to nodes on the communication network 108. Each CDN POP 116, 122, 128 includes a DNS component 118, 124, 130 made up of a number of DNS server computing devices for resolving DNS queries from the client computers 102. Each CDN POP 116, 122, 128 also includes a resource cache component 120, 126, 132 made up of a number of cache server computing devices for storing resources from content providers and transmitting various requested resources to various client computers. The DNS components 118, 124 and 130 and the resource cache components 120, 126, 132 may further include additional software and/or hardware components that facilitate communications including, but not limited to, load balancing or load sharing software/hardware components.

[0021] In an illustrative embodiment, the DNS component 118, 124, 130 and resource cache component 120, 126, 132 are considered to be logically grouped, regardless of whether the components, or portions of the components, are physically separate. Additionally, although the CDN POPs 116, 122, 128 are illustrated in FIGURE 1 as logically associated with the CDN Provider 106, the CDN POPs will be geographically distributed throughout the communication network 108 in a manner to best serve various demographics of client computing devices 102. Additionally, one skilled in the relevant art will appreciate that the CDN service provider 106 can be associated with various additional computing resources, such additional computing devices for administration of content and resources, and the like.

[0022] With further continued reference to FIGURE 1, the content delivery environment 100 can also include a network computing and storage provider 107 in communication with the one or more client computing devices 102, the CDN service provider 106, and the content provider 104 via the communication network 108. The network computing and storage provider 107 illustrated in FIGURE 1 also corresponds to a logical association of one or more computing devices associated with a network computing and storage provider. Specifically, the network computing and storage provider 107 can include a number of Point of Presence (“POP”) locations 134, 142, 148 that correspond to nodes on the communication network 108. Each POP 134, 142, 148 includes a network computing component (NCC) 136, 144, 150 for hosting

applications, such as data streaming applications, via a number of instances of a virtual machine, generally referred to as an instance of an NCC. One skilled in the relevant art will appreciate that NCC 136, 144, 150 would include physical computing device resources and software to provide the multiple instances of a virtual machine or to dynamically cause the creation of instances of a virtual machine. Such creation can be based on a specific request, such as from a client computing device, or the NCC can initiate dynamic creation of an instance of a virtual machine on its own. Each NCC POP 134, 142, 148 also includes a storage component 140, 146, 152 made up of a number of storage devices for storing resources from content providers which will be processed by an instance of an NCC 136, 144, 150 and transmitted to various client computers. The NCCs 136, 144, 150 and the storage components 140, 146, 152 may further include additional software and/or hardware components that facilitate communications including, but not limited to, load balancing or load sharing software/hardware components for selecting instances of a virtual machine supporting a requested application and/or providing information to a DNS nameserver to facilitate request routing.

[0023] In an illustrative embodiment, NCCs 136, 144, 150 and the storage components 140, 146, 152 are considered to be logically grouped, regardless of whether the components, or portions of the components, are physically separate. Additionally, although the NCC POPs 134, 142, 148 are illustrated in FIGURE 1 as logically associated with the network computing and storage provider 107, the NCC POPs will be geographically distributed throughout the communication network 108 in a manner to best serve various demographics of client computing devices 102. Additionally, one skilled in the relevant art will appreciate that the network computing and storage provider 107 can be associated with various additional computing resources, such as additional computing devices for administration of content and resources, and the like. Even further, one skilled in the relevant art will appreciate that the components of the network computing and storage provider 107 and components of the CDN service provider 106 can be managed by the same or different entities.

[0024] With reference now to FIGURES 2-6, the interaction between various components of the content delivery environment 100 of FIGURE 1 will be illustrated. For purposes of the example, however, the illustration has been simplified such that

many of the components utilized to facilitate communications are not shown. One skilled in the relevant art will appreciate that such components can be utilized and that additional interactions would accordingly occur without departing from the spirit and scope of the present disclosure.

[0025] With reference to FIGURE 2, an illustrative interaction for registration of a content provider 104 with the CDN service provider 106 will be described. As illustrated in FIGURE 2, the CDN content registration process begins with registration of the content provider 104 with the CDN service provider 106. In an illustrative embodiment, the content provider 104 utilizes a registration application program interface ("API") to register with the CDN service provider 106 such that the CDN service provider 106 can provide content on behalf of the content provider 104. The registration API includes the identification of the origin server 112 of the content provider 104 that will provide requested resources to the CDN service provider 106.

[0026] One skilled in the relevant art will appreciate that upon identification of appropriate origin servers 112, the content provider 104 can begin to direct requests for content from client computing devices 102 to the CDN service provider 106. Specifically, in accordance with DNS routing principles, a client computing device request corresponding to a resource identifier would eventually be directed toward a CDN POP 116, 122, 128 associated with the CDN service provider 106. In the event that the resource cache component 120, 126, 132 of a selected CDN POP, or the storage component 140, 146, 152 of a subsequently selected instance of an NCC 136, 144, 150 as will be described further below, does not have a copy of a resource requested by a client computing device 102, the resource cache component 120, 126, 132, or the storage component 140, 146, 152, will request the resource from the origin server 112 previously registered by the content provider 104.

[0027] With continued reference to FIGURE 2, upon receiving the registration API, the CDN service provider 106 obtains and processes the registration information. In an illustrative embodiment, the CDN service provider 106 can then generate additional information that will be used by the client computing devices 102 as part of the content requests. The additional information can include, without limitation, client identifiers, such as client identification codes, content provider identifiers, such as content provider identification codes, executable code for processing resource

identifiers, such as script-based instructions, the like. In another embodiment, in addition or alternatively, the additional information can include file type identifiers and/or application identifiers which can include file type information, as well as information pertaining to a type of application for processing the requested content or a specific instance of an application desired for processing the requested content. Application identifiers may also include or be associated with other additional information or requirements for selecting an instance of an application for processing the requested content, such as quality of service criteria which can include information as to compression rates, processing power, processing speed, and/or bandwidth of the NCC, and the like. One skilled in the relevant art will appreciate that various types of additional information may be generated by the CDN service provider 106 and that the additional information may be embodied in any one of a variety of formats.

[0028] The CDN service provider 106 returns an identification of applicable domains for the CDN service provider (unless it has been previously provided) and any additional information to the content provider 104. In turn, the content provider 104 can then process the stored content with content provider specific information. In one example, as illustrated in FIGURE 2, the content provider 104 translates resource identifiers originally directed toward a domain of the origin server 112 to a domain corresponding to the CDN service provider. The translated URLs are embedded into requested content in a manner such that DNS queries for the translated URLs will resolve to a DNS server corresponding to the CDN service provider 106 and not a DNS server corresponding to the content provider 104. Although the translation process is illustrated in FIGURE 2, in some embodiments, the translation process may be omitted in a manner described in greater detail below.

[0029] Generally, the identification of the resources originally directed to the content provider 104 will be in the form of a resource identifier that can be processed by the client computing device 102, such as through a browser software application. In an illustrative embodiment, the resource identifiers can be in the form of a uniform resource locator ("URL"). Because the resource identifiers are included in the requested content directed to the content provider, the resource identifiers can be referred to generally as the "content provider URL." For purposes of an illustrative example, the content provider URL can identify a domain of the content provider 104

(e.g., contentprovider.com), generally referred to as a DNS portion of the URL, a name of the resource to be requested (e.g., "resource.jpg") and a path where the resource will be found (e.g., "path"), the path and resource generally referred to as a path portion of the URL. In this illustrative example, the content provider URL has the form of:

<http://www.contentprovider.com/path/resource.jpg>

[0030] During an illustrative translation process, the content provider URL is modified such that requests for the resources associated with the translated URLs resolve to a CDN POP associated with the CDN service provider 106. In one embodiment, the translated URL identifies the domain of the CDN service provider 106 (e.g., "cdnprovider.com"), the same name of the resource to be requested (e.g., "resource.jpg") and the same path where the resource will be found (e.g., "path"). Additionally, the translated URL can include additional processing information (e.g., "additional information") in the DNS portion of the URL. The translated URL would have the form of:

http://additional_information.cdnprovider.com/path/resource.jpg

[0031] In another embodiment, the information associated with the CDN service provider 106 is included in a modified URL, such as through prepending or other techniques, such that the translated URL can maintain all of the information associated with the original URL. In this embodiment, the translated URL would have the form of:

http://additional_information.cdnprovider.com/www.contentprovider.com/path/resource.jpg

[0032] With reference now to FIGURE 3, after completion of the registration and translation processes illustrated in FIGURE 2, a client computing device 102 subsequently generates a content request that is received and processed by the content provider 104, such as through the Web server 110. In accordance with an illustrative

embodiment, the request for content can be in accordance with common network protocols, such as the hypertext transfer protocol (“HTTP”). Upon receipt of the content request, the content provider 104 identifies the appropriate responsive content. In an illustrative embodiment, the requested content can correspond to a Web page that is displayed on the client computing device 102 via the processing of information, such as hypertext markup language (“HTML”), extensible markup language (“XML”), and the like. The requested content can also include a number of embedded resource identifiers, described above, that corresponds to resource objects that should be obtained by the client computing device 102 as part of the processing of the requested content. The embedded resource identifiers can be generally referred to as original resource identifiers or original URLs.

[0033] Upon receipt of the requested content, the client computing device 102, such as through a browser software application, begins processing any of the markup code included in the content and attempts to acquire the resources identified by the embedded resource identifiers. Accordingly, the first step in acquiring the content corresponds to the issuance, by the client computing device 102 (through its local DNS resolver), of a DNS query for the original URL resource identifier that results in the identification of a DNS server authoritative to the “.” and the “com” portions of the translated URL. After resolving the “.” and “com” portions of the embedded URL, the client computing device 102 then issues a DNS query for the resource URL that results in the identification of a DNS server authoritative to the “.cdnprovider” portion of the embedded URL. The issuance of DNS queries corresponding to the “.” and the “com” portions of a URL are well known and have not been illustrated.

[0034] With reference now to FIGURE 4, in an illustrative embodiment, the successful resolution of the “.cdnprovider” portion of the original URL identifies a network address, such as an IP address, of a DNS server associated with the CDN service provider 106. In one embodiment, the IP address can be a specific network address unique to a DNS server component of a POP. In another embodiment, the IP address can be shared by one or more CDN POPs. In this embodiment, a further DNS query to the shared IP address utilizes a one-to-many network routing schema, such as anycast, such that a specific POP will receive the request as a function of network topology. For example, in an anycast implementation, a DNS query issued by a client

computing device 102 to a shared IP address will arrive at a DNS server component logically having the shortest network topology distance, often referred to as network hops, from the client computing device. The network topology distance does not necessarily correspond to geographic distance. However, in some embodiments, the network topology distance can be inferred to be the shortest network distance between a client computing device 102 and a POP. It will be appreciated by one skilled in the relevant art that a number of ways exist to determine network topology distance.

[0035] With continued reference to FIGURE 4, in either of the above identified embodiments (or any other embodiment), a specific DNS server in the DNS component 118 of a CDN POP 116 receives the DNS query corresponding to the original URL from the client computing device 102. Once one of the DNS servers in the DNS component 118 receives the request, the specific DNS server attempts to resolve the request. In one illustrative embodiment as shown in FIGURE 4, a specific DNS server resolves the DNS query by identifying an IP address of a cache server component or an instance of an NCC that will process the request for the requested resource. As described above and as will be described further below in reference to FIGURES 6A and 6B, a selected resource cache component or a storage component associated with an instance of an NCC can process the request by either providing the requested resource if it is available or attempting to obtain the requested resource from another source, such as a peer cache server computing device or the origin server 112 of the content provider 104.

[0036] In further reference to FIGURE 4, the specific DNS server can utilize a variety of information in selecting a resource cache component or an instance of an NCC. In one illustrative embodiment, the DNS server can use the additional information in the DNS portion of the resource identifier (which is used to resolve the DNS query by the DNS server) to return an IP address of a resource cache component or an instance of an NCC. As generally described above, the additional information can include, without limitation, client identifiers, such as client identification codes, content provider identifiers, such as content provider identification codes, executable code for processing resource identifiers, such as script-based instructions, and the like. In addition or alternatively, the additional information can include file type identifiers and/or application identifiers which can include file type information, as well as

information pertaining to a type of hosted application for processing the requested content or a specific instance of a hosted application desired for processing the requested content.

[0037] As will be further described below in reference to FIGURE 8, where the requested content corresponds to a streaming media file, for example, the DNS server can use an application identifier included in a DNS portion of the first resource identifier to select an instance of an NCC for processing the streaming media file. In one embodiment, the application identifier can specify file type information for the content to be processed, and the DNS server selects an instance of an NCC that has a data streaming application capable of processing a request for the identified file type, e.g., an MPEG or Flash media file. In another embodiment, the application identifier can specify a type of hosted application, e.g., an Adobe Flash server streaming application or a Real Network Helix server streaming application, to be used to process the requested content. Based on that information, the DNS server resolves the DNS query by identifying an instance of an NCC that corresponds to the identified type of application for processing the requested content or that can dynamically cause creation of such an instance. Still further, in another embodiment, the application identifier can specify a specific instance of an application, e.g., Company's Flash server, specified by a content provider for example. Based on that information, the DNS server then resolves the DNS query by identifying the IP address of a specific instance of an NCC that has the application required to process the requested content. Yet further, in another embodiment, the DNS portion of the first resource identifier can have a separate file type identifier which provides the file type information for use by the DNS server in selecting an appropriate instance of an NCC device for servicing the requested content.

[0038] Even further, the DNS server can also use information obtained directly from a client computing device (such as information provided by the client computing device or ISP) or indirectly (such as inferred through a client computing device's IP address) to select a resource cache component or an instance of an NCC. Such client computing device information can, for example, be geographic information. Still further, the DNS server components can utilize network performance metrics or measurements to assign specific resource cache components or instances of an NCC.

The IP address selected by a DNS server component may correspond to a specific caching server in the resource cache or a specific instance of an NCC device. Alternatively, the IP address can correspond to a hardware/software selection component (such as a load balancer) at a specific CDN POP or NCC POP for selecting a specific cache component or instance of an NCC.

[0039] Yet still further, for selection of an instance of an NCC, the DNS server components can utilize additional selection information provided from network computing and storage provider 107 to CDN service provider 106. Such selection information can include information typically related to quality of service, such as computing capacity measurements of NCCs, compression rates, processing power, processing speed, bandwidth, and the like, which can be indirectly related to the cost associated with creating and/or using a particular instance of an NCC. This additional selection information can be provided over a communication channel between the network computing and storage provider 107 and CDN service provider 106, as generally illustrated in FIGURE 4, at a variety of times. Moreover, as will be appreciated by one skilled in the relevant art, the additional selection information may be transmitted in any of a number of ways, such as upon individual requests from the CDN service provider 106, batch processing initiated by the CDN service provider or network computing and storage provider, and the like.

[0040] In further reference to FIGURE 4, once an IP address is identified, the DNS server 118 provides the IP address of the resource cache component or the instance of an NCC to the client computing device 102 for further processing, as will be discussed further below in reference to FIGURES 6A and 6B.

[0041] In another illustrative embodiment, and as will be described further in reference to FIGURES 5A-5C, the additional information included in the DNS portion of the resource identifier can be used to identify a DNS server component which can then default to a selection of a resource cache component of the same CDN POP or an instance of an NCC of an associated NCC POP. The association can correspond to geographic or network proximity.

[0042] With reference to FIGURES 5A-5C, as an alternative to selecting a resource cache component or an instance of an NCC upon receipt of a DNS query as described in reference to FIGURE 4, the CDN service provider 106 can maintain sets of

various alternative resource identifiers which correspond to DNS server components associated with a cache component or an instance of an NCC that satisfy the criteria provided in the DNS portion of the resource identifier, e.g., in an application identifier, and that will be further selected based on additional request routing criteria. The alternative resource identifiers can be provided by the CDN service provider 106 to the client computing device 102 such that a subsequent DNS query on the alternative resource identifier will resolve to a different DNS server component within the CDN service provider's network. In an illustrative embodiment, the alternative resource identifiers are in the form of one or more canonical name ("CNAME") records. In one embodiment, each CNAME record identifies a domain of the CDN service provider 106 (e.g., "cdnprovider.com" or "cdnprovider-1.com"). As will be explained in greater detail below, the domain in the CNAME does not need to be the same domain found in original URL or in a previous CNAME record. Additionally, each CNAME record includes additional information, such as request routing information, (e.g., "request routing information"). An illustrative CNAME record can have the form of:

request_routing_information.cdnprovider.com

[0043] In an illustrative embodiment, the CNAME records are generated and provided by the DNS servers to identify a more appropriate DNS server of the CDN service provider 106.

[0044] In one embodiment, if the first DNS server determines that it is the appropriate DNS server (e.g., it can resolve the DNS query by providing an IP address of a cache component or instance of an NCC), the DNS server identifies a default, associated cache component or instance of an NCC POP to resolve the request. Alternatively, even though the DNS server can resolve the request, the DNS server can determine if another DNS server can also service the request based on additional request routing criteria, which will be described in detail below.

[0045] In one example, the CDN service provider 106 can also attempt to direct a DNS query to DNS servers according to geographic criteria. The geographic criteria can correspond to geographic-based regional service plans contracted between the CDN

service-provider 106 and the content provider 104 in which various CDN service provider 106 POPs are grouped into geographic regions. Accordingly, a client computing device 102 DNS query received in a region not corresponding to the content provider's regional plan may be better processed by a DNS server in a region corresponding to the content provider's regional plan. In this example, the DNS server component 118 may also obtain geographic information from the client directly (such as information provided by the client computing device or ISP) or indirectly (such as inferred through a client computing device's IP address).

[0046] In another example, the CDN service provider 106 can also attempt to direct a DNS query to DNS servers according to service level criteria. The service level criteria can correspond to service or performance metrics contracted between the CDN service provider 106 and the content provider 104. Examples of performance metrics can include latencies of data transmission between the CDN service provider POPs and the client computing devices 102, total data provided on behalf of the content provider 104 by the CDN service provider POPs, error rates for data transmissions, and the like.

[0047] In still a further example, the CDN service provider 106 can attempt to direct a DNS query to DNS servers according to network performance criteria. The network performance criteria can correspond to measurements of network performance for transmitting data from the CDN service provider POPs to the client computing device 102. Examples of network performance metrics can include network data transfer latencies (measured by the client computing device or the CDN service provider 106, network data error rates, and the like.

[0048] In yet a further example, the receiving DNS server can attempt to direct a DNS query to another appropriate DNS query for load balancing or load sharing purposes. The receiving DNS may obtain an identification of other CDN POPs that define an order for such load balancing/load sharing. The list can be dynamic or periodically updated based on current network conditions.

[0049] In yet another example, the CDN service provider 106 can attempt to direct a DNS query associated with a resource request for processing by an instance of an NCC to DNS servers according to selection information criteria provided by the network computing and storage provider 107. As similarly set forth above, such

selection information can include information typically related to quality of service, such as computing capacity measurements of NCCs, compression rates, processing power, processing speed, bandwidth, and the like, which can be indirectly related to the cost associated with creating and/or using a particular instance of an NCC. This selection information can be provided over a communication channel between the network computing and storage provider 107 and CDN service provider 106 at a variety of times. Moreover, as will be appreciated by one skilled in the relevant art, the selection information may be transmitted in any of a number of ways, such as upon individual requests from the CDN service provider 106, batch processing initiated by the CDN service provider or network computing and storage provider, and the like.

[0050] In accordance with an illustrative embodiment, the DNS server maintains a data store that defines CNAME records for various incoming DNS queries. If a DNS query corresponding to a particular DNS query matches an entry in the data store, the DNS server returns a CNAME record as defined in the data store. In an illustrative embodiment, the data store can include multiple CNAME records corresponding to a particular DNS query. The multiple CNAME records would define a set of potential candidates that can be returned to the client computing device. In such an embodiment, the DNS server, either directly or via a network-based service, selects one of the CNAME records defined in the data store as more appropriate routing information based on logic that factors information contained in a DNS portion of the first resource identifier, as described above. It will be appreciated by one skilled in the art and others that the DNS server can implement further additional logic in selecting an appropriate CNAME from a set of possible of CNAMEs. In an illustrative embodiment, each DNS server component 118, 124, 130 maintains the same data stores that define CNAME records, which can be managed centrally by the CDN service provider 106. Alternatively, each DNS server component 118, 124, 130 can have POP specific data stores that define CNAME records, which can be managed centrally by the CDN service provider 106 or locally at the POP 116, 122, 128. Still further, each DNS server computing device within the DNS server components 118, 124, 130 can utilize shared data stores managed by a respective POP or a local data store specific to an individual DNS server computing device.

[0051] The returned CNAME can also include request routing information that is different from or in addition to the information provided in the URL/CNAME of the current DNS query. For example, if the CNAME selection is further based on a service level plan, in addition to information associated with an application identifier included in the first resource identifier, a specific further identifier can be included in the “request_routing_information” portion of the CNAME record. An illustrative CNAME record in this specific example can have the form of:

serviceplan.appID.cdnprovider.com

[0052] In another embodiment, request routing information can be found in the identification of a CDN service provider 106 domain different from the domain found in the current URL/CNAME. For example, if the CNAME is based on a regional plan, a specific regional plan domain (e.g., “cdnprovider-region1.com”) could be used in the domain name portion of the specific CNAME record. One skilled in the relevant art will appreciate that additional or alternative techniques and/or combination of techniques may be used to include the additional request routing information in the CNAME record that is selected by the DNS server component 118.

[0053] With continued reference to FIGURE 5A, in accordance with the embodiments set forth above, the DNS server may select (or otherwise obtain) a CNAME record that is intended to resolve to a more appropriate DNS server of the CDN service provider 106. It may be possible, however, that the same DNS server would also be authoritative for the subsequent DNS query for the CNAME to be provided to the client computing device. For example, a specific DNS server may be authoritative for both a specific regional plan and a service level plan. Thus, returning a CNAME would still result in the DNS query arriving at the same DNS query (which may also be due in part to the client computing device’s geography). In such an embodiment, the DNS server, such as DNS server component 118, may choose to resolve the future DNS query in advance.

[0054] With reference now to FIGURE 5B, upon receipt of the CNAME from the DNS server component 118, the client computing device 102 generates a subsequent DNS query corresponding to the CNAME. As previously discussed with

regard to FIGURE 4, the DNS query process could first start with DNS queries for the “.” and “com” portions, followed by a query for the “cdnprovider” portion of the CNAME. To the extent, however, that the results of a previous DNS queries can be cached (and remain valid), the client computing device 102 can utilize the cached information and does not need to repeat the entire process. However, at some point, depending on whether the CNAME provided by DNS server component 118 (FIGURE 5A) and the previous URL/CNAME share common CDN service provider domains, the current CNAME DNS query resolves to a different POP provided by the CDN service provider 106. As illustrated in FIGURE 5B, the DNS server component 124 of CDN POP 122 receives the current CNAME based on the different information in the current CNAME previously provided by the DNS server component 118. As previously described, the DNS server component 124 can then determine whether to resolve the DNS query on the CNAME with an IP address of a cache component or an instance of an NCC that will process the content request or whether to provide another alternative resource identifier selected in the manners described above.

[0055] For purposes of illustration, assume that the DNS server component 124 determines that the DNS query corresponding to the current CNAME (provided by DNS server component 118) also corresponds to a CNAME record in its data store. In such an example, the DNS server component 124 would do any necessary processing to select a specific CNAME as described above and return the CNAME to the client computing device. With reference now to FIGURE 5C, the client computing device 102 would now transmit a second subsequent DNS query corresponding to the CNAME provided by DNS server component 124 (FIGURE 5B). In accordance with DNS query processes already described, the DNS query would illustratively be received by the DNS server component 130 of CDN POP 128. Again, the DNS server component 130 can then determine whether to resolve the DNS query on the CNAME with an IP address of a cache component or an instance of a NCC that will process the content request or whether to provide another alternative resource identifier selected in the manners described above. In this example, the DNS server component 130 returns an IP address.

[0056] With continued reference to FIGURE 5C, in an illustrative embodiment, the DNS server components, such as DNS server component 130, can utilize a variety of information in selecting a resource cache component or an instance of an NCC, as similarly set forth above. In one example, the DNS server component can default to a selection of a resource cache component of the same CDN POP or an instance of an NCC associated with the same CDN POP.

[0057] With reference now to FIGURES 6A, in an illustrative example, assume that the DNS server component 130 has selected the resource cache component 120 of POP 116. Upon receipt of the IP address for the resource cache component 120, the client computing device 102 transmits a request for the content to the resource cache component 120. The resource cache component 120 processes the request in a manner described above and the requested content is transmitted to the client computing device 102.

[0058] With reference now to FIGURES 6B, in a further illustrative example, assume that the DNS server component 130 has selected an instance of NCC 136 of NCC POP 134. Upon receipt of the IP address for the instance of NCC 136, the client computing device 102 transmits a request for the requested content to the instance of NCC 136. The instance of NCC 136 processes the request and the requested content is transmitted to the client computing device 102. For example, for streaming media, the instance of the NCC will begin streaming the content.

[0059] With reference now to FIGURE 7, a request routine 700 implemented by the CDN provider 106 will be described. One skilled in the relevant art will appreciate that actions/steps outlined for routine 700 may be implemented by one or many computing devices/components that are associated with the CDN service provider 106. Accordingly, routine 700 has been logically associated as being performed by the CDN service provider 106.

[0060] At block 702, one of the DNS server components 118, 124, 130 receives a DNS query corresponding to a resource identifier. As previously discussed, the resource identifier can be a URL that has been embedded in content requested by the client computing device 102 and previously provided by the content provider 104. Alternatively, the resource identifier can also correspond to a CNAME provided by a content provider DNS server in response to a DNS query previously received from the

client computing device 102. At decision block 704, a test is conducted to determine whether the current DNS server is authoritative to resolve the DNS query. In an illustrative embodiment, the DNS server can determine whether it is authoritative to resolve the DNS query if there are no CNAME records corresponding to the received resource identifier. Additional methodologies may also be practiced to determine whether the DNS server is authoritative.

[0061] If the current DNS server is authoritative (including a determination that the same DNS server will be authoritative for subsequent DNS queries), the current DNS server resolves the DNS query by returning the IP address of a default cache server component or a default instance of an NCC. In a non-limiting manner, a number of methodologies for selecting an appropriate resource cache component or instance of an NCC have been previously discussed. Additionally, as described above, the IP address can correspond to a hardware/software selection component (such as a load balancer) at a specific CDN POP or NCC POP for selecting a specific cache component or instance of an NCC.

[0062] Alternatively, if at decision block 704 the DNS server is not authoritative, at block 808, the DNS server component selects and transmits an alternative resource identifier. As described above, the DNS server component can utilize a data store to identify an appropriate CNAME as a function of the current DNS query, including for example the application identifier in the DNS portion of the URL or CNAME. Additionally, the DNS server component can also implement additional logical processing to select from a set of potential CNAMEs. At block 710, different DNS server components 118, 124, 130 receive a DNS query corresponding to the CNAME. The routine 700 then returns to decision block 704 and continues to repeat as appropriate.

[0063] With reference now to FIGURE 8, a request routine 800 implemented by the CDN provider 106 will be described. One skilled in the relevant art will appreciate that actions/steps outlined for routine 800 may be implemented by one or many computing devices/components that are associated with the CDN service provider 106. Accordingly, routine 800 has been logically associated as being performed by the CDN service provider 106.

[0064] At block 802, one of the DNS server components 118, 124, 130 receives a DNS query corresponding to a resource identifier. Again, as previously discussed, the resource identifier can be a URL that has been embedded in content requested by the client computing device 102 and previously provided by the content provider 104. Alternatively, the resource identifier can also correspond to a CNAME provided by a content provider DNS server in response to a DNS query previously received from the client computing device 102. At block 804, the current DNS server selects a resource cache component or an instance of an NCC to service the request associated with the DNS query. As described above, the DNS server makes this selection as a function of a DNS portion of the resource identifier associated with the DNS query. As an example, the DNS portion of the resource identifier can specify information associated with the file type of the requested resource or an application type or specific instance of an application for processing the requested resource. As further set forth above, a number of other factors may additionally be taken into consideration for selecting the appropriate cache component or instance of an NCC for servicing the resource request.

[0065] Next, at block 806, the DNS server provides information to the client computing device identifying the selected resource cache component or instance of an NCC. In one embodiment, this information includes the IP address of the resource cache component or instance of an NCC. Thereafter, the client computing device can use that information to request and receive the requested resource. Routine 800 ends at block 806.

[0066] While illustrative embodiments have been disclosed and discussed, one skilled in the relevant art will appreciate that additional or alternative embodiments may be implemented within the spirit and scope of the present invention. Additionally, although many embodiments have been indicated as illustrative, one skilled in the relevant art will appreciate that the illustrative embodiments do not need to be combined or implemented together. As such, some illustrative embodiments do not need to be utilized or implemented in accordance with scope of variations to the present disclosure.

[0067] Conditional language, such as, among others, "can," "could," "might," or "may," unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other

embodiments do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without user input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment.

[0068] Any process descriptions, elements, or blocks in the flow diagrams described herein and/or depicted in the attached FIGURES should be understood as potentially representing modules, segments, or portions of code which include one or more executable instructions for implementing specific logical functions or steps in the process. Alternate implementations are included within the scope of the embodiments described herein in which elements or functions may be deleted, executed out of order from that shown or discussed, including substantially concurrently or in reverse order, depending on the functionality involved, as would be understood by those skilled in the art. It will further be appreciated that the data and/or components described above may be stored on a computer-readable medium and loaded into memory of the computing device using a drive mechanism associated with a computer readable storing the computer executable components such as a CD-ROM, DVD-ROM, or network interface further, the component and/or data can be included in a single device or distributed in any manner. Accordingly, general purpose computing devices may be configured to implement the processes, algorithms and methodology of the present disclosure with the processing and/or execution of the various data and/or components described above.

[0069] Clause 1. A method for request routing comprising:

obtaining a DNS query from a client computing device at a first DNS server, wherein the DNS query corresponds to a requested resource associated with an original resource identifier provided by a content provider and wherein the first DNS server corresponds to a content delivery network service provider that is different from the content provider;

determining whether the first DNS server is authoritative to the DNS query;

obtaining an alternative resource identifier based on an application identifier included in the original resource identifier if the DNS server is determined not to be authoritative, wherein the alternative resource identifier includes information for causing a DNS query to resolve to a domain corresponding to the content delivery service provider and request routing information not included in the original resource identifier;

transmitting the alternative resource identifier to the client computing device;

obtaining a subsequent DNS query from the client computing device at a second DNS server, wherein the subsequent DNS query corresponds to the alternative resource identifier provided by the content delivery network service provider and wherein the second DNS server corresponds to the content delivery network service provider that is different from the content provider;

determining whether the second DNS server is authoritative to the subsequent DNS query; and

selecting a network computing component for processing content associated with the original resource request if the second DNS server is authoritative to the second DNS query; and

transmitting information identifying the selected network computing component if the second DNS server is determined to be authoritative.

[0070] Clause 2. The method as recited in Clause 1, wherein the original resource identifier corresponds to a uniform resource locator provided by the content provider such that the DNS query resolves to a domain corresponding to the content delivery network service provider.

[0071] Clause 3. The method as recited in Clause 1, wherein the alternative resource identifier corresponds to a canonical name record identifier.

[0072] Clause 4. The method as recited in Clause 1, wherein the original resource identifier and the alternative resource identifier include information identifying a same domain corresponding to the content delivery service provider.

[0073] Clause 5. The method as recited in Clause 1, wherein the original resource identifier includes information identifying a first domain corresponding to the content delivery service provider and wherein the alternative resource identifier

includes information identifying a second domain corresponding to the content delivery service provider.

[0074] Clause 6. The method as recited in Clause 1, wherein the network computing component comprises an application for processing the requested resource.

[0075] Clause 7. The method as recited in Clause 6, wherein the application is a data streaming application.

[0076] Clause 8. The method as recited in Clause 1, wherein the network computing component is operable to dynamically cause the creation of an instance of a virtual machine for processing the requested resource.

[0077] Clause 9. The method as recited in Clause 1, wherein the application identifier includes information associated with the file type of the requested resource.

[0078] Clause 10. The method as recited in Clause 1, wherein the application identifier includes information associated with a type of application for processing the requested resource.

[0079] Clause 11. The method as recited in Clause 1, wherein the application identifier includes information associated with an instance of an application for processing the requested resource.

[0080] Clause 12. The method as recited in Clause 1, wherein the application identifier includes network computing component selection information associated with a type of an application or an instance of an application for processing the requested resource.

[0081] Clause 13. The method as recited in Clause 12, wherein the selection information includes quality of service information.

[0082] Clause 14. The method as recited in Clause 1, wherein obtaining an alternative resource identifier is further based on information associated with a geographic location of the client computing device.

[0083] Clause 15. The method as recited in Clause 1, wherein obtaining an alternative resource identifier is further based on information associated with the content provider, wherein the information includes regional service information and wherein the alternative resource identifier includes information identifying a regional service plan provided by the content delivery network service provider.

[0084] Clause 16. The method as recited in Clause 1, wherein obtaining an alternative resource identifier is further based on information associated with the content provider, wherein the information includes service level information and wherein the alternative resource identifier includes information identifying a service level plan provided by the content delivery network service provider.

[0085] Clause 17. The method as recited in Clause 1, wherein obtaining an alternative resource identifier is further based on information associated with network computing resources, wherein the information includes network computing component selection information.

[0086] Clause 18. The method as recited in Clause 17, wherein the network computing component selection information includes performance metrics associated with one or more network computing components.

[0087] Clause 19. The method as recited in Clause 1, further comprising:
obtaining a request from the client computing device for the requested content; and
processing and transmitting the requested content from the selected network computing component in response to the request from the client computing device.

[0088] Clause 20. The method as recited in Clause 1, wherein the information identifying the selected network computing component includes a network address corresponding to the selected network computing component.

[0089] Clause 21. A method for request routing comprising:
obtaining a DNS query from a client computing device at a first DNS server, wherein the DNS query corresponds to a requested resource associated with a first resource identifier and wherein the first DNS server corresponds to a content delivery network service provider;

determining that the first resource identifier is associated with an alternative resource identifier;

transmitting the alternative resource identifier to the client computing device, wherein the alternative resource identifier includes information for causing a DNS query to resolve to a domain corresponding to the content delivery service

provider and supplemental request routing information not included in the first resource identifier;

obtaining a subsequent DNS query from the client computing device at a second DNS server, wherein the subsequent DNS query corresponds to the alternative resource identifier provided by the content delivery network service provider and wherein the second DNS server corresponds to the content delivery network service provider;

resolving the subsequent DNS query, at the second DNS server, to identify a network computing component for processing and providing content associated with the original resource request; and

transmitting information identifying the identified network computing component to the client computing device.

[0090] Clause 22. The method as recited in Clause 21, wherein first resource identifier corresponds to a uniform resource locator such that the DNS query resolves to a domain corresponding to the content delivery network service provider.

[0091] Clause 23. The method as recited in Clause 21, wherein the alternative resource identifier corresponds to a canonical name record identifier.

[0092] Clause 24. The method as recited in Clause 21, wherein the network computing component comprises an application for processing the requested resource.

[0093] Clause 25. The method as recited in Clause 24, wherein the application is a data streaming application.

[0094] Clause 26. The method as recited in Clause 21, wherein the network computing component is operable to dynamically cause the creation of an instance of a virtual machine for processing the requested resource.

[0095] Clause 27. The method as recited in Clause 21, where determining that the first resource identifier is associated with an alternative resource identifier comprises:

obtaining resource routing information corresponding to an application identifier included in the first resource identifier;

selecting an alternative resource identifier as a function of the resource routing information corresponding to the application identifier.

[0096] Clause 28. The method as recited in Clause 27, wherein the application identifier includes information associated with the file type of the requested resource and wherein the alternative resource identifier is selected as a function of the file type of the requested resource.

[0097] Clause 29. The method as recited in Clause 27, wherein the application identifier includes information associated with a type of application for processing the requested resource and wherein the alternative resource identifier is selected as a function of the application type.

[0098] Clause 30. The method as recited in Clause 27, wherein the application identifier includes information identifying an instance of an application for processing the requested resource and wherein the alternative resource identifier is selected as a function of the application instance.

[0099] Clause 31. The method as recited in Clause 27, wherein the application identifier includes network computing component selection information associated with a type of an application or an instance of an application for processing the requested resource.

[0100] Clause 32. The method as recited in Clause 31, wherein the selection information includes quality of service information.

[0101] Clause 33. The method as recited in Clause 27 further comprising obtaining resource routing information corresponding to the client computing device associated with the DNS query, wherein the information corresponding to the client computing device is associated with a geographic location of the client computing device and wherein selecting an alternative resource identifier is further based on information associated with a geographic location of the client computing device.

[0102] Clause 34. A method for request routing comprising:

obtaining a DNS query from a client computing device at a first DNS server, wherein the DNS query corresponds to a requested resource associated with a first resource identifier and wherein the first DNS server corresponds to a content delivery network service provider;

selecting a network computing component for processing the requested resource from a plurality of network computing components based on the first resource identifier; and

transmitting information identifying the selected network computing component from the first DNS server to the client computing device.

[0103] Clause 35. The method as recited in Clause 34, wherein the network computing component comprises an application for processing the requested resource.

[0104] Clause 36. The method as recited in Clause 34, wherein the application is a data streaming application.

[0105] Clause 37. The method as recited in Clause 34, wherein the network computing component is operable to dynamically cause the creation of an instance of a virtual machine for processing the requested resource.

[0106] Clause 38. The method as recited in Clause 34, wherein the first resource identifier includes a first portion for identifying information associated with a service plan provided by the content provider and wherein a network computing component is selected based on the first portion.

[0107] Clause 39. The method as recited in Clause 34, wherein the first resource identifier includes a first portion for identifying information associated with the file type of the requested resource and wherein a network computing component is selected at the first DNS server based on the file type.

[0108] Clause 40. The method as recited in Clause 34, wherein the first resource identifier includes a DNS portion and a path portion, wherein each portion includes a file identifier and wherein a network computing component is selected at the first DNS server based on the file identifier included in the DNS portion of the first resource identifier.

[0109] Clause 41. The method as recited in Clause 34, wherein the first resource identifier includes a first portion for identifying information associated with network computing component selection information, wherein the selection information includes quality of service information.

[0110] Clause 42. The method as recited in Clause 34, wherein selecting a network computing component is further based on information associated with the geographic location of the client computing device.

[0111] Clause 43. The method as recited in Clause 34, wherein selecting a network computing component is further based on performance metrics associated with one or more network computing components.

[0112] Clause 44. A method for request routing comprising:

obtaining a DNS query from a client computing device at a first DNS server, wherein the DNS query corresponds to a requested resource associated with a first resource identifier, wherein the first resource identifier includes a first portion with DNS information and a second portion with path information, and wherein the DNS server corresponds to a content delivery network service provider;

selecting a network computing component for processing the requested resource based on the first portion of the first resource identifier; and

transmitting information identifying the selected network computing component from the first DNS server to the client computing device.

[0113] Clause 45. The method as recited in Clause 44, wherein the network computing component comprises an application for processing the requested resource.

[0114] Clause 46. The method as recited in Clause 44, wherein the application is a data streaming application.

[0115] Clause 47. The method as recited in Clause 44, wherein the network computing component is operable to dynamically cause the creation of an instance of a virtual machine for processing the requested resource.

[0116] Clause 48. The method as recited in Clause 44, wherein the first and second portions of the first resource identifier each include a file identifier.

[0117] Clause 49. The method as recited in Clause 44, wherein the first portion of the first resource identifier includes a file identifier and wherein the network computing component is selected as a function of the file identifier.

[0118] Clause 50. The method as recited in Clause 44, wherein the first portion of the first resource identifier includes an application identifier and wherein the network computing component is selected as a function of the application identifier.

[0119] Clause 51. The method as recited in Clause 50, wherein the application identifier identifies a type of application for processing the requested resource.

[0120] Clause 52. The method as recited in Clause 50, wherein the application identifier identifies an instance of an application for processing the requested resource.

[0121] Clause 53. The method as recited in Clause 50, wherein the application identifier includes network computing component selection information associated with a type of an application or an instance of an application for processing the requested resource.

[0122] Clause 54. The method as recited in Clause 53, wherein the selection information includes quality of service information.

[0123] Clause 55. The method as recited in Clause 44, wherein selecting a network computing component is further based on information associated with the geographic location of the client computing device.

[0124] Clause 56. The method as recited in Clause 44, wherein the information identifying the selected network computing component is the IP address of the selected network computing component.

[0125] It should be emphasized that many variations and modifications may be made to the above-described embodiments, the elements of which are to be understood as being among other acceptable examples. All such modifications and variations are intended to be included herein within the scope of this disclosure and protected by the following claims.

WHAT IS CLAIMED IS:

1. A method for request routing comprising:
 - obtaining a DNS query from a client computing device at a first DNS server, wherein the DNS query corresponds to a requested resource associated with a first resource identifier and wherein the first DNS server corresponds to a content delivery network service provider;
 - determining that the first resource identifier is associated with an alternative resource identifier;
 - transmitting the alternative resource identifier to the client computing device, wherein the alternative resource identifier includes information for causing a DNS query to resolve to a domain corresponding to the content delivery service provider and supplemental request routing information not included in the first resource identifier;
 - obtaining a subsequent DNS query from the client computing device at a second DNS server, wherein the subsequent DNS query corresponds to the alternative resource identifier provided by the content delivery network service provider and wherein the second DNS server corresponds to the content delivery network service provider;
 - resolving the subsequent DNS query, at the second DNS server, to identify a network computing component for processing and providing content associated with the original resource request; and
 - transmitting information identifying the identified network computing component to the client computing device.
2. The method as recited in Claim 1, wherein first resource identifier corresponds to a uniform resource locator such that the DNS query resolves to a domain corresponding to the content delivery network service provider.
3. The method as recited in Claim 1, wherein the alternative resource identifier corresponds to a canonical name record identifier.
4. The method as recited in Claim 1, wherein the network computing component comprises an application for processing the requested resource.
5. The method as recited in Claim 4, wherein the application is a data streaming application.

6. The method as recited in Claim 1, wherein the network computing component is operable to dynamically cause the creation of an instance of a virtual machine for processing the requested resource.

7. The method as recited in Claim 1, wherein determining that the first resource identifier is associated with an alternative resource identifier comprises:

obtaining resource routing information corresponding to an application identifier included in the first resource identifier;

selecting an alternative resource identifier as a function of the resource routing information corresponding to the application identifier.

8. A method for request routing comprising:

obtaining a DNS query from a client computing device at a first DNS server, wherein the DNS query corresponds to a requested resource associated with a first resource identifier and wherein the first DNS server corresponds to a content delivery network service provider;

selecting a network computing component for processing the requested resource from a plurality of network computing components based on the first resource identifier; and

transmitting information identifying the selected network computing component from the first DNS server to the client computing device.

9. The method as recited in Claim 8, wherein the network computing component comprises an application for processing the requested resource.

10. The method as recited in Claim 9, wherein the first resource identifier includes a first portion for identifying information associated with a service plan provided by the content provider and wherein a network computing component is selected based on the first portion.

11. The method as recited in Claim 9, wherein the first resource identifier includes a DNS portion and a path portion, wherein each portion includes a file identifier and wherein a network computing component is selected at the first DNS server based on the file identifier included in the DNS portion of the first resource identifier.

12. The method as recited in Claim 9, wherein selecting a network computing component is further based on performance metrics associated with one or more network computing components.

13. A method for request routing comprising:

obtaining a DNS query from a client computing device at a first DNS server, wherein the DNS query corresponds to a requested resource associated with a first resource identifier, wherein the first resource identifier includes a first portion with DNS information and a second portion with path information, and wherein the DNS server corresponds to a content delivery network service provider;

selecting a network computing component for processing the requested resource based on the first portion of the first resource identifier; and

transmitting information identifying the selected network computing component from the first DNS server to the client computing device.

14. The method as recited in Claim 13, wherein the first portion of the first resource identifier includes a file identifier and wherein the network computing component is selected as a function of the file identifier.

15. The method as recited in Claim 13, wherein the first portion of the first resource identifier includes an application identifier and wherein the network computing component is selected as a function of the application identifier.

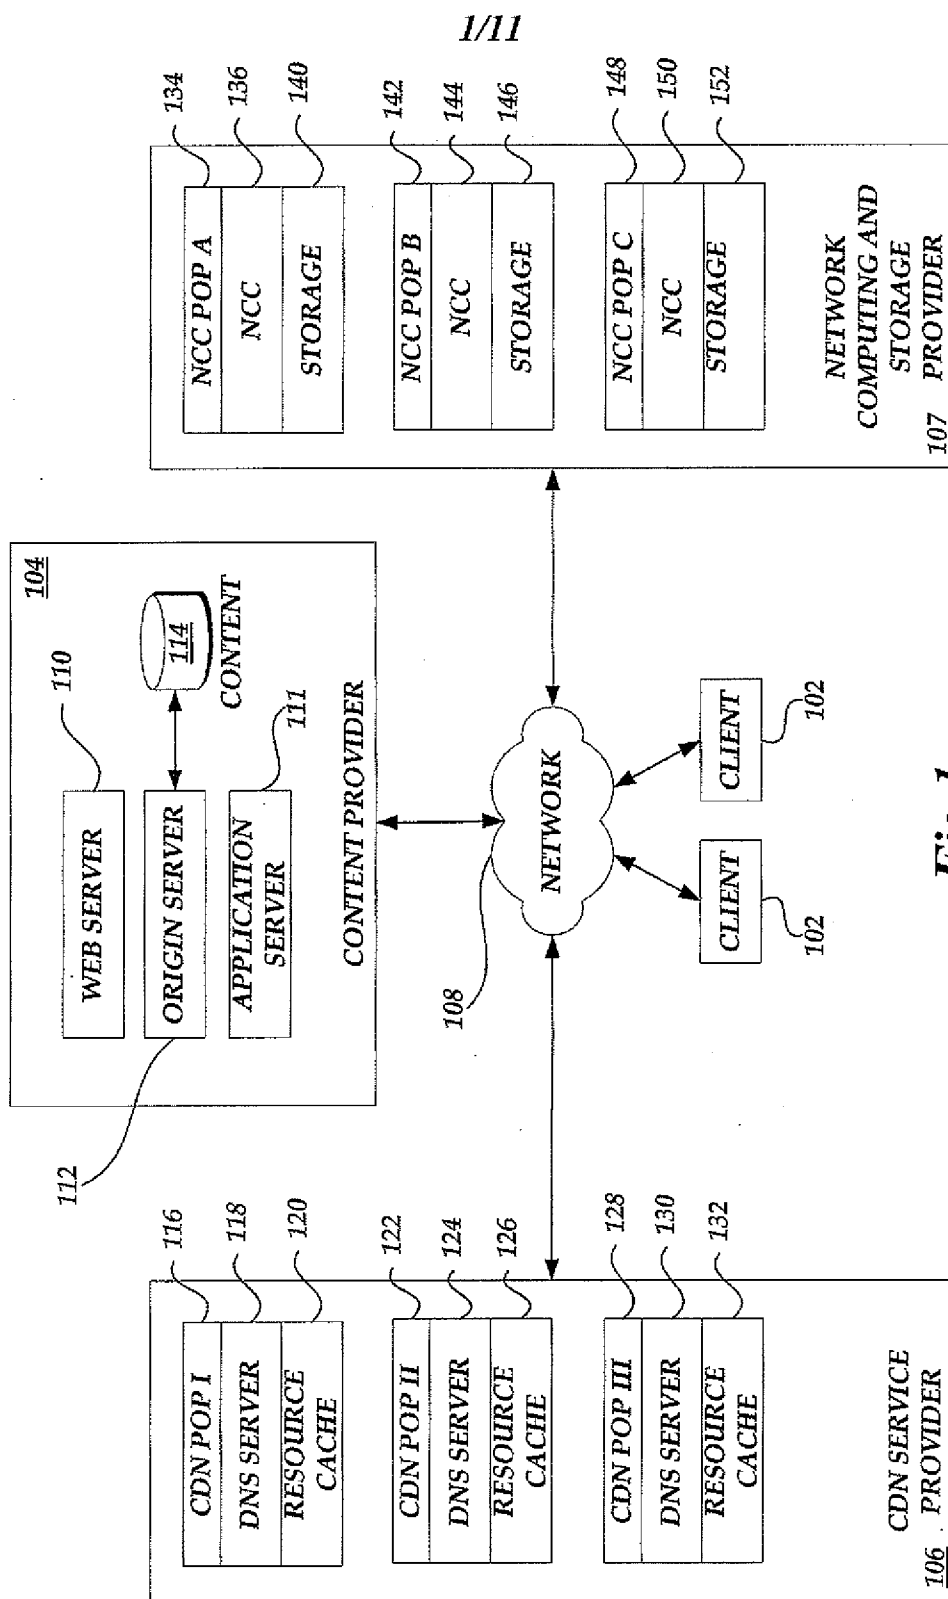


Fig. 1.

2/11

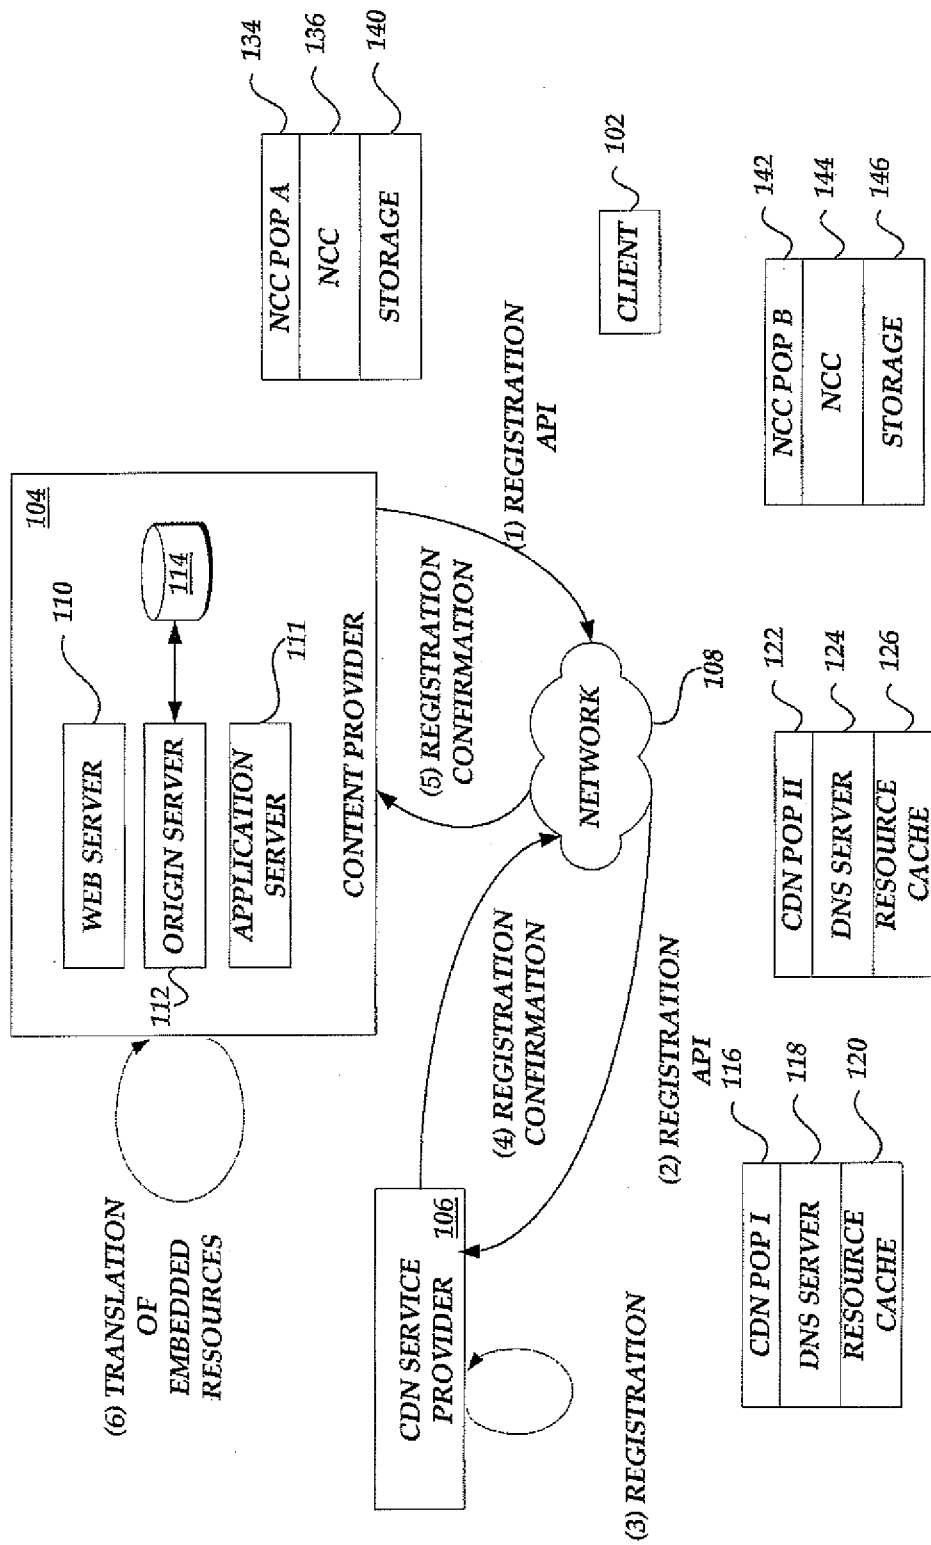


Fig.2.

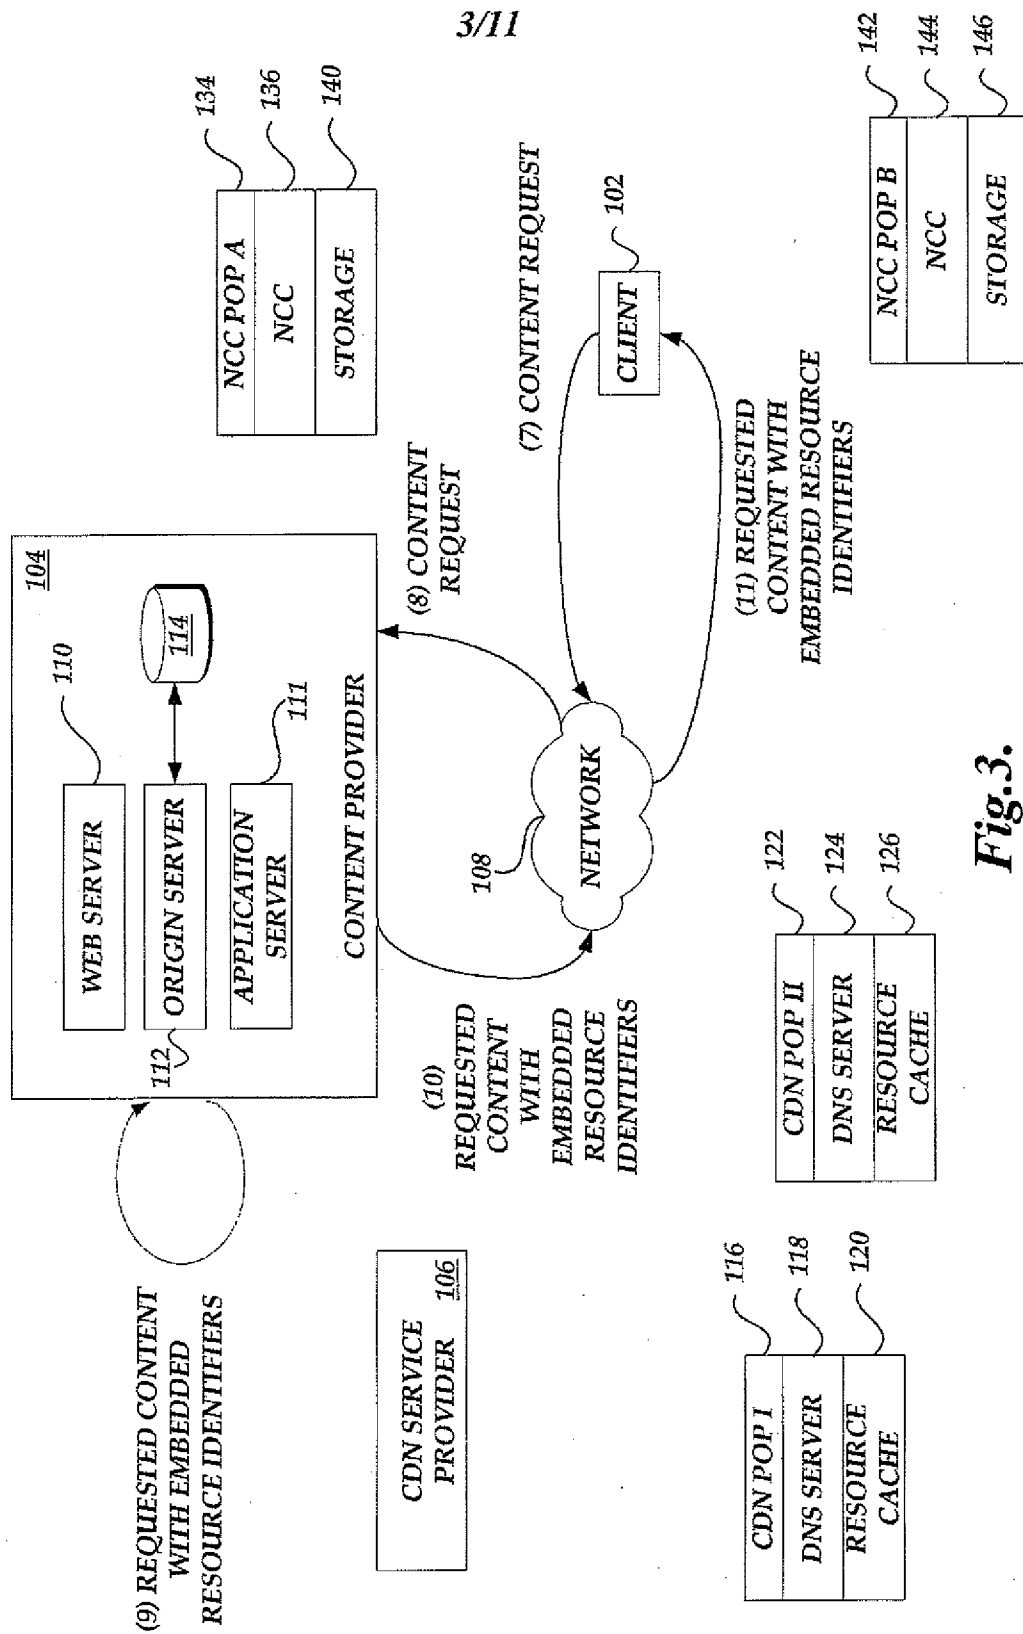


Fig.3.

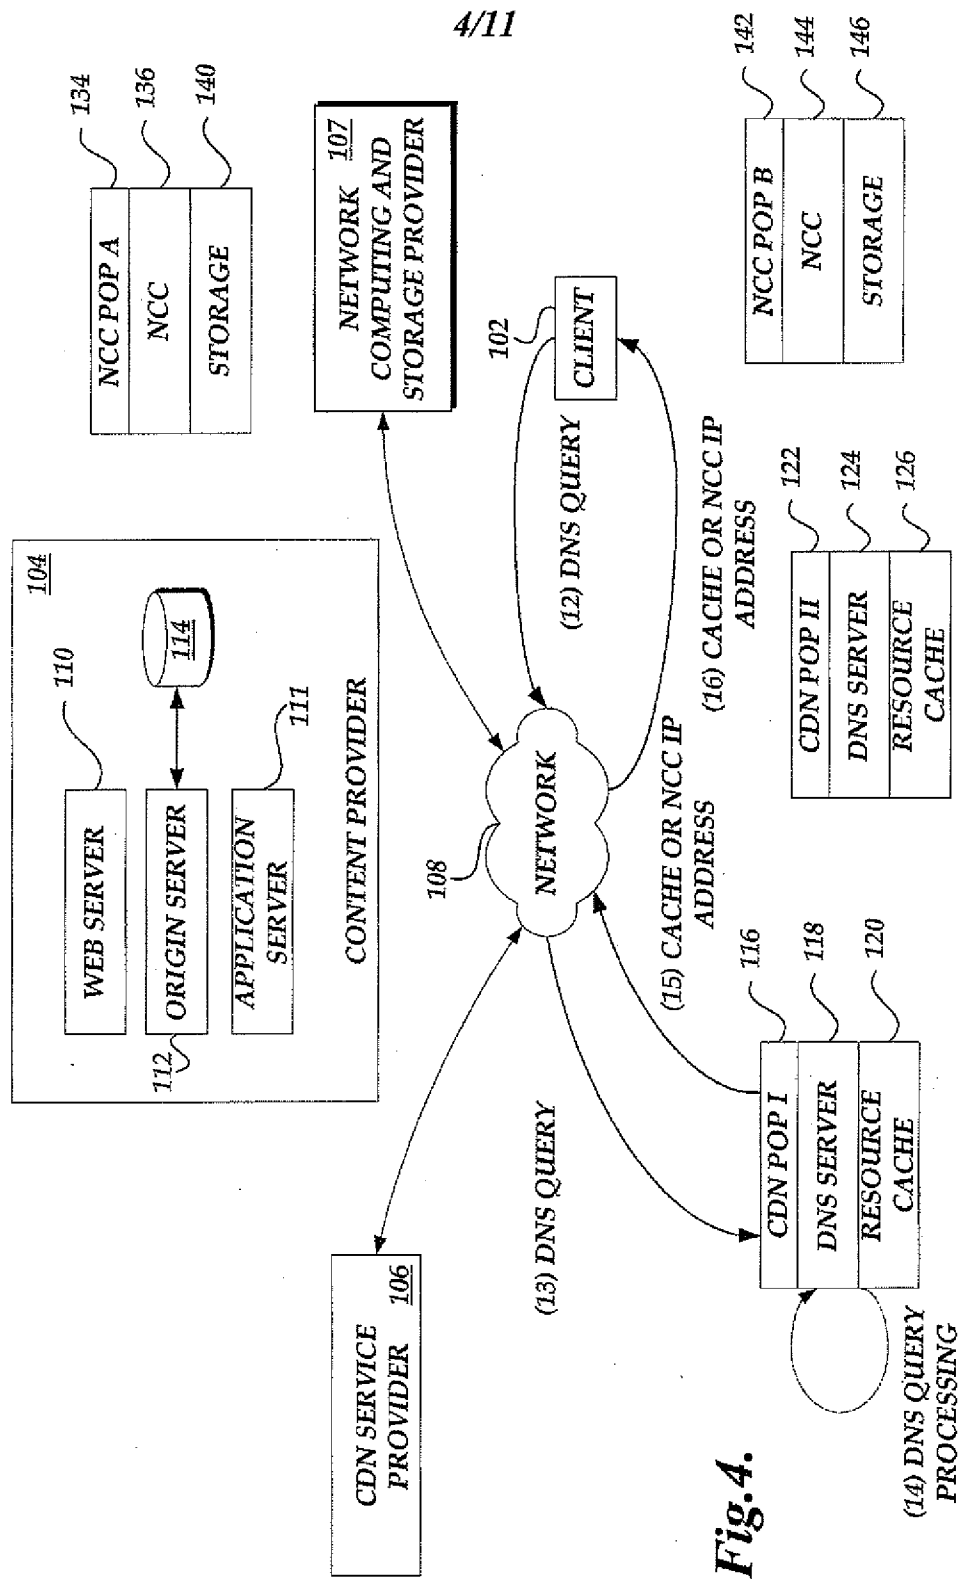


Fig. 4.

5/11

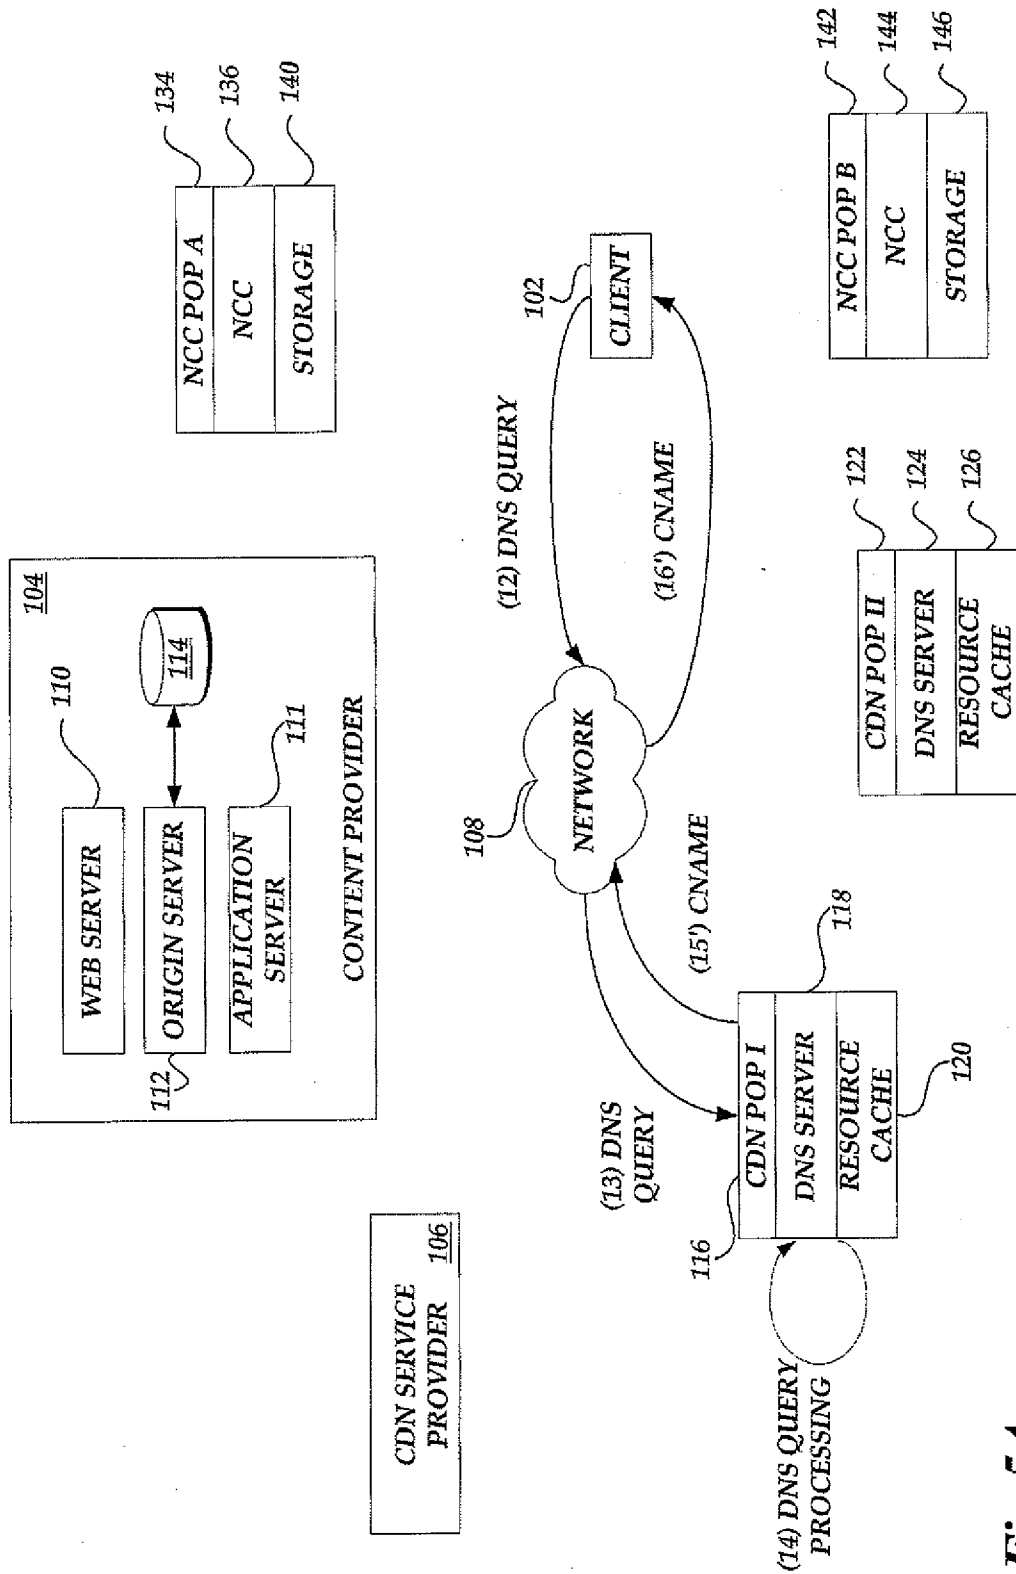


Fig. 5A.

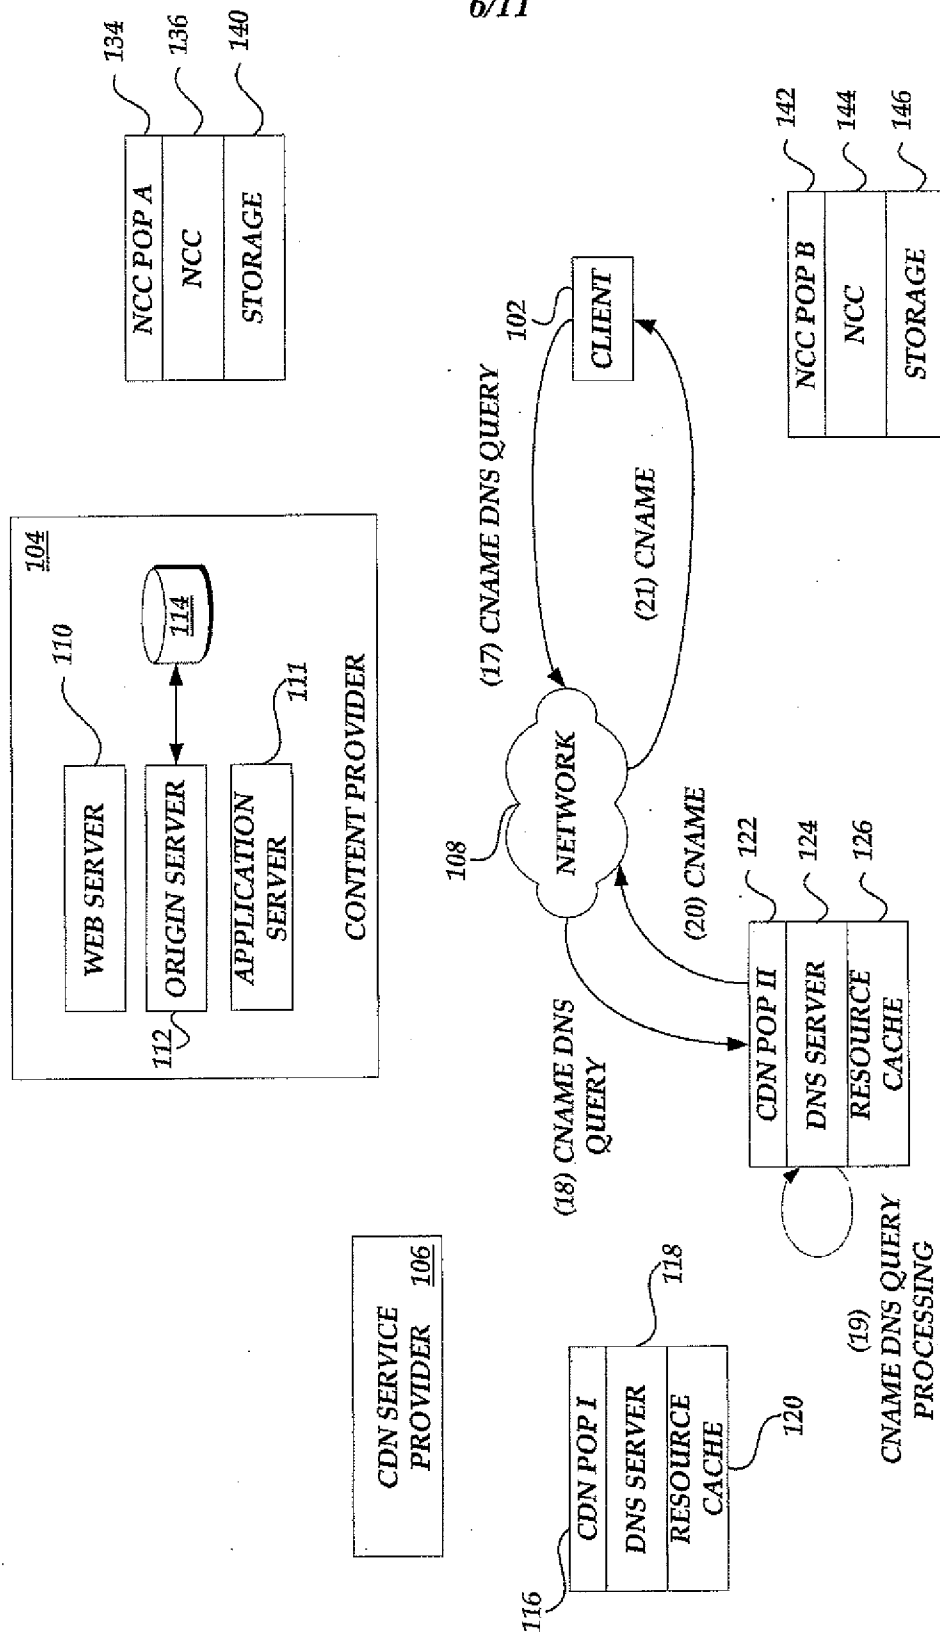


Fig. 5B.

7/11

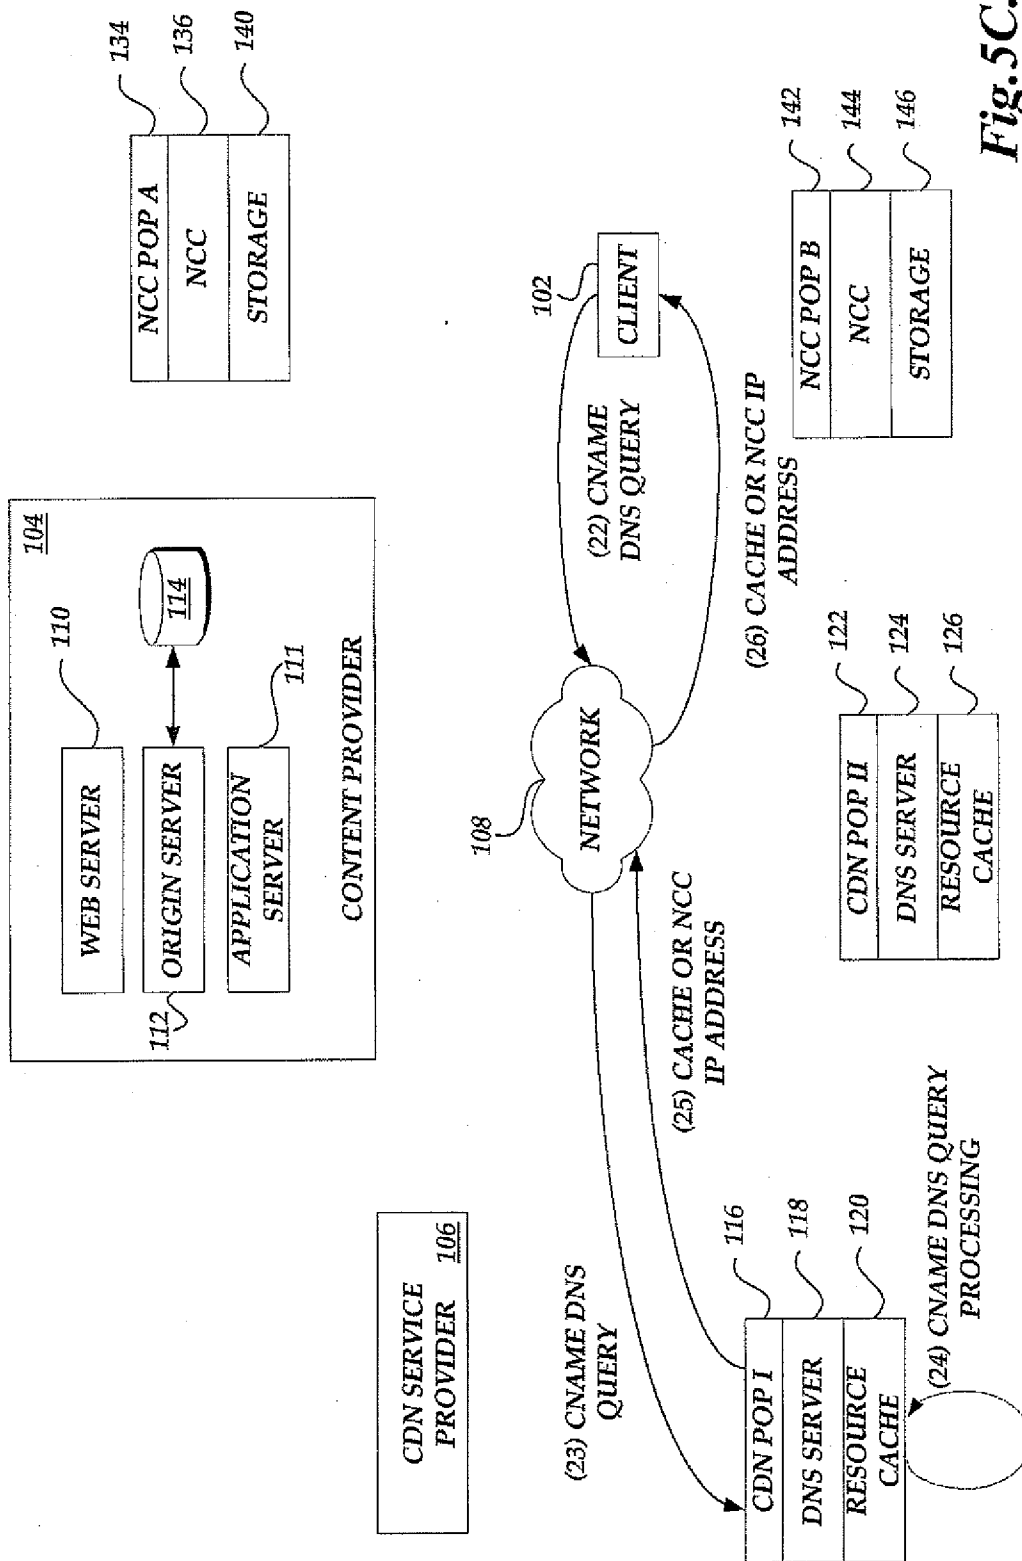
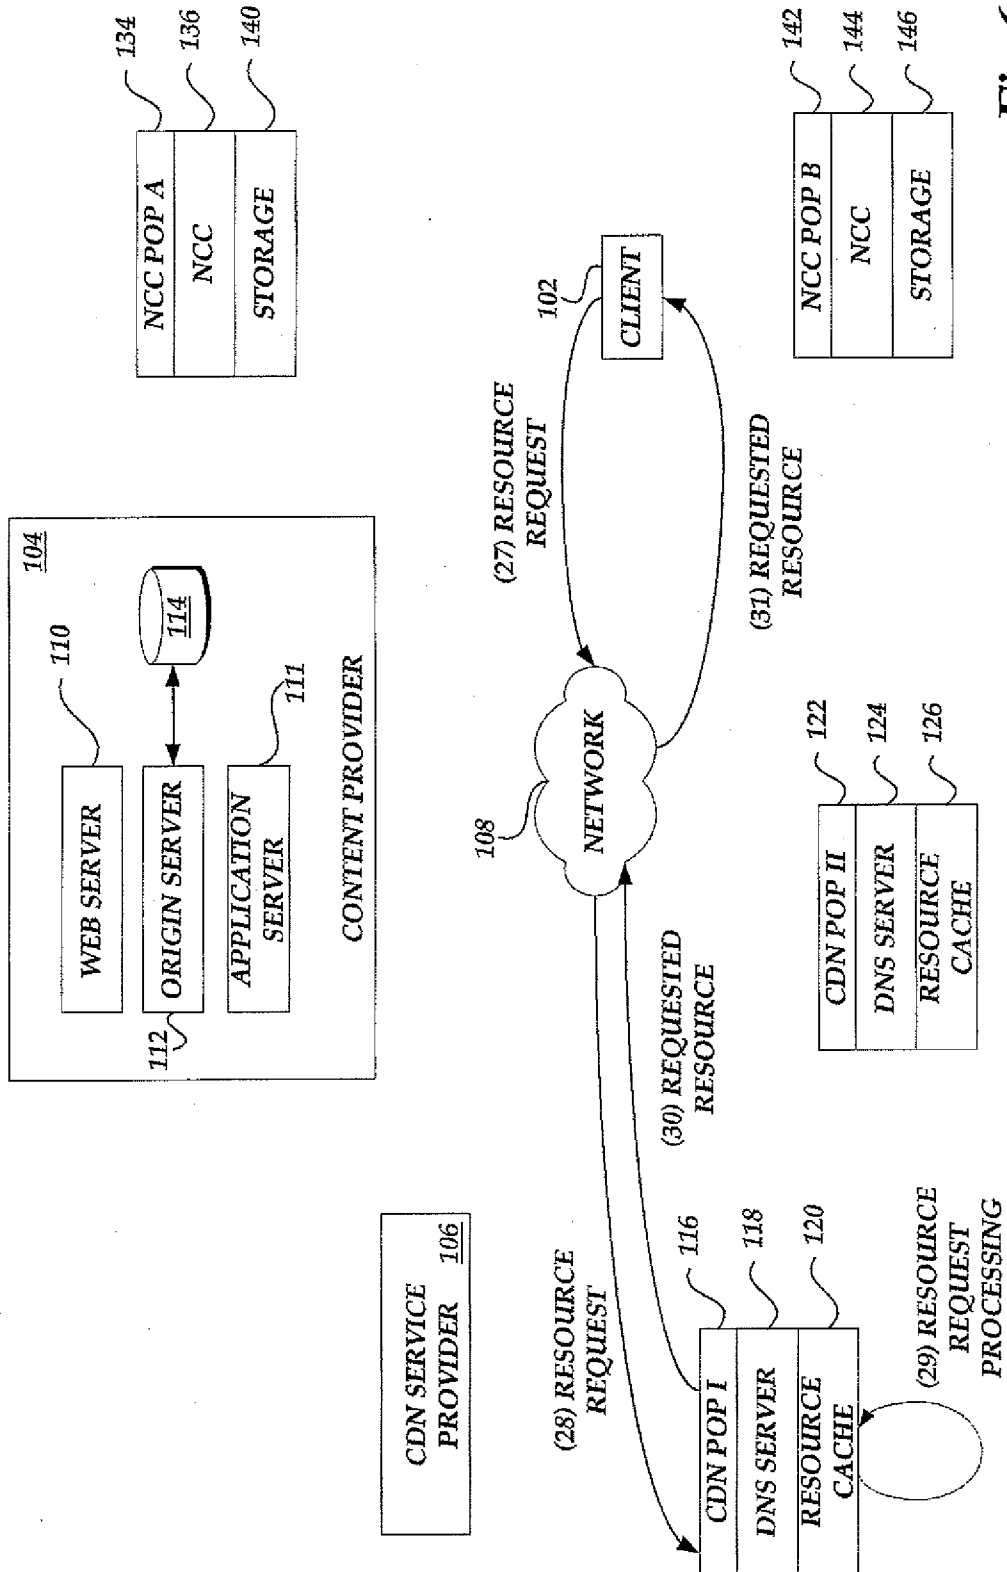


Fig.5C.



8/11

Fig. 6A.

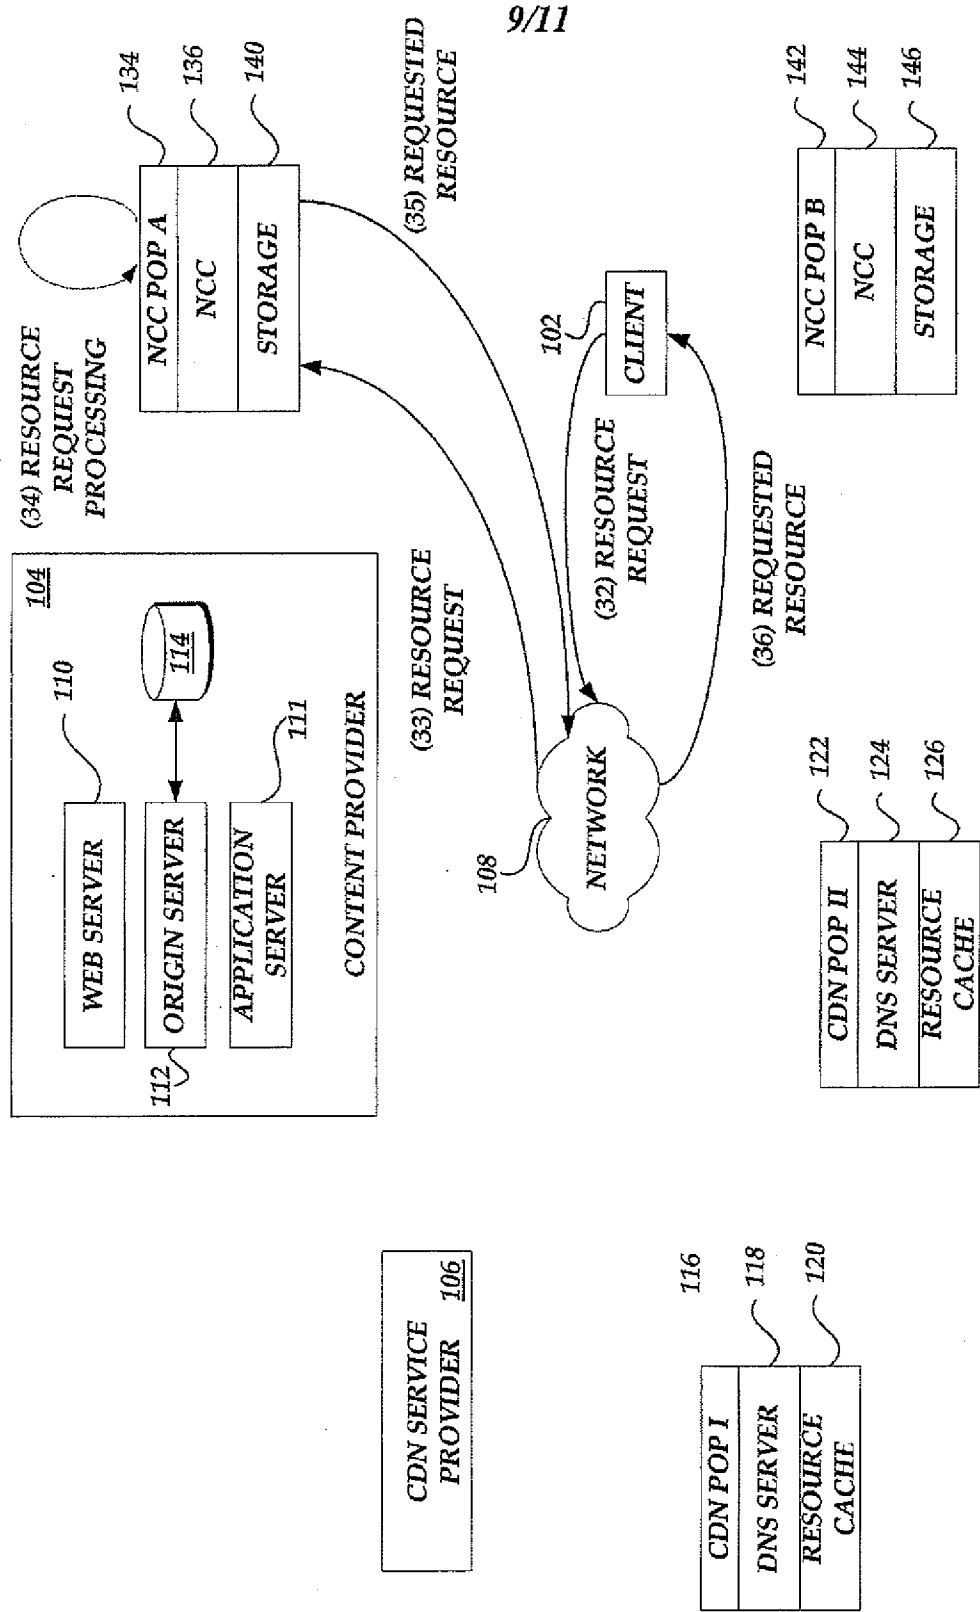


Fig. 6B.

10/11

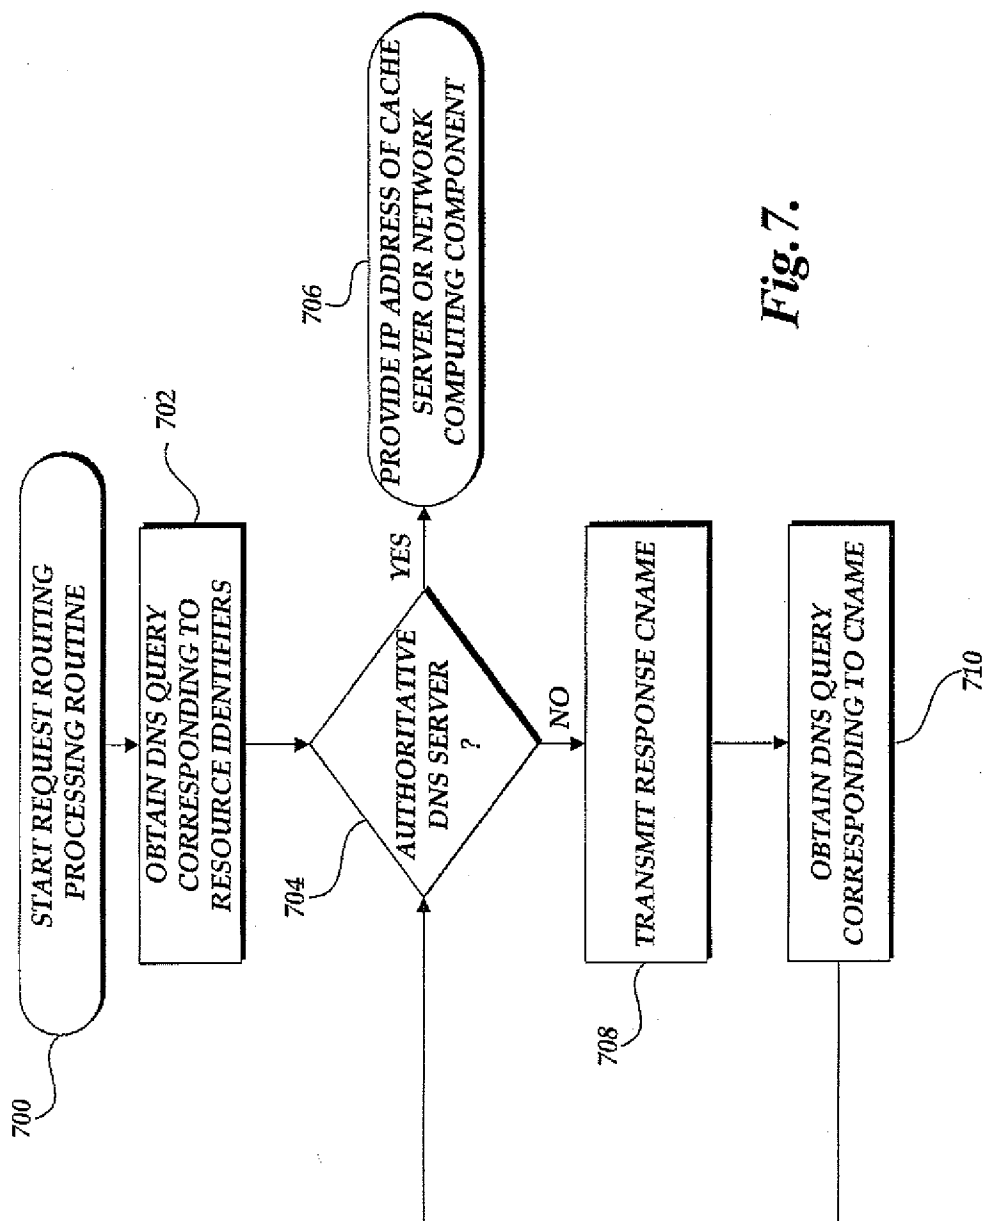
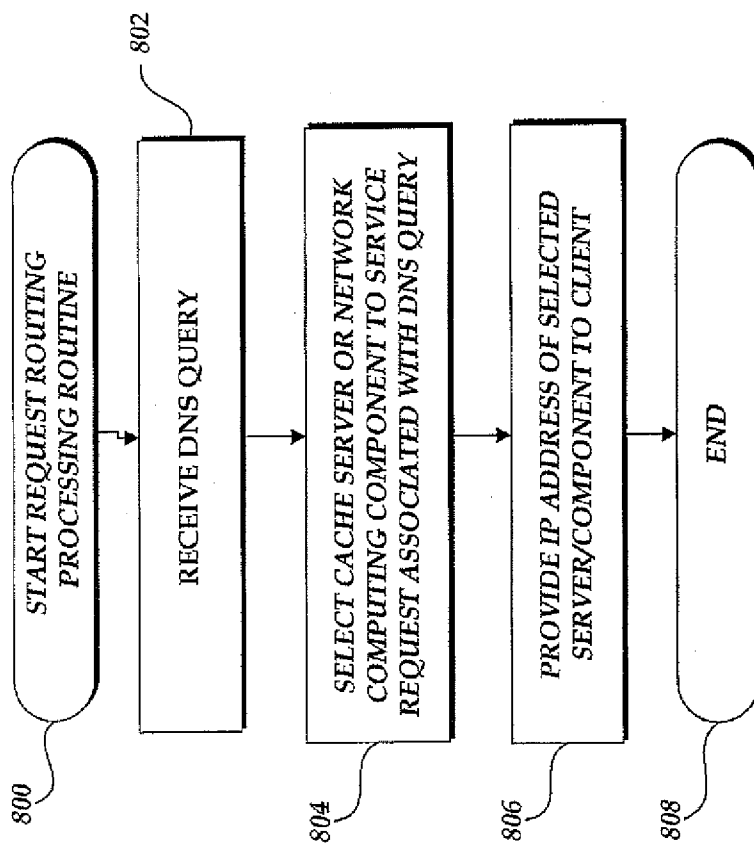


Fig. 7.

11/11

*Fig. 8.*

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 09/47832

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 15/16 (2009.01)

USPC - 709/203

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)

USPC: 709/203

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

USPC: 709/201,238-244

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

USPTO WEST (PGPB, USPT, EPAB, JPAB); GoogleScholar

Search Terms Used: DNS, CDN, cname, domain, name, system, content, delivery, routing, alternative, supplemental, canon, virtual, machine, resource, identifier, client, server

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X ---- Y	US 2008/0086574 A1 (RACIBORSKI et al.) 10 April 2008 (10.04.2008), entire document, especially para [0004], [0018]-[0019], [0025]-[0026], [0032], [0035]-[0036], [0041]-[0046] and abstract	8-14 ----- 1-7, 15
Y	US 2004/0249971 A1 (KLINKER) 09 December 2004 (09.12.2004), entire document, especially para [0055], [0070]-[0074], [0078]	1-7, 15
A	US 2006/0075139 A1 (JUNGCK) 06 April 2006 (06.04.2006)	1-15
A	US 2008/0008089 A1 (BORNSTEIN et al.) 10 January 2008 (10.01.2008)	1-15

☐ Further documents are listed in the continuation of Box C.


* 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

16 July 2009 (16.07.2009)

Date of mailing of the international search report

28 JUL 2009

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

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