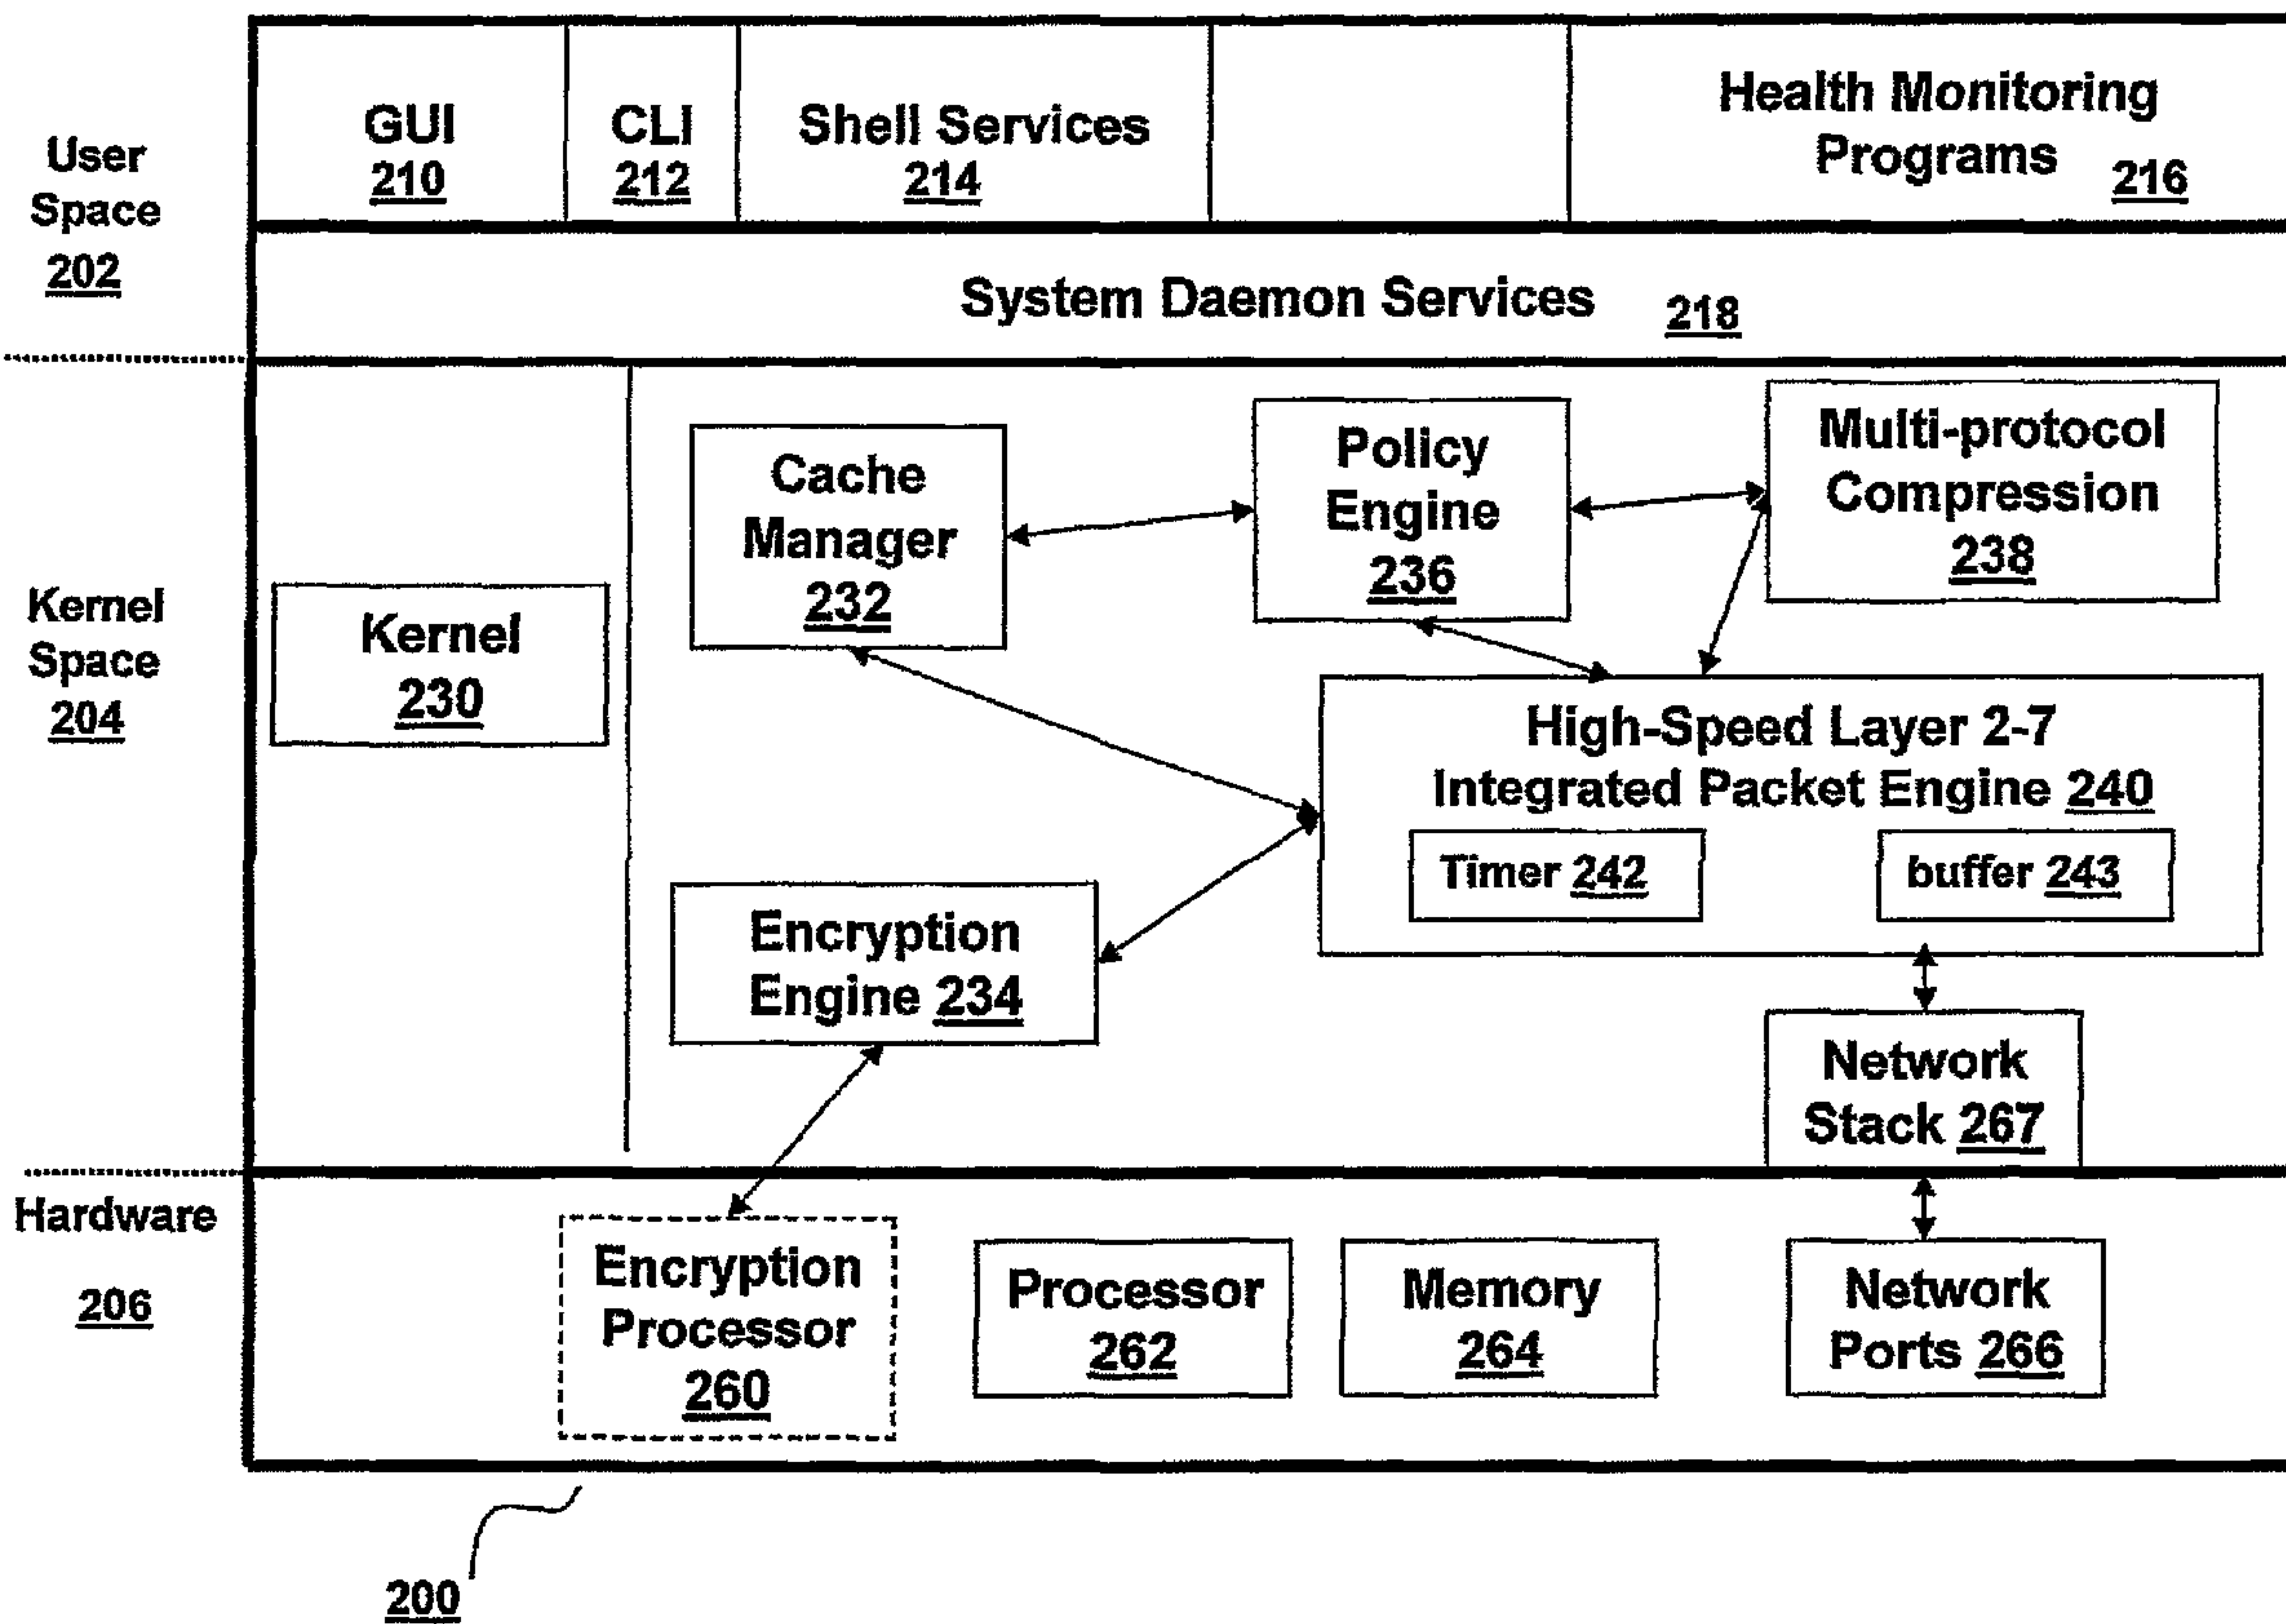




(86) Date de dépôt PCT/PCT Filing Date: 2005/12/30
 (87) Date publication PCT/PCT Publication Date: 2006/08/03
 (85) Entrée phase nationale/National Entry: 2007/07/17
 (86) N° demande PCT/PCT Application No.: US 2005/047435
 (87) N° publication PCT/PCT Publication No.: 2006/081032
 (30) Priorités/Priorities: 2005/01/24 (US11/039,946);
 2005/06/29 (US11/169,002)

(51) Cl.Int./Int.Cl. *H04L 29/08* (2006.01),
G06F 17/30 (2006.01)
 (71) Demandeur/Applicant:
 CITRIX SYSTEMS, INC., US
 (72) Inventeurs/Inventors:
 KHEMANI, PRAKASH, US;
 SUNDARRAJAN, PRABAKAR, US;
 KAILASH, KAILASH, US;
 SONI, AJAY, US;
 SINHA, RAJIV, US;
 ANNAMALAISAMI, SARAVANA, US;
 BHUSHAN, K R, US
 (74) Agent: BERESKIN & PARR

(54) Titre : SYSTEMES ET PROCEDES POUR METTRE EN MEMOIRE CACHE DES OBJETS PRODUITS
 DYNAMIQUEMENT DANS UN RESEAU
 (54) Title: SYSTEMS AND METHODS FOR PERFORMING CACHING OF DYNAMICALLY GENERATED OBJECTS IN A
 NETWORK



(57) Abrégé/Abstract:

The solution of the present invention increases the ability of cache memories to store and serve dynamically generated data. The present invention also enables the cache to effectively deal with a variety of different application request types, thereby increasing

(57) **Abrégé(suite)/Abstract(continued):**

application performance and easing the administrative complexity of preserving freshness of data served from the cache. The present invention provides an effective approach to caching dynamic content by the use of heuristics to effectively predict the behavior of such applications servers in addition to incorporating the ability to understand and process data in a way that does not duplicate processing carried out by the application server that originally generates the object. These techniques of the present invention, in turn, increase the use of dynamic caching, and thereby contribute to the improvement of the performance of both the network, as well as underlying application infrastructure.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
3 August 2006 (03.08.2006)

PCT

(10) International Publication Number
WO 2006/081032 A3

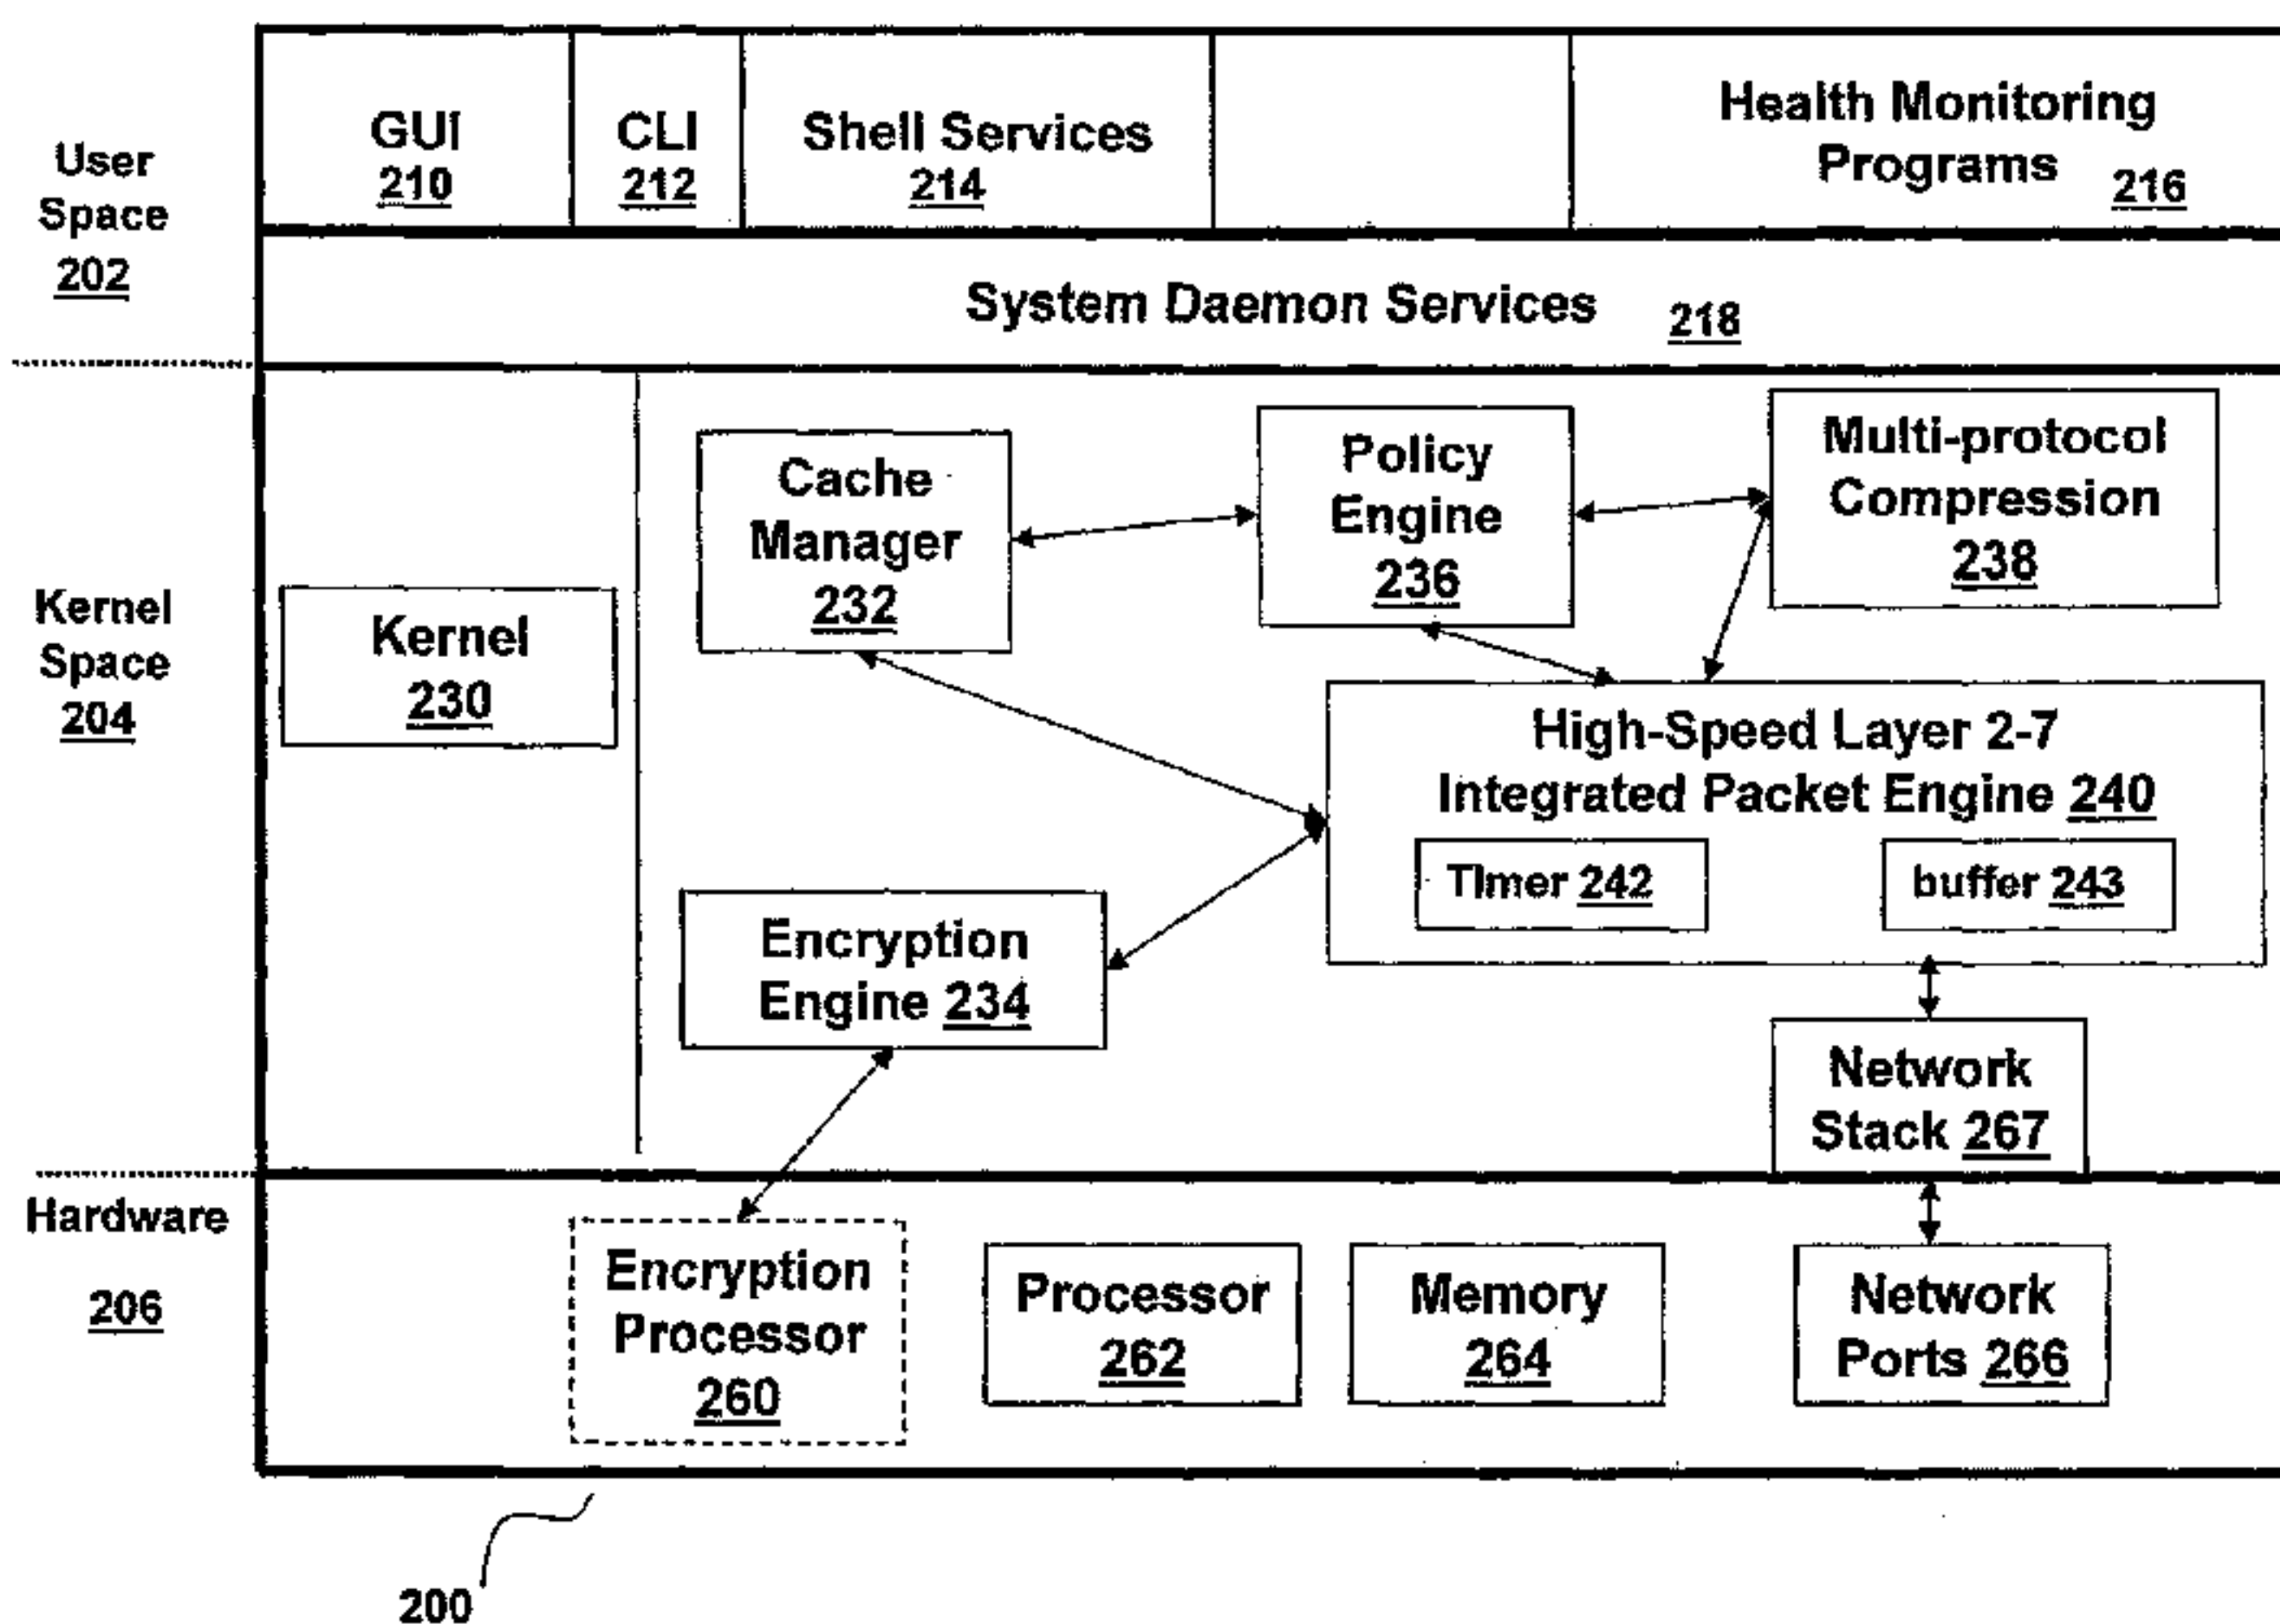
- (51) International Patent Classification:
H04L 29/08 (2006.01) *G06F 17/30* (2006.01)
- (21) International Application Number:
PCT/US2005/047435
- (22) International Filing Date:
30 December 2005 (30.12.2005)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/039,946 24 January 2005 (24.01.2005) US
11/169,002 29 June 2005 (29.06.2005) US
- (71) Applicant (for all designated States except US): **CITRIX SYSTEMS, INC.** [US/US]; 851 West Cypress Creek Road, Fort Lauderdale, Florida 33309 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **KHEMANI, Prakash** [IN/US]; C/O CITRIX APPLICATION NETWORKING GROUP, 180 Baytech Drive, San Jose, California 95134 (US). **SUNDARRAJAN, Prabakar** [US/US]; C/O CITRIX APPLICATION NETWORKING GROUP, 180 Baytech Drive, San Jose, California 95134 (US). **KAILASH, Kailash** [IN/US]; C/O CITRIX APPLICATION NETWORKING GROUP, 180 Baytech Drive, San Jose, California 95134 (US). **SONI, Ajay** [IN/US];

C/O CITRIX APPLICATION NETWORKING GROUP, 180 Baytech Drive, San Jose, California 95134 (US). **SINHA, Rajiv** [US/US]; C/O CITRIX APPLICATION NETWORKING GROUP, 180 Baytech Drive, San Jose, California 95134 (US). **ANNAMALAISAMI, Saravana** [IN/US]; C/O CITRIX APPLICATION NETWORKING GROUP, 180 Baytech Drive, San Jose, California 95134 (US). **BHUSHAN, K R** [IN/US]; C/O CITRIX APPLICATION NETWORKING GROUP, 180 Baytech Drive, San Jose, California 95134 (US).

- (74) Agent: **LANZA, John, D.**; CHOATE, HALL & STEWART, Two International Place, Boston, Massachusetts 02110 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

[Continued on next page]

(54) Title: SYSTEMS AND METHODS FOR PERFORMING CACHING OF DYNAMICALLY GENERATED OBJECTS IN A NETWORK



(57) Abstract: The solution of the present invention increases the ability of cache memories to store and serve dynamically generated data. The present invention also enables the cache to effectively deal with a variety of different application request types, thereby increasing application performance and easing the administrative complexity of preserving freshness of data served from the cache. The present invention provides an effective approach to caching dynamic content by the use of heuristics to effectively predict the behavior of such applications servers in addition to incorporating the ability to understand and process data in a way that does not duplicate processing carried out by the application server that originally generates the object. These techniques of the present invention, in turn, increase the use of dynamic caching, and thereby contribute to the improvement of the performance of both the network, as well as underlying application infrastructure.

WO 2006/081032 A3

WO 2006/081032 A3

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, HU, IE, IS, IT, LT, LU, LV, MC, NL, 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*

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

(88) Date of publication of the international search report:
12 October 2006

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEMS AND METHODS FOR PERFORMING CACHING OF DYNAMICALLY GENERATED OBJECTS IN A NETWORK

Related Applications

[0001] This present application claims priority to U.S. Patent Application Number 11/169,002 entitled "METHOD AND DEVICE FOR PERFORMING INTEGRATED CACHING IN A DATA COMMUNICATION NETWORK," filed June 29, 2005, and U.S. Patent Application Number 11/039,946, entitled "SYSTEM AND METHOD FOR ESTABLISHING A VIRTUAL PRIVATE NETWORK," filed January 24, 2005, both of which are incorporated herein by reference.

Field of the Invention

[0002] The present invention relates generally to caching data in a network. In particular, the present invention relates to a device, systems, and methods for performing caching of dynamically generated objects in a network.

Background

[0003] The growth rate of network traffic continues to strain the infrastructure that carries that traffic. Various solutions have arisen to permit network operators to handle this increasing problem, including the development of caching technology. With traditional caching, static content can be reused and served to multiple clients without burdening server infrastructure. Additionally, cache memories permit static content to be stored closer to the end user, thereby improving response time while at the same time reducing server infrastructure burden. Lowered response times and lowered server infrastructure load reduces bandwidth and the processing requirements of such infrastructure.

[0004] However, an increasing amount of the content delivered across networks is dynamically generated, including a large percentage of network traffic created by enterprise computing solutions and complex internet applications. Dynamically generated content is content generated by the server at the time an object is requested,

and is often based on inputs received from the client. Therefore, it frequently changes both through time and with respect to inputs made to the generating system. Common examples of dynamic content include where a stock quotation request made by a client or database searches. In each instance, the response object is generated in real time following receipt of a specific, client request.

[0005] The challenges to caching dynamically generated content are manifold. For example, there are no generally-accepted standards or specifications for caching dynamically generated content. Since there exists no standard for designating whether a dynamically generated object may be cached, such objects are typically treated as non-cacheable. Another challenge is determining the validity of “freshness” of a dynamically generated object because changes to the underlying data used to generate such objects may be irregular and unpredictable.

[0006] In addition to the above difficulties, requests for dynamically generated content are also typically more complex than requests for static content. Dynamic requests often contain a string of information that needs to be processed or parsed by the destination application to identify applicable parameters that will be used by the cache to identify the appropriate object related to such request. These parameters, however, are rarely placed in the request in a logical or consistent order by the client. To determine which of the multitude of dynamically generated objects is identified by the request, each such request must be normalized (i.e., place the parameters in non-arbitrary order).

[0007] Furthermore, matching a request to a dynamically generated object becomes a much more complex task with dynamically generated content because certain processing done by the application may need to be duplicated or otherwise anticipated, by making an educated guess. This duplication or guessing is necessary to decide whether an object stored by the cache is appropriate for serving to a particular incoming request. The complexity arises as a result of the complexity of the applications themselves, and also because the contents of the response can be a function of both the contents of the request, as well as certain other external variables like the user-identity (which may or may not be present in the request), the time of the day, the current state of the user database and a myriad of other factors.

[0008] In summary, caching originally developed around the caching of static objects. As the Internet and applications becomes more and more dependent upon delivering dynamically generated content, the need has arisen for a solution that extends the

benefits of caching to dynamic content, and that solves the variety of challenges such content presents for traditional caching technology.

Brief Summary of the Invention

[0009] The solution of the present invention increases the ability of cache memories to store and serve dynamically generated data. The present invention also enables the cache to effectively deal with a variety of different application request types, thereby increasing application performance and easing the administrative complexity of preserving freshness of data served from the cache. The present invention provides an effective approach to caching dynamic content by the use of heuristics to effectively predict the behavior of such applications servers in addition to incorporating the ability to understand and process data in a way that does not duplicate processing carried out by the application server that originally generates the object. These techniques of the present invention, in turn, increase the use of dynamic caching, and thereby contribute to the improvement of the performance of both the network, as well as underlying application infrastructure.

[0010] In one embodiment, the present invention is directed towards a method and system for caching and maintaining dynamically generated objects in a cache. The techniques of the present invention include receiving an invalidation command at the cache to invalidate an object, such as a dynamically generated object previously served from an originating server and stored in the cache. The dynamically generated object may not be identified as cacheable from the originating server. The invalidation command received by the cache identifies the cached dynamically generated object. In response to the invalidation command, the cache marks the cached dynamically generated object as invalid, and flushes the object from the cache.

[0011] In another embodiment, the present invention also provides techniques for identifying cached dynamically generated objects using an object determinant. The cache may intercept a communication between the client and server, and parse the communication to identify an object determinant. The object determinant may identify an object previously served and stored in the cache. The cache determines from the object determinant whether a change has occurred or will occur, in the object identified by the object determinant at the originating server. If a change has occurred or will occur, the cache marks as invalid in the cache the object identified by the

object determinant. Once the object has been marked as invalid, the cache may flush the invalid object and retrieve the object from the originating server.

[0012] Further embodiments of the present invention apply the above methods to groups of dynamically generated objects. For example, a group of previously served dynamically generated objects are formed in the cache. The group of objects is associated with at least one object determinant. A record of the group is maintained in the cache and may be associated with the object determinant. The group of previously served objects is marked as invalid if the identified object determinant of a communication intercepted by the cache indicates a change has occurred or will occur in one or more objects of the group at the originating server.

[0013] In one aspect, the present invention is related to a method for caching a dynamically generated object not identified as cacheable. The method includes storing in a cache a dynamically generated object served from an originating server and not identified as cacheable, and receiving, by the cache, a request to invalidate the cached dynamically generated object. In response to the request, the cache marks the cached dynamically generated object as invalid. The cache may be operated on any device, such as an appliance, a network device or a computing device in communications between the originating server and a client.

[0014] In some embodiments, the method of the present invention includes requesting, by the originating server to invalidate the cached dynamically generated object. In a further embodiment, the originating server automatically requests to invalidate the cached dynamically generated object in response to a change to the dynamically generated object in the originating server. In another embodiment, the client is in communication with the originating server to receive the dynamically generated object, and the client requests to invalidate the cached dynamically generated object. In yet another embodiments, an external administrative control requests invalidation of the cached dynamically generated object.

[0015] In one embodiment, the method of the present invention includes flushing from the cache the cached dynamically generated object marked as invalid. In another embodiment, the cache receives the request to invalidate within a very short time period, for example, ten milliseconds or less, of caching the dynamically generated object. In some embodiments, the cache invalidates the cached dynamically generated object responsive to an expiration of a very short time expiry, such as an expiry of 10 milliseconds or less of the cached object.

- [0016] In another aspect, the present invention is related to a method for caching a group of dynamically generated objects. In some embodiments, the group of objects has at least one object not identified as cacheable. The method includes identifying, in a cache, a group of dynamically generated objects previously served from an originating server. The cache associates the group with an object determinant. The method further includes intercepting, by the cache, a communication identifying the object determinant of the group and indicating a change is about to occur or has occurred on the originating server to one of the objects of the group. In one embodiment, the method of the present invention includes marking, by the cache, the group of dynamically generated objects as invalid in response to intercepting the communication or identifying the object determinant.
- [0017] Furthermore, in some embodiments of the present invention, the cache may flush from the cache the group of objects marked as invalid. In other embodiments, the group of dynamically generated objects is pre-designated. In other embodiments, the method automatically identifies the group of dynamically generated objects and associates the object determinant with the group according to a rule.
- [0018] In one aspect, the present invention is related to another method for maintaining a cache of dynamically generated object. This method of the present invention includes intercepting, by a cache, a communication between a client and an originating server, for example, a client's request to an originating server for a dynamically generated object. The dynamically generated object may have been previously served from the originating server and stored in the cache. The method identifies, by the cache, an object determinant in the communication indicating one of a change has occurred or will occur in a dynamically generated object at the originating server, and marks, by the cache, the cached dynamically generated object as invalid. The method then obtains the requested dynamically generated object from the originating server.
- [0019] In one embodiment of the method, the cache flushes the invalid dynamically generated object. In another embodiment, the cache associates the dynamically generated object with a group of dynamically generated objects previously served from the originating server, associates the group with the object determinant, and marks the group of dynamically generated objects as invalid in response to the request. In some embodiments, the dynamically generated object is not identified as cacheable. In further embodiments, the cache flushes the group of dynamically

generated objects marked as invalid. The group of dynamically generated objects may be pre-designated or other automatically identified via an object determinant according to a rule.

[0020] In some embodiments of the method, the client embeds in the communication the object determinant as a pre-defined string. In other embodiments, the cache identifies the object determinant based on a pre-defined heuristic rule associated with the dynamically generated object. In one embodiment, the cache selects the object determinant from one of the following elements of the communication: 1) USERID, 2) IP address, 3) TCP port, 4) HTTP header, 5) custom HTTP header, 6) client URL, 7) cookie header, 8) URL query string, and 9) POST body. The cache may also extract from the communication the object determinant based on a user-configured invalidation policy.

[0021] In another embodiment, the method of the present invention includes maintaining a table in the cache to associate the object determinant with one or more objects or groups of objects stored in the cache. In some embodiments, the method includes examining, by an intelligent statistical engine, communications from the client to identify a set of dynamically generated objects to associate as a group in the cache. The cache may associate objects stored in the cache into a content group. The content group may be represented by a hash table having an incarnation number as an index. In other embodiments, the cache performs a hash algorithm on the dynamically generated object identified via the request to determine a change to the dynamically generated object.

[0022] In some aspects, the present invention is related to a system for caching a dynamically generated object not identified as cacheable, the system includes means for storing in a cache a dynamically generated object not identified as cacheable, the dynamically generated object served from an originating server, and a means for receiving, by the cache, a request to invalidate the cached dynamically generated object. The system also includes means for marking, by the cache in response to the request, the cached dynamically generated object as invalid.

[0023] In other aspects, the presented is related to a system for caching a group of objects, such as dynamically generated objects. In some embodiments, the dynamically generated objects has at least one object not identified as cacheable. The system includes means for identifying, in a cache, a group of dynamically generated objects previously served from an originating server, at least one of the dynamically

generated objects not identified as cacheable, and means for associating, by the cache, the group with an object determinant. The system also includes means for intercepting, by the cache, a communication identifying the object determinant of the group and indicating a change is about to occur or has occurred on the originating server to one of the objects of the group.

[0024] In one aspect, the present invention is related to a system for maintaining a cache of dynamically generated objects. The system includes means for intercepting, by a cache, a communication between a client and an originating server. The system also includes means for identifying, by the cache, an object determinant in the communication that indicates a change has occurred or will occur in a dynamically generated object at the originating server, and means for marking, by the cache, the cached dynamically generated object as invalid. In some embodiments, the system also includes means for obtaining the dynamically generated object from the originating server.

[0025] In another aspect, the present invention is directed towards a method and system for providing granular invalidation of dynamically generated objects stored in a cache. The techniques of the present invention incorporates the ability to configure the expiration time of objects stored by the cache to fine granular time intervals, such as the granularity of time intervals provided by a packet processing timer of a packet processing engine. As such, the present invention can cache objects with expiry times down to very small intervals of time. This characteristic is referred to as "invalidation granularity." By providing this fine granularity in expiry time, the cache of the present invention can cache and serve objects that frequently change, sometimes even many times within a second. One technique is to leverage the packet processing timers used by the device of the present invention that are able to operate at time increments on the order of milliseconds to permit invalidation or expiry granularity down to 10 ms. As the integrated cache of the present invention performs cache operations responsive to the packet processing timer of a packet processing engine, invalidation of objects can occur within time periods in the order of milliseconds and with objects having an expiry in the order of the time intervals provided by the packet processing timer, such as 10 ms.

[0026] In one aspect, the present invention is related to a method for invalidating a cached dynamically generated object in the cache responsive to a signal of a packet processing timer. The method may be performed in a packet processing engine of a

device, in which the packet processing engine processes network packets having at least one dynamically generated object and storing the dynamically generated object in a cache. The method includes receiving a first signal from a packet processing timer to process a network packet, processing the network packet responsive to the signal and transmitting a second signal to a cache to indicate a cached dynamically generated object is invalid. In one embodiment, the second signal provides an invalidation command to invalidate the cached dynamically generated object. In another embodiment, the second signal provides expiration of an expiry of the cached dynamically generated object.

[0027] In some embodiments, the second signal is transmitted during processing of the network packet, while in other embodiments, the second signal is transmitted upon completion of processing of the network packet. In one embodiment of the present invention, the packet processing timer operates at a time increment of one or more milliseconds. The packet processing timer operates at a time increment to invalidate an object stored in cache at an expiry time of 10 milliseconds or less.

The method may include marking, by the cache, the cached dynamically generated object as invalid. In one embodiment, the method includes transmitting from an originating server the network packet comprising dynamically generated object.

[0028] The device having the packet processing engine may include any of the following: 1) a bridge, 2) a router, 3) a switch, and a 4) SSL VPN device. In some embodiments, the cache, the packet processing engine, or the packet processing timer operates in a kernel space of the device. In some embodiments, the method includes providing, via the packet processing timer, a time interval for expiry of the cached dynamically generated object a time equal to or less than a duration of time to dynamically generate and serve the dynamically generated object from an originating server. The dynamically generated object may be identified as not cacheable when served from the originating server.

[0029] In another aspect, the present invention is related to a device for invalidating a cached dynamically generated object responsive to a signal of a packet processing timer. The device processes network packets having at least one dynamically generated object and storing the dynamically generated object in a cache. The device includes a packet processing timer generating a first signal, and a packet processing engine in communication with the packet processing timer. The packet processing engine processes the network packet responsive to receipt of the first signal of the

packet processing timer. The device also includes a cache in communication with the packet processing engine and having a dynamically generated object cached in a storage element. The cache receives a second signal from the packet processing engine indicating the cached dynamically generated object is invalid. In one embodiment, the second signal comprises an invalidation command to invalidate the cached dynamically generated object. In another embodiment, the second signal comprises an expiration of an expiry of the cached dynamically generated object.

[0030] In one embodiment of the device of the present invention, the packet processing engine transmits the second signal to the cache during processing of the network packet by the packet processing engine. In another embodiment, the packet processing engine transmits the second signal to the cache upon completion of processing of the network packet. In some embodiments, the packet processing timer operates at a time interval of one or more milliseconds. In one embodiment, the packet processing timer operates at a time interval to trigger the second signal to the cache to invalidate the cached object at a first time interval of 10 milliseconds or less.

[0031] The device of the present invention may include one of the following: 1) a bridge, 2) a router, 3) a switch, and 4) a SSL VPN device. The cache, the packet processing engine, or the packet processing timer may operate in a kernel space of the device. In some embodiments, the packet processing timer provides a time interval for expiry of the cached dynamically generated object less than or equal to a duration of time to dynamically generate and serve the dynamically generated object from an originating server. The dynamically generated object is not identified as cacheable from the originating server. The originating server may transit the network packet having the dynamically generated object.

[0032] In another aspect, the present invention is directed towards a method and system for providing a technique referred to as flash caching to respond to requests for an object, such as a dynamically generated object, from multiple clients. This technique of the present invention uses a dynamically generated object stored in a buffer for transmission to a client, for example in response to a request from the client, to also respond to additional requests for the dynamically generated object from other clients while the object is stored in the buffer. Using this technique, the present invention is able to increase cache hit rates for extremely fast changing objects, such as dynamically generated objects that would not otherwise be cacheable.

[0033] In one aspect, the present invention is related to a method for responding to a request for a dynamically generated object from a plurality of clients. The method includes receiving, by a device, a response from an originating server to a first request of a first client for a dynamically generated object. The response includes the dynamically generated object. The method also includes requesting, by the device, to a packet processing engine to transmit the response to the first client. The response is held in a buffer while waiting to be transmitted. The method further includes receiving, by the device, from a second client a second request for the dynamically generated object prior to completing transmission of the response to the first client, and requesting, by the device, to the packet processing engine to transmit the dynamically generated object of the response to the first client held in the buffer to the second client.

[0034] In one embodiment, the method of the present invention includes transmitting, by the packet processing engine, the response to the first client. In another embodiment, the method includes transmitting, by the packet processing engine, the dynamically generated object to the second client in response to the second request. In some embodiments, the method includes removing, by the packet processing engine, the response from the buffer after transmitting the response to one of the first client or the second client. In some embodiments, the buffer used by the present invention comprises a TCP/IP buffer or a portion of a TCP/IP stack of the device. The device may have a single TCP/IP stack for transmitting to the first client and the second client, and the buffer is associated with the single TCP/IP stack.

[0035] In some embodiments of the present invention, the method includes receiving, by the device, multiple requests from multiple clients for the dynamically generated object, and requesting, by the device, to transmit the dynamically generated object of the response to the first client held in the buffer to the multiple clients. In one embodiment, the packet processing engine holds the response to the first client based on one or more characteristics related to a connection of the first client to a network. In another embodiment, the device queues the second request instead of transmitting the second request to the originating server. In one embodiment, the device responds to the second request using the response of the first request held in the buffer.

[0036] In some embodiments, the method of the present invention removes the dynamically generated object from the buffer after transmitting the dynamically generated object to the first client and the second client. In one embodiment, the

dynamically generated object is identified as non-cacheable, such as when served from an originating server. In other embodiments of the method, the device determines the second request is for the dynamically generated object of the first request via an object determinant associated with the second request.

[0037] In some embodiments, the device may be an appliance, a network device or a computing device in communication between the first client and the originating server. In further embodiments, the device comprise a cache in communication with the packet processing engine. In one embodiment, the method includes requesting, by the cache, the packet processing engine to transmit one of the response to the first client or the dynamically generated object of the first client to the second client.

[0038] In another aspect, the present invention is related to a system for responding to a request for a dynamically generated object from a plurality of clients. The system includes a device receiving a response from an originating server to a first request of a first client for a dynamically generated object, the response comprising the dynamically generated object, the device comprising, and a buffer for holding one or more network packets waiting to be transmitted from the device. The device also includes a packet processing engine receiving a request from the device to transmit the response to the first client; the buffer holding one or more network packets comprising the response. The device receives from a second client a second request for the dynamically generated object prior to completing transmission of the response to the first client, and requests the packet processing engine to transmit the dynamically generated object of the response held in the buffer to the second client.

[0039] In one embodiment of the system of the presenting invention, the packet processing engine transmits from the buffer the response to the first client. In another embodiment, wherein the packet processing engine transmits from the buffer the dynamically generated object to the second client in response to the second request. In some embodiments of the system, the packet processing engine removes the response from the buffer after transmitting the response to one of the first client or the second client. Additionally, in other embodiments, the buffer is a TCP/IP buffer or otherwise is a portion of a TCP/IP stack of the device. In some embodiments, the device has a single TCP/IP stack for transmitting to the first client and the second client the buffer associated with the single TCP/IP stack.

[0040] In another embodiment of the present invention, the device receives from multiple clients multiples requests for the dynamically generated object, and wherein the device requests the packet processing engine to transmit the dynamically generated object of the response to the first client held in the buffer to the multiple clients. In one embodiment, the packet processing engine holds the response to the first client based on one or more characteristics related to a connection of the first client to a network. In another embodiment, the device queues the second request instead of transmitting the second request to the originating server. In some embodiments, the device queues the second request and responds to the second request with the response held in the buffer for the first client. In additional embodiments, the packet processing engine removes the dynamically generated object from the buffer after transmitting the dynamically generated object to the first client and the second client. In one embodiment, the dynamically generated object is identified as non-cacheable, such as by an originating server.

[0041] In another embodiment of the present invention, the device determines the second request is for the dynamically generated object of the first request via an object determinant associated with the second request. In some embodiments, the device is an appliance, a network device or a computing device in communication between the first client and the originating server. The device may also have a cache in communication with the packet processing engine. Additionally, the cache may request the packet processing engine to transmit the response to the first client or the dynamically generated object of the first client to the second client.

[0042] In another aspect, the present invention is directed towards a "flash crowd" technique for handling situations where the cache receives additional requests, e.g., nearly simultaneous requests, for the same object during the time the server is processing and returning the response object for a first requestor. Once all such nearly simultaneous requests are responded to by the cache, the object is flushed from the cache, with no additional expiry time or invalidation action needed. This technique of the present invention enables data to be cached and served for very small amounts of time for objects that would otherwise be considered non-cacheable. This approach yields a significant improvement in applications that serve fast changing data to a large volume of concurrent users, such, for example, as real time stock quotes, or a fast evolving news story.

[0043] In one aspect, the present invention is related to a method in a network device for responding to a request for a dynamically generated object from a plurality of clients. The method includes receiving from a first client a first request for a dynamically generated object from an originating server and transmitting the first request to the originating server. Prior to responding to the first request of the first client, the method includes receiving from a second client a second request for the dynamically generated object. The method further includes receiving the response to the first request from the originating server, the response having the dynamically generated object, and transmitting the dynamically generated object to the first client in response to the first request, and the dynamically generated object to the second client in response to the second request.

[0044] In one embodiment of the present invention, the method includes receiving multiple requests from multiple clients for the dynamically generated object, and transmitting the dynamically generated object to the multiple clients in response to each client's respective request. In some embodiments, the method includes receiving a third request for the dynamically generated object from a third client prior to completing transmission of the dynamically generated object to the second client, and transmitting the dynamically generated object to the third client in response to the third request.

[0045] In another embodiment, the method of the present invention receives the first request and the second request nearly simultaneously. In one embodiment, the method queues the second request instead of transmitting the second request to the originating server. In some embodiments, the method flushes the dynamically generated object from a cache after transmitting the dynamically generated object to the first client and the second client. In one embodiment, the dynamically generated object is identified as non-cacheable, such as when served from the originating server. The method may include determining the second request is for the dynamically generated object of the first request via an object determinant associated with the second request. In yet another embodiment, the method modifies the response to include an entity tag header or a cache-control header and transmitting the modified response to the first client and/or the second client

[0046] In another aspect, the present invention is related to a network device for responding to a request for a dynamically generated object from a plurality of clients. The network device includes means for receiving from a first client a first request for

a dynamically generated object from an originating server, and means for transmitting the first request to the originating server. The network device also includes means for receiving from a second client a second request for the dynamically generated object prior to responding to the first request of the first client, and for receiving the response to the first request from the originating server, the response comprising the dynamically generated object, The system also includes means for transmitting the dynamically generated object to the first client in response to the first request, and the dynamically generated object to the second client in response to the second request.

[0047] In one embodiment of the present invention, the network device further includes means for receiving from multiple clients multiple requests for the dynamically generated object, and transmitting the dynamically generated object to the multiple clients in response to each client's request. In another embodiment, the network device includes means for receiving a third request for the dynamically generated object from a third client prior to completing transmission of the dynamically generated object to the second client, and transmitting the dynamically generated object to the third client in response to the third request. In other embodiments, the network devices has means for receiving the first request and the second request nearly simultaneously. In some embodiments, the network device queues the second request, for example, instead of transmitting the second request too the server.

[0048] In one embodiment, the network device of the present invention flushes the dynamically generated object from a cache after transmitting the dynamically generated object to the first client and the second client. In some embodiments, the dynamically generated object is not identified as cacheable. In other embodiments, the network device includes means for determining the second request is for the dynamically generated object of the first request via an object determinant associated with the second request. In another embodiment, the network device also includes means for modifying the response to include one of an entity tag header or a cache-control header and transmitting the modified response to one of the first client or the second client. The network device may be an appliance or a computing device in communication between the client and the originating server.

[0049] In yet another aspect, the present invention is directed towards a method and system for modifying by a cache responses from a server that do not identify a dynamically generated object as cacheable to identify the dynamically generated

object to a client as cacheable in the response. In some embodiments, such as an embodiment handling HTTP requests and responses for objects, the techniques of the present invention insert an entity tag, or "etag" into the response to provide cache control for objects provided without entity tags and/or cache control information from an originating server. This technique of the present invention provides an increase in cache hit rates by inserting information, such as entity tag and cache control information for an object, in a response to a client to enable the cache to check for a hit in a subsequent request.

[0050] In one aspect, the present invention is related to a method of using an entity tag for caching a dynamically generated object not identified as cacheable. The method includes receiving by a cache a response to a request of a client for a dynamically generated object from an originating server. The dynamically generated object is not identified as cacheable. The method includes generating by the cache an entity tag for the dynamically generated object stored in the cache, and modifying the response to include the generated entity tag. The method further includes transmitting, by the cache, to the client the modified response identifying the dynamically generated object as cacheable.

[0051] In one embodiment of the method of the present invention, the response received from the originating server does not comprise an entity tag. In some embodiments, the cache modifies the response to identify a duration for the client to cache the dynamically generated object. In other embodiments, the cache modifies the response to indicate to the client to store the response and serve the response after checking if the dynamically generated object has changed with one of the cache or the originating server. In one embodiment, the response from the originating server does not have cache control information for the dynamically generated object, and the cache modifies the response transmitted to the client to include cache control information.

[0052] In some embodiments, the method of the present invention also include generating by the cache a number for the entity tag. The cache may store the generated entity tag with the dynamically generated object. In another embodiment, the cache may store cache control information with the dynamically generated object. In one embodiment, the cache serves the cached dynamically generated object with the generated entity tag and cache control information.

In some embodiments of the present invention, the cache is operated on an appliance, network device or computing device. In other embodiments, the request or the response comprises Hypertext Markup Language (HTML). In another embodiment, the cache modifies the response to include an etag header identifying the generated entity, while still in another embodiment, the cache modifies the response to include a cache-control header. In some embodiments, the method includes receiving by the cache, from the client a second request having an If-None-Match header, and in response to the second request, comparing the etag value associated with the If-None-Match header to the entity tag stored in the cache with the dynamically generated object. In one further embodiment, the cache transmits to the client a not modified response to the second request if the etag value matches the entity tag; otherwise in another embodiment, the cache transmits to the client an updated dynamically generated object with a second entity tag if the etag value does not match the entity tag.

[0054] The details of various embodiments of the invention are set forth in the accompanying drawings and the description below.

Brief Description of the Drawings

[0055] The accompanying drawings, which are incorporated herein and form part of the specification, illustrate the present invention and, together with the description, further serve to explain the principles of the invention and to enable a person skilled in the relevant art(s) to make and use the invention.

[0056] FIG. 1 is a block diagram illustrating an example network environment in which an embodiment of the present invention may be implemented;

[0057] FIG. 2 is a block diagram illustrating an example architecture of an appliance that performs integrated caching in accordance with an embodiment of the present invention;

[0058] FIG. 3A is a flow diagram of steps taken in an embodiment of a method of the present invention for integrating device operations with packet processing and the packet processing timer;

FIG. 3B is a flow diagram of steps taken in an embodiment of a method of the present invention for practicing invalidation granularity techniques in view of FIG.

3A;

[0060] FIG. 4A is a flow diagram of steps taken in an embodiment of a method of the present invention using invalidation commands to invalidate stale objects;

[0061] FIG. 4B is a flow diagram of steps taken in an embodiment of a method of the present invention incorporating invalidation of groups of objects;

[0062] FIG. 4C is a flow diagram of steps taken in an embodiment of a method of the present invention wherein a client request is parsed for object determinants;

[0063] FIG. 4D is a flow diagram of steps taken in an embodiment of a method of the present invention incorporating invalidation of groups of objects using object determinants;

[0064] FIG. 5. is a flow diagram of steps taken in an embodiment of a method of the present invention for providing a flash cache technique;

[0065] FIG. 6. is a flow diagram of steps taken in an embodiment of a method of the present invention for providing a flash crowd control technique;

[0066] FIGs. 7A and 7B are flow diagrams of steps taken in an embodiment of a method of the present invention for providing entity tag and cache control for an object; and

[0067] FIGs. 8A and 8B are block diagrams of embodiments of a computing device for practicing an illustrative embodiment of the present invention.

[0068] The features and advantages of the present invention will become more apparent from the detailed description set forth below when taken in conjunction with the drawings, in which like reference characters identify corresponding elements throughout. In the drawings, like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements.

Detailed Description

A. Example Network Environment

[0069] FIG. 1 illustrates an example network environment 100 in which an embodiment of the present invention may be practiced. As shown in FIG. 1, example

network environment 100 includes a plurality of clients 102a-102n, a plurality of servers 106a-106n., and an appliance 104, which may also referred to as a cache appliance, device, or cache. The servers 106a-106n originate and manage databases, such as object or relational databases, that provide requested content to the clients 102a-102n. For this reason, the servers 106a-106n are sometimes referred to herein as "originating servers" because they typically, though not necessarily, originate the objects forming the requested content. Each of the clients 102a-102n and servers 106a-106n may be any type and form of computing device, such as the computing device 800 described in more detail later in conjunction with FIGs. 8A and 8B. For example, any of the client 102a-102n may be a mobile computing device, such as a telecommunication device, e.g., cellphone or personal digital assistant, or a laptop or notebook computer in addition to any type of desktop computer.

[0070] Each of the clients 102a-102n are communicatively coupled to appliance 104 via a public data communication network 108, while appliance 104 is communicatively coupled to servers 106a-106n via a private data communication network 110. In one embodiment, public data communication network 108 comprises the Internet and private data communication network 110 comprises an enterprise network. The public data communication network 108 and private data communication network 110 can be any type and form of network, public, private or otherwise, and in some cases, may be the same network.

[0071] Although FIG. 1 shows a network 108 and a network 110 between the clients 102a-102n and the servers 106a-106n, the clients 102a-102n and the servers 106a-106n may be on the same network 108 or 110. The networks 108 and 110 can be the same type of network or different types of networks. The network 108 and/or the network 110 can be a local-area network (LAN), such as a company Intranet, a metropolitan area network (MAN), or a wide area network (WAN), such as the Internet or the World Wide Web. The network 108 and/or 110 may be any type and/or form of network and may include any of the following: a point to point network, a broadcast network, a wide area network, a local area network, a telecommunications network, a data communication network, a computer network, an ATM (Asynchronous Transfer Mode) network, a SONET (Synchronous Optical Network) network, a SDH (Synchronous Digital Hierarchy) network, a wireless network and a wireline network. The topology of the network 108 and/or 110 may be a bus, star, or ring network topology. The network 108 and/or 110 and network

topology may be of any such network or network topology as known to those ordinarily skilled in the art capable of supporting the operations of the present invention described herein.

[0072] As shown in FIG. 1, the appliance 104 is shown between the public data communication network 108 and the private data communication network 110. In other embodiments, the appliance 104 may be located on the public data communication network 108, or on the private data communication network 110. In other embodiments, the appliance 104 could be an integral part of any individual client 102a-102n or any individual server 106a-106n on the same or different network 108, 110 as the client 102a-102n. As such, the appliance 104 may be located at any point in the network or network communications path between a client 102a-102n and a server 106a-106n.

[0073] In accordance with an embodiment of the present invention, the appliance 104 includes cache management logic and also includes or has access to a storage medium which it utilizes to implement a cache memory. Using these features, appliance 104 monitors object requests made by clients 102a-102n to any of the servers 106a-106n. Objects returned from servers 106a-106n in response to these object requests are stored in the cache memory by appliance 104. Subsequent requests for the same object from any of clients 102a-102n are intercepted by appliance 104, which attempts to deliver the object from the cache rather than passing the request on to servers 106a-106n. This provides the dual benefit of reducing both the time required to respond to requests from clients 102a-102n and the load on the infrastructure supporting servers 106a-106n.

[0074] In summary, the network environment 100 depicted in FIG. 1 is presented by way of example only and is not intended to be limiting. Based on the teachings provided herein, persons skilled in the relevant art(s) will readily appreciate that the present invention may be implemented in any network environment in which object requests and responses are transferred between nodes of one or more network(s).

B. Example Appliance or Device Architecture

[0075] As will be described in more detail herein, in an embodiment of the present invention, the appliance 104 integrates caching functionality at the kernel level of the operating system with one or more other processing tasks, including but not limited to

decryption, decompression, or authentication and/or authorization. Such an implementation is illustrated in the commonly owned and co-pending U.S. Patent Application No. 11/169,002 entitled "Method and Device for Performing Integrated Caching in a Data Communications Network," filed June 29, 2005, which is incorporated by reference herein. Such an example architecture is described herein in accordance with FIG. 2, but the present invention is not so limited and other architectures may be used in practicing the operations of the present invention described herein.

[0076] FIG. 2 illustrates an example architecture 200 of an appliance 104. As noted above, architecture 200 is provided by way of illustration only and is not intended to be limiting. As shown in FIG. 2, example architecture 200 consists of a hardware layer 206 and a software layer divided into a user space 202 and a kernel space 204.

[0077] Hardware layer 206 provides the hardware elements upon which programs and services within kernel space 204 and user space 202 are executed. Hardware layer 206 also provides the structures and elements which allow programs and services within kernel space 204 and user space 202 to communicate data both internally and externally with respect to appliance 104. As shown in FIG. 2, the hardware layer 206 includes a processing unit 262 for executing software programs and services, a memory 264 for storing software and data, network ports 266 for transmitting and receiving data over a network, and an encryption processor 260 for performing functions related to Secure Sockets Layer processing of data transmitted and received over the network. In some embodiments, the central processing unit 262 may perform the functions of the encryption processor 260 in a single processor. Additionally, the hardware layer 206 may comprise multiple processors for each of the processing unit 262 and the encryption processor 260. Although the hardware layer 206 of appliance 104 is generally illustrated with an encryption processor 260, processor 260 may be a processor for performing functions related to any encryption protocol, such as the Secure Socket Layer (SSL) or Transport Layer Security (TLS) protocol. In some embodiments, the processor 260 may be a general purpose processor (GPP), and in further embodiments, may be have executable instructions for performing processing of any security related protocol.

[0078] Although the hardware layer 206 of appliance 104 is illustrated with certain elements in FIG. 2, the hardware portions or components of appliance 104 may comprise any type and form of elements, hardware or software, of a computing

device, such as the computing device 800 illustrated and discussed in conjunction with FIGs. 8A and 8B further herein. In some embodiments, the appliance 104 may comprise a server, gateway, router, switch, bridge or other type of computing or network device, and have any hardware and/or software elements associated therewith.

[0079] The operating system of appliance 104 allocates, manages, or otherwise segregates the available system memory into kernel space 204 and user space 204. In example software architecture 200, the operating system may be any type and/or form of Unix operating system although the invention is not so limited. As such, the appliance 104 can be running any operating system such as any of the versions of the Microsoft® Windows operating systems, the different releases of the Unix and Linux operating systems, any version of the Mac OS® for Macintosh computers, any embedded operating system, any network operating system, any real-time operating system, any open source operating system, any proprietary operating system, any operating systems for mobile computing devices or network devices, or any other operating system capable of running on the appliance 104 and performing the operations described herein.

[0080] The kernel space 204 is reserved for running the kernel 230, including any device drivers, kernel extensions or other kernel related software. As known to those skilled in the art, the kernel 230 is the core of the operating system, and provides access, control, and management of resources and hardware-related elements of the application 104. In accordance with an embodiment of the present invention, the kernel space 204 also includes a number of network services or processes working in conjunction with a cache manager 232. sometimes also referred to as the integrated cache, the benefits of which are described in detail further herein. Additionally, the embodiment of the kernel 230 will depend on the embodiment of the operating system installed, configured, or otherwise used by the device 104.

[0081] In one embodiment, the device 104 comprises one network stack 267, such as a TCP/IP based stack, for communicating with the client 102a-102b and/or the server 106a-106n. In one embodiment, the network stack 267 is used to communicate with a first network, such as network 108, and a second network 110. In some embodiments, the device 104 terminates a first transport layer connection, such as a TCP connection of a client 102a-102n, and establishes a second transport layer connection to a server 106a-106n for use by the client 102a-102n, e.g., the second transport layer connection

is terminated at the appliance 104 and the server 106a-106n. The first and second transport layer connections may be established via a single network stack 267. In other embodiments, the device 104 may comprise multiple network stacks, for example 267 and 267', and the first transport layer connection may be established or terminated at one network stack 267, and the second transport layer connection on the second network stack 267'. For example, one network stack may be for receiving and transmitting network packet on a first network, and another network stack for receiving and transmitting network packets on a second network. In one embodiment, the network stack 267 comprises a buffer 243 for queuing one or more network packets for transmission by the appliance 104.

[0082] As shown in FIG. 2, the kernel space 204 includes the cache manager 232, a high-speed layer 2-7 integrated packet engine 240, an encryption engine 234, a policy engine 236 and multi-protocol compression logic 238. Running these components or processes 232, 240, 234, 236 and 238 in kernel space 204 or kernel mode instead of the user space 202 improves the performance of each of these components, alone and in combination. Kernel operation means that these components or processes 232, 240, 234, 236 and 238 run in the core address space of the operating system of the device 104. For example, running the encryption engine 234 in kernel mode improves encryption performance by moving encryption and decryption operations to the kernel, thereby reducing the number of transitions between the memory space or a kernel thread in kernel mode and the memory space or a thread in user mode. For example, data obtained in kernel mode may not need to be passed or copied to a process or thread running in user mode, such as from a kernel level data structure to a user level data structure. In another aspect, the number of context switches between kernel mode and user mode are also reduced. Additionally, synchronization of and communications between any of the components or processes 232, 240, 235, 236 and 238 can be performed more efficiently in the kernel space 204.

[0083] In some embodiments, any portion of the components 232, 240, 234, 236 and 238 may run or operate in the kernel space 204, while other portions of these components 232, 240, 234, 236 and 238 may run or operate in user space 202. In one embodiment, the present invention uses a kernel-level data structure providing access to any portion of one or more network packets, for example, a network packet comprising a request from a client 102a-102n or a response from a server 106a-106n. In some embodiments, the kernel-level data structure may be obtained by the packet

engine 240 via a transport layer driver interface or filter to the network stack 267. The kernel-level data structure may comprise any interface and/or data accessible via the kernel space 204 related to the network stack 267, network traffic or packets received or transmitted by the network stack 267. In other embodiments, the kernel-level data structure may be used by any of the components or processes 232, 240, 234, 236 and 238 to perform the desired operation of the component or process. In one embodiment, a component 232, 240, 234, 236 and 238 is running in kernel mode 204 when using the kernel-level data structure, while in another embodiment, the component 232, 240, 234, 236 and 238 is running in user mode when using the kernel-level data structure. In some embodiments, the kernel-level data structure may be copied or passed to a second kernel-level data structure, or any desired user-level data structure.

[0084] The cache manager 232 may comprise software, hardware or any combination of software and hardware to provide cache access, control and management of any type and form of content, such as objects or dynamically generated objects served by the originating servers 106a-106n. The data, objects or content processed and stored by the cache manager 232 may comprise data in any format, such as a markup language, or communicated via any protocol. In some embodiments, the cache manager 232 duplicates original data stored elsewhere or data previously computed, generated or transmitted, in which the original data may require longer access time to fetch, compute or otherwise obtain relative to reading a cache memory element. Once the data is stored in the cache memory element, future use can be made by accessing the cached copy rather than refetching or recomputing the original data, thereby reducing the access time. In some embodiments, the cache memory element may comprise a data object in memory 264 of device 104. In other embodiments, the cache memory element may comprise memory having a faster access time than memory 264. In another embodiment, the cache memory element may comprise any type and form of storage element of the device 104, such as a portion of a hard disk. In some embodiments, the processing unit 262 may provide cache memory for use by the cache manager 232 of the present invention. In yet further embodiments, the cache manager 232 may use any portion and combination of memory, storage, or the processing unit for caching data, objects, and other content.

[0085] Furthermore, the cache manager 232 of the present invention includes any logic, functions, rules, or operations to perform any embodiments of the techniques of

the present invention described herein. For example, the cache manager 232 includes logic or functionality to invalidate objects based on the expiration of an invalidation time period or upon receipt of an invalidation command from a client 102a-102n or server 106a-106n. In some embodiments, the cache manager 232 may operate as a program, service, process or task executing in the kernel space 204, and in other embodiments, in the user space 202. In one embodiment, a first portion of the cache manager 232 executes in the user space 202 while a second portion executes in the kernel space 204. In some embodiments, the cache manager 232 can comprise any type of general purpose processor (GPP), or any other type of integrated circuit, such as a Field Programmable Gate Array (FPGA), Programmable Logic Device (PLD), or Application Specific Integrated Circuit (ASIC).

[0086] The policy engine 236 may include, for example, an intelligent statistical engine or other programmable application(s). In one embodiment, the policy engine 236 provides a configuration mechanism to allow a user to identifying, specify, define or configure a caching policy. Policy engine 236, in some embodiments, also has access to memory to support data structures such as lookup tables or hash tables to enable user-selected caching policy decisions. In other embodiments, the policy engine 236 may comprise any logic, rules, functions or operations to determine and provide access, control and management of objects, data or content being cached by the appliance 104 in addition to access, control and management of security, network traffic, network access, compression or any other function or operation performed by the appliance 104. Further examples of specific caching policies are further described herein.

[0087] The encryption engine 234 comprises any logic, business rules, functions or operations for handling the processing of any security related protocol, such as SSL or TLS, or any function related thereto. For example, the encryption engine 234 encrypts and decrypts network packets, or any portion thereof, communicated via the appliance 104. The encryption engine 234 may also setup or establish SSL or TLS connections on behalf of the client 102a-102n, server 106a-106n, or appliance 104. As such, the encryption engine 234 provides offloading and acceleration of SSL processing. In one embodiment, the encryption engine 234 uses a tunneling protocol to provide a virtual private network between a client 102a-102n and a server 106a-106n. In some embodiments, the encryption engine 234 is in communication with the

Encryption processor 260. In other embodiments, the encryption engine 234 comprises executable instructions running on the Encryption processor 260.

[0088] The multi-protocol compression engine 238 comprises any logic, business rules, function or operations for compressing one or more protocols of a network packet, such as any of the protocols used by the network stack 267 of the device 104. In one embodiment, multi-protocol compression engine 238 compresses bi-directionally between clients 102a-102n and servers 106a-106n any TCP/IP based protocol, including Messaging Application Programming Interface (MAPI) (email), File Transfer Protocol (FTP), HyperText Transfer Protocol (HTTP), Common Internet File System (CIFS) protocol (file transfer), Independent Computing Architecture (ICA) protocol, Remote Desktop Protocol (RDP), Wireless Application Protocol (WAP), Mobile IP protocol, and Voice Over IP (VoIP) protocol. In other embodiments, multi-protocol compression engine 238 provides compression of Hypertext Markup Language (HTML) based protocols and in some embodiments, provides compression of any markup languages, such as the Extensible Markup Language (XML). In one embodiment, the multi-protocol compression engine 238 provides compression of any high-performance protocol, such as any protocol designed for appliance 104 to appliance 104 communications. In another embodiment, the multi-protocol compression engine 238 compresses any payload of or any communication using a modified transport control protocol, such as Transaction TCP (T/TCP), TCP with selection acknowledgements (TCP-SACK), TCP with large windows (TCP-LW), a congestion prediction protocol such as the TCP-Vegas protocol, and a TCP spoofing protocol.

[0089] As such, the multi-protocol compression engine 238 of the present invention accelerates performance for users accessing applications via desktop clients, e.g., Microsoft Outlook and non-Web thin clients, such as any client launched by popular enterprise applications like Oracle, SAP and Siebel, and even mobile clients, such as the Pocket PC. In some embodiments, the multi-protocol compression engine 238 by executing in the kernel mode 204 and integrating with packet processing engine 240 accessing the network stack 267 is able to compress any of the protocols carried by the TCP/IP protocol, such as any application layer protocol.

[0090] High speed layer 2-7 integrated packet engine 240, also generally referred to as a packet processing engine or packet engine, is responsible for managing the kernel-level processing of packets received and transmitted by appliance 104 via

network ports 266. The high speed layer 2-7 integrated packet engine 240 may comprise a buffer for queuing one or more network packets during processing, such as for receipt of a network packet or transmission of a network packer. Additionally, the high speed layer 2-7 integrated packet engine 240 is in communication with one or more network stacks 267 to send and receive network packets via network ports 266. The high speed layer 2-7 integrated packet engine 240 works in conjunction with encryption engine 234, cache manager 232, policy engine 236 and multi-protocol compression logic 238. In particular, encryption engine 234 is configured to perform SSL processing of packets, policy engine 236 is configured to perform functions related to traffic management such as request-level content switching and request-level cache redirection, and multi-protocol compression logic 238 is configured to perform functions related to compression and decompression of data.

[0091] The high speed layer 2-7 integrated packet engine 240 includes a packet processing timer 242. In one embodiment, the packet processing timer 242 provides one or more time intervals to trigger the processing of incoming, i.e., received, or outgoing, i.e., transmitted, network packets. In some embodiments, the high speed layer 2-7 integrated packet engine 240 processes network packets responsive to the timer 242. The packet processing timer 242 provides any type and form of signal to the packet engine 240 to notify, trigger, or communicate a time related event, interval or occurrence. In many embodiments, the packet processing timer 242 operates in the order of milliseconds, such as for example 100ms, 50ms or 25ms. For example, in some embodiments, the packet processing timer 242 provides time intervals or otherwise causes a network packet to be processed by the high speed layer 2-7 integrated packet engine 240 at a 10 ms time interval, while in other embodiments, at a 5 ms time interval, and still yet in further embodiments, as short as a 3, 2, or 1 ms time interval. The high speed layer 2-7 integrated packet engine 240 may be interfaced, integrated or in communication with the encryption engine 234, cache manager 232, policy engine 236 and multi-protocol compression engine 238 during operation. As such, any of the logic, functions, or operations of the encryption engine 234, cache manager 232, policy engine 236 and multi-protocol compression logic 238 may be performed responsive to the packet processing timer 242 and/or the packet engine 240. Therefore, any of the logic, functions, or operations of the encryption engine 234, cache manager 232, policy engine 236 and multi-protocol compression logic 238 may be performed at the granularity of time intervals provided via the

packet processing timer 242, for example, at a time interval of less than or equal to 10ms. For example, in one embodiment, the cache manager 232 may perform invalidation of any cached objects responsive to the high speed layer 2-7 integrated packet engine 240 and/or the packet processing timer 242. In another embodiment, the expiry or invalidation time of a cached object can be set to the same order of granularity as the time interval of the packet processing timer 242, such as at every 10 ms

[0092] In contrast to kernel space 204, user space 202 is the memory area or portion of the operating system used by user mode applications or programs otherwise running in user mode. A user mode application may not access kernel space 204 directly and uses service calls in order to access kernel services. As shown in FIG. 2, user space 202 of appliance 104 includes a graphical user interface (GUI) 210, a command line interface (CLI) 212, shell services 214, health monitoring program 216, and daemon services 218. GUI 210 and CLI 212 provide a means by which a system administrator or other user can interact with and control the operation of appliance 104, such as via the operating system of the appliance 104 and either is user space 202 or kernel space 204. The GUI 210 may be any type and form of graphical user interface and may be presented via text, graphical or otherwise, by any type of program or application, such as a browser. The CLI 212 may be any type and form of command line or text-based interface, such as a command line provided by the operating system. For example, the CLI 212 may comprise a shell, which is a tool to enable users to interact with the operating system. In some embodiments, the CLI 212 may be provided via a bash, csh, tcsh, or ksh type shell. The shell services 214 comprises the programs, services, tasks, processes or executable instructions to support interaction with the appliance 104 or operating system by a user via the GUI 210 and/or CLI 212.

[0093] Health monitoring program 216 is used to monitor, check, report and ensure that network systems are functioning properly and that users are receiving requested content over a network. Health monitoring program 216 comprises one or more programs, services, tasks, processes or executable instructions to provide logic, rules, functions or operations for monitoring any activity of the appliance 104. In some embodiments, the health monitoring program 216 intercepts and inspects any network traffic passed via the appliance 104. In other embodiments, the health monitoring program 216 interfaces by any suitable means and/or mechanisms with one or more of

the following: the encryption engine 234, cache manager 232, policy engine 236, multi-protocol compression logic 238, packet engine 240, daemon services 218, and shell services 214. As such, the health monitoring program 216 may call any application programming interface (API) to determine a state, status, or health of any portion of the appliance 104. For example, the health monitoring program 216 may ping or send a status inquiry on a periodic basis to check if a program, process, service or task is active and currently running. In another example, the health monitoring program 216 may check any status, error or history logs provided by any program, process, service or task to determine any condition, status or error with any portion of the appliance 104.

[0094] Daemon services 218 are programs that run continuously or in the background and handle periodic service requests received by appliance 104. In some embodiments, a daemon service may forward the requests to other programs or processes, such as another daemon service 218 as appropriate. As known to those skilled in the art, a daemon service 218 may run unattended to perform continuous or periodic system wide functions, such as network control, or to perform any desired task. In some embodiments, one or more daemon services 218 run in the user space 202, while in other embodiments, one or more daemon services 218 run in the kernel space.

C. Caching of Dynamically Generated objects

[0095] Dynamic content, such as one or more dynamically generated objects, may be generated by servers, referred to as application or originating servers 106a-106n and/or back-end databases (not shown) that process object requests from one or more clients 102a-102n, local or remote, as depicted in FIG. 1. As those applications or databases process data, including data related to inputs received from clients, the response objects served by these databases and applications may change. Prior objects generated by those applications or databases in an originating server will no longer be fresh and therefore should no longer be stored by a cache. For example, given the same set of inputs a dynamically generated object of a first instance may be different than a dynamically generated object of a second instance. In another example, the same object may be dynamically generated with a different set of inputs

such that a first instance of the object is generated differently from a second instance of the object.

[0096] In order to achieve improved network performance, the appliance 104 is designed and configured to address the problems that arise in caching dynamically generated content through a variety of methods, as described in detail below. In some embodiments of the present invention described herein, the appliance 104 incorporates a set of one or more techniques for making the invalidation of dynamically generated content stored in the cache more efficient and effective. Furthermore, the appliance may incorporate techniques for performing control and caching for flash crowds. Cache memories typically store every response to a request for an object as long as such response is not marked as non-cacheable. As described herein, efficient caching of dynamically generated contents requires techniques that enable the timely invalidation of objects in the cache memory that have undergone a change at the originating server. Timely invalidation allows the cache to avoid serving stale content--a task of particular concern with dynamically generated content, especially where changes to the content occur irregularly. Set forth below are a number of techniques to ensure timely invalidation of dynamically generated content.

1. Integrated Functionality

[0097] In one aspect, the present invention is related to techniques of integrating functions, logic, or operations of the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238 with packet processing operations of the high-speed layer 2-7 integrated packet engine 240 responsive to the packet processing timer 242. For example, the operations of the cache manager 232 can be performed within the time intervals of the packet processing timer 242 used for packet processing operations, such as on a receipt or transmit of a network packet. In one embodiment, by integrating with the packet processing operations and/or using the packet processing timer, the cache manager 232 of the present invention can cache objects with expiry times down to very small intervals of time, as will be described in further detail below. In other embodiments, the cache manager 232 responsive to the packet processing timer 242 can also receive an invalidation command to invalidate an object within a very short time period of caching the object.

The method 300 depicted in FIG. 3A illustrates one embodiment of a technique of the present invention for requesting the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238 to perform an operation during processing or in association with the time intervals for processing a network packet by the high-speed layer 2-7 integrated packet engine or packet processing engine 240. In brief overview, at step 310 of method 300, the device 104 receives a network packet or is requested to transmit a network packet. At step 315, the device 104 requests the packet processing engine 240 to process the network packet responsive to the packet processing timer 242. As part of, or associated with, packet processing operations, at step 320, the packet processing engine 240 requests the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238 to perform an operation on a cached object. At step 325, the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238 performs the requested operation, which may include any one or combination of the techniques of the present invention described herein. In one embodiment, the cache manager 232 determines invalidation of a cached object, and marks the cached object invalid. In some embodiments, the cache manager 232 flushes the invalid object in response to a request by the packet processing engine 240. As the cache manager 232 is performing these operations responsive to the packet processing timer 242, invalidation of objects can occur within time periods in the order of milliseconds and with objects having an expiry in the order of the time intervals provided by the packet processing timer 242, such as 10 ms.

[0099] In further detail of method 300 of the present invention, at step 310, the appliance 104 receives one or more network packets, and/or transmits one or more network packets. In some embodiments, the appliance 104 requests to transmit one or more network packets over the network 108 or network 110. In another embodiment, the appliance 104 receives a network packet on one port 266 and transmits a network packet on the same port 266 or a different port 266'. In some embodiments, the packet engine 240 of the appliance 104 transmits or requests to transmit one or more network packets. In one embodiment, the appliance 104 receives or transmits a packet on a first network 108, while in another embodiment, the appliance 104 receives or transmits a packet on a second network 110. In other embodiments, the appliance 104 receives and transmits packets on the same network. In some

embodiments, the appliance 104 receives and/or transmits networks packets to one or more clients 102a-102n. In other embodiments, the appliance 104 receives and/or transmits networks packets to one or more servers 106a-106n.

[00100] At step 315, the device 104 may request or trigger packet processing operations of the packet processing engine 240 upon receipt of a network packet at the network port 266 of the device 104 or upon request to transmit a network packet from the device 104, or upon any combination of receipt and/or transmit of one or more network packets. In some embodiments, the packet processing operations of the packet processing engine 240 are triggered via a signal provided by a packet processing timer 242. In one embodiment, the packet processing timer 242 may provide interrupt-driven or event-driven timer functionality related to the receipt and/or transmission of one or more network packets. In some embodiments, the packet processing timer 242 is driven by a rate of receipt and/or transmit of network packets via the device 104, or by the rate by which each packet or a batch of packets are processed. As such, the packet processing timer 242 may be triggered and reset after each set of one or more packet processing operations. In another embodiment, the packet processing timer 242 provides time intervals, either equal or variable time intervals, to trigger, wake-up, or signal the packet processing engine 240 to perform a function or operation, such as handling a received packet or transmitting a submitted packet. As discussed above in connection with the device 104 of FIG. 2, the packet processing timer 242 may operate in the order of milliseconds, such as causing time intervals or triggering of packet processing operations at intervals of 10ms or less. The granular timer functionality of the packet processing timer of the present invention may be provided in various ways and used in operations of the packet processing operations of the packet processing engine 240.

[00101] At step 320 of method 300 of the present invention, the packet processing engine 240 requests one or more of the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238 to perform an operation. In one embodiment, the packet processing engine 240 or packet processing timer 242 generates a signal or signals to one or more of the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238. The packet processing engine 240 may request or signal the operation at any point before, during, or after a packet processing operation of a network packet, or one or more packets. In one embodiment, the packet processing

engine 240 makes the request upon trigger of the packet processing timer 242 or expiration of a time interval provided by the packet processing timer 242, and before performing a packet processing operation on a network packet. In another embodiment, during the course of performing one or more packet processing operations, the packet processing engine 240 makes the request. For example, during execution of an operation, such as within a function call, the packet processing engine 240 may make an application programming interface (API) call to one of the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238. In other embodiments, the packet processing engine 240 makes the request upon completion of a network packet processing operation.

[00102] At step 325, the requested operation is performed by one or more of the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238. In some embodiments, any functionality or operation provided via the kernel 204 may be requested to be executed, such as via a kernel application programming interface (API). As such, any of the functions of the device 104 may be performed in conjunction with the timing or timing intervals of packet processing via the packet processing timer 232. In some embodiments, the requested operation is performed synchronously and in conjunction with the packet processing operations of the packet processing engine 240. For example, the packet processing operations wait and continue upon a completion of, or response from, the requested operation. In other embodiments, the requested operation is performed asynchronously with the packet processing operations. For example, the packet processing engine 240 sends a request to perform the operation but does not block or wait to receive a response from the operation. As will be discussed in further detail in conjunction with method 350 of the present invention depicted in FIG. 3B, the packet processing engine 240 may request the cache manager 232 to perform any cache management function, such as checking for expiry or invalidation of objects, marking objects as invalid, or flushing invalid or expired objects.

[00103] In some embodiments, the packet processing engine 240 at step 320 sends multiple requests, such as a first request to the cache manager 232 and a second request to the encryption engine 234. In other embodiments, the packet processing engine 240, at step 320, sends a single request comprising multiple requests to be distributed by the device 104, such as via the kernel 230 to the intended component of the device 104. In one embodiment, the requests are communicated subsequent to

each other. In another embodiment, requests may be dependent on the status, result, success, or completion of a previous request. For example a first request to the policy engine 236 may be used to determine a policy for processing a network packet from another device or a user associated with the network packet. Based on a policy of the policy engine 236, a second request to the cache may be made or not made depending on a result of the first request. With the cache manager 232, policy engine 236, encryption engine 234, and/or the multi-protocol compression engine 238 integrated in the kernel space 204 of the device 104 with the packet processing engine 240, there are various operations of the device 104 as described herein that may be triggered by and integrated with packet processing operations.

2. Invalidation Granularity

[00104] In another aspect, the present invention is related to and incorporates the ability to configure the expiration time of objects stored by the cache to fine granular time intervals, such as the granularity of time intervals provided by the packet processing timer. This characteristic is referred to as "invalidation granularity." As such, in one embodiment, the present invention can cache objects with expiry times down to very small intervals of time. In other embodiments, the cache manager responsive to a packet processing timer can also receive an invalidation command to invalidate an object within a very short time period of caching the object. By providing this fine granularity in expiry time, the cache of the present invention can cache and serve objects that frequently change, sometimes even many times within a second. One technique is to leverage the packet processing timer used by the device of the present invention that is able to operate at time increments on the order of milliseconds to permit invalidation or expiry granularity down to 10 ms or less. Traditional caches, by contrast to the present invention, typically do not set or have expiry or invalidation granularity of less than one second.

[00105] Referring now to FIG. 3B, an embodiment of a method 350 of the present invention is depicted for invalidating or expiring a cached object responsive to the packet processing timer 242 and/or packet processing engine 240. As such, in some embodiments of the present invention, cached objects can be invalidated or expired in the order of milliseconds, such as 10ms or less. In overview, at step 355 of method 350, the cache manager 232 receives a signal or request to perform an operation via

the packet processing engine 240 in response to the packet processing timer 242. At step 360, the cache manager 232 determines if a cached object, such as a dynamically generated object, is invalid or expired. At step 365, if the object is invalid, the cache manager 232 marks the object as invalid, and at step 370, flushes the invalid object from the cache manager 232.

[00106] In further detail of step 355, in some embodiments, the cache manager 232 may be signaled or requested to perform a cache related operation at any point of time during network packet processing. In one embodiment, at step 355, the cache manager 232 receives an operation request prior to the processing of a network packet received or to be transmitted by the device 104. In another embodiment, the cache manager 232 receives an operation request upon the completion of processing of a network packet. For example, the packet processing engine 240 completes processing of a network packet, and before either waiting for the next time interval of the timer 242 or before processing the next packet, requests the cache to perform an operation. In other embodiments, during an operation of packet processing, the packet processing engine 240 communicates an operation request to the cache manager 232. In another embodiment, the cache manager 232 receives a signal, such as from the packet processing engine 240 or packet processing timer 242 to trigger the cache manager 232 to perform an operation. In some embodiments, the signal indicates to invalidate a cached object or to expire an expiry of a cached object.

[00107] In some embodiments, the cache manager 232 may receive a request to perform a cache operation from an entity external to the cache manager 232, such as a request to invalidate an object communicated by a server 106a-106n, and processed by the packet processing engine 240. In one embodiment, the cache manager 232 may receive an invalidation request within 10 ms or less of caching the object, while in another embodiment, as short as 5ms, 2ms or 1ms. In other embodiments, the cache manager 232 may perform a cache operation responsive to the operations or functionality of the cache manager 232, such as the expiration of a timer to cause an object to be invalidated or during the processing of any cache command. In other embodiments, the cache manager 232 uses the packet processing timer 242 of the device 104 to trigger cache operations. For example, the timer 242 may trigger or signal the cache to check for invalidation or expiry of a cached object at any time interval capable of being set by the timer 242. In one embodiment, the timer 242 may be set to trigger or signal the cache within 10ms or less of being set, or in another

embodiment, as short as 5ms, 2ms, or 1ms of being set. In some embodiments, the originating server 106a-106n may set the expiry time of the object. In other embodiments, the appliance 104 or client 102a-102n may set the expiry time of the object.

[00108] At step 360, the cache manager 232 determines the invalidation or expiry of an object stored in cache. In some embodiments, an object in cache is invalidated based on the expiration of a timer. In one embodiment, the cache manager 232 may issue an invalidation command on an object based on the expiration of a timer. In another embodiment, the object stored in cache is automatically invalidated by the cache manager 232 responsive to the expiration of a timer, such as a timer set with the packet processing timer 242. In some embodiments, responsive to the packet processing timer 242, the cache manager 232 checks for the expiration of any timers for cached objects. In one embodiment, the cache manager 232 determines an object timer has expired, while in another embodiment, the cache manager 232 determines the object timer has not expired. In a further embodiment, the cache manager 232 responsive to a second trigger or second timer interval of the packet processing timer 242 will check a second time if a previously checked object timer has expired.

[00109] In some embodiments, the cache manager 232 parses, interprets, accesses, reads or otherwise processes an invalidation command or request to identify the object to invalidate in the cache. In one embodiment, an entity external to the cache manager 232 issues an invalidation command to the cache manager 232 to invalidate the object. In another embodiment, the external entity may issue the invalidation command responsive to a packet processing timer 242. If the object is valid and/or has not been invalidated, the cache manager 232 invalidates the object responsive to the request. In some embodiments, the invalidation request processed by the cache manager 232 is responsive to the packet processing operations of the packet processing engine 240 processing the request, which in turn may also be responsive to the packet processing timer 242.

[00110] At step 365, the cache manager 232 marks the object as invalid. The cache manager 232 may mark each object as invalid in any suitable or desired manner. In one embodiment, an object is marked as invalid by setting a flag, attribute, or property of the stored object. For example, a flag may be set to any value identifying to the cache manager 232 the object is invalid. In another embodiment, an object may be marked as invalid by moving the object to an area or portion of the cache for storing

invalid objects. In other embodiments, the cache manager 232 may identify or track the invalid and/or valid state of a stored object by a database or a linked list or any type and form of data structure. In some embodiments, the cache manager 232 uses one or more objects to identify or track the validity or invalidity of one or more objects stored in cache. In another embodiment, the object is marked as invalid by changing, modifying or altering the stored object, for example deleting or removing a portion of the object so that it may not be used, or by changing or mangling the name of the object.

[00111] At step 370, the cache manager 232, in some embodiments, flushes from the cache those objects marked as invalid. In another embodiment, the cache manager 232 flushes the invalid object from cache upon request for the object, such as by a client 102a-102n. In some embodiments, the cache manager 232 overwrites the invalid object with an updated copy or version of the object received after invalidation or expiration of the object. In another embodiment, the cache manager 232 reuses the cache memory occupied by the invalid object by storing another to the same portion of cache memory. In yet another embodiment, the cache manager 232 does not flush the object marked as invalid but keeps the object stored in memory or storage of the cache.,

[00112] Although method 350 describes invalidation and flushing of cached objects responsive to a packet processing timer and/or in conjunction with packet processing operations to provide invalidation granularity, any operation of the cache and any techniques of cache management as well as any other operation of the device 104 described herein may be executed at fine granular time intervals provided by the packet processing timer. In some embodiments, the invalidation or expiration of cached objects can occur as short as a 100ms time interval, while in another embodiment, as short as a 50ms time interval. In some embodiments, the invalidation or expiration of cached objects can occur as short as 25 ms time interval, and in other embodiments, as short as a 10 ms time interval. While in other embodiments, the invalidation or expiration of cached objects can occur as short as a 5 ms time interval, and still yet in further embodiments, as short as a 3, 2, or 1 ms time interval.

[00113] By incorporating the capacity to invalidate objects after the elapse of very small increments of time as described in methods 300 and 350 in conjunction with FIGs. 3A and 3B above, improved caching of dynamically generated content is enabled. Some dynamic content is in fact amenable to being stored and served from a

cache for very short periods of time. To successfully cache such content, however, an approach in accordance with an embodiment of the present invention provides caching objects for very short periods of time before the object is invalidated and flushed from the cache memory. For example, certain dynamically generated objects may be cacheable for as long as 1 second but anything longer is frequently unacceptable for content that is constantly changing. In an embodiment, the approach of the present invention included invalidating or expiring cached content after small fractions of a second. As an example, if an application 100 takes milliseconds to generate a dynamic response, then the cache can store and serve that response for a duration of less than or equal to the period of 100 milliseconds, without compromising the freshness of the data. There will not be a new object generated during that 100 millisecond period because it is shorter than the time it takes to generate a new object. The appliance 104 can thus be set up to serve the prior object during that duration. The ability of the appliance 104 to invalidate down to very small increments of time is frequently very useful for application environments where the database transaction isolation level is set to allow Repeatable Reads or Serialized Reads.

3. Invalidation Commands

[00114] Traditional caching technology invalidates stored content based on a pre-defined expiry time for the content, which is typically configured either by the administrator or is received from the server that served the object. Described below is another technique of the present invention for invalidating content in order to more efficiently cache dynamically generated content. The technique of the present invention includes the ability to receive at the appliance 104 an invalidation command that identifies one or more of the previously stored objects in the cache as invalid in real time. For example, the invalidation command may be communicated via a network packet transmitted to the client or an application programming interface (API) call made by a server to the appliance. This differs from the traditional approach by which the server simply sets a cache expiry time that it includes in the object header at the time the object is served.

[00115] The technique of the present invention is more specifically illustrated in FIGS. 4A and 4B. FIG. 4A is a flow chart illustrating a method for maintaining a cache,

such as a computer memory cache. In brief overview and according to step 410, dynamically generated objects previously served from an originating server 106a-106n are stored in the cache. For example, the dynamically generated object may not be identified as cacheable or otherwise include any cache or cache control information. At step 420, an invalidation command is received at the cache or cache manager 232. The invalidation command identifies one or more previously served objects as invalid. As step 430, in response to the invalidation command, the cache or cache manager 232 marks the identified object as invalid.

[00116] In further detail at step 410, the cache manager 232 stores in a cache memory element a dynamically generated object received, obtained or communicate from any source. In some embodiments, the dynamically generated object may be generated and served from a server 106a-106n. In other embodiments, the dynamically generated object may be generated and communicated by a client 102a-102n. In some embodiments, another portion, component or process of the appliance 104 generates the object and stores the object in the cache. In further embodiments, the dynamically generated object may be generated by another appliance 104 or another computing device on the network and transmitted or communicated to the appliance 104. In some embodiments, the dynamically generated object is not identified as cacheable or identified as non-cacheable. In other embodiments, the dynamically generated object is identified as cacheable or is under cache control.

[00117] At step 420, the cache manager 232 receives an invalidation command identifying an object to invalidate, such a dynamically generated object stored in the cache. In one embodiment, the invalidation command may comprise any type of directive or instruction indicating to the cache that an object is invalid or otherwise may be stale. In some embodiments, the invalidation command identifies the object and may also identify the time at which the object is invalid as well as what portions of the object may be invalid. In one embodiment, the cache manager 232 provides an application programming interface (API) that may be called remotely by an originating server 106a-106n. In some embodiments, the cache manager 232 may provide any type and form of protocol for receiving commands and replying to commands via one or more network packets. In one embodiment, the cache manager 232 or device 104 provides an Extensible Markup Language (XML) API interface for receiving and processing invalidation commands. For example, the cache manager 232 may provide a web service interface. In some embodiments, the cache manager

232 replies to the invalidation command by sending an acknowledgement, status or other response to the originating server 106a-106n. In other embodiments, the cache manager 232 does not reply to the invalidation command. In one embodiment, an object is marked as invalid if an application running in an originating server 106a-106n performed an action that made the stored object stale, such as by generated a new or updated version of the object. This could occur, for example, when news editors make changes to a fast developing news story and therefore want to be assured the most recent version of the story is being served to clients.

[00118] Invalidation commands may be issued from an originating server by the application that generated the object, by another server 106a-106n or another appliance 104. In one embodiment, the originating server 106a-106n issues or communicates an invalidation command to the cache 232 automatically in response to a change to the dynamically generated object on the originating server 106a-106n. The invalidation command can also be generated by an administrative control outside or external to the server 106a-106n and the appliance 104. For example, the administrative control may be any type and form of program or application running on the network and in communication with the appliance 104, such as administrator console. Furthermore, a client 102a-102n could issue or communicate an invalidation command to the appliance 104 or cache manager 232. For example if the client were to take action that the client 102a-102n recognizes would cause a change to the requested objects at the originating server, the client may communicate the invalidation command. Any object stored in the cache can be invalidated by the transmission to the cache of a user command executed locally at the cache or invoked remotely using the XML API infrastructure.

[00119] According to step 430, an object stored in cache, e.g., a previously served dynamically generated object, that has been identified as invalid is marked as such in response to the invalidation command. An invalid object will not be provided to a requesting client from the cache, but instead would be served directly from the originating server. The cache manager 232 may mark each object as invalid in any suitable or desired manner. In one embodiment, an object is marked as invalid by setting a flag, attribute, or property of the stored object. For example, a flag may be set to any value identifying to the cache manager 232 the object is invalid. In another embodiment, an object may be marked as invalid by moving the object to an area or portion of the cache for storing invalid objects. In other embodiments, the cache

manager 232 may identify or track the invalid and/or valid state of a stored object by a database or a linked list or any type and form of data structure. In some embodiments, the cache manager 232 uses one or more objects to identify or track the validity or invalidity of one or more objects stored in cache. In another embodiment, the object is marked as invalid by changing, modifying or altering the stored object, for example deleting or removing a portion of the object so that it may not be used, or by changing or mangling the name of the object.

[00120] In some embodiments, the appliance 104 subsequently flushes from the cache those objects marked as invalid. In another embodiment, the appliance 104 flushes the invalid object from cache upon request for the object, such as by a client 102a-102n. In some embodiments, the appliance 104 overwrites the invalid object with an updated copy or version of the object. In another embodiment, the appliance 104 reuses the cache memory occupied by the invalid object by storing another dynamically generated object to the same portion of cache memory.

[00121] With the command invalidation API of the cache manager 232 of the present invention, any computing device or user in communication with the appliance 104 may request to invalidate an object, such as a dynamically generated object, stored in the cache. As such, the invalidation of objects stored in cache can be controlled real-time instead of using pre-determined configuration expiry or invalidation time periods. Thus, using these techniques the longevity of the cached objects can be controlled from external application processing nodes such as databases or originating application servers. For example, the appliance 104 can be configured to work with a database such that a change to the database automatically triggers an invalidation command from the database (or application) to the appliance 104 to flush a particular object or objects.

4. Invalidation of Groups Using Invalidation Command

[00122] In a further embodiment of the present invention, the appliance 104 identifies and invalidates at the same time a group of objects stored by the cache. Objects stored in a traditional cache memory are each treated individually and separately by the cache in determining whether the object is stale. As each object reaches its specified expiry time (generally as set by the server and stored by the cache in a table) that item is flushed from cache memory. This traditional approach is inefficient and

ultimately insufficient, however, to successfully handle the challenges that arise in attempting to cache dynamically generated content.

[00123] FIG. 4B illustrates another embodiment of a method of the present invention for maintaining a cache, such as a computer memory cache, wherein the appliance 104 has the ability to create, store, and invalidate groups of related objects that have been previously served from an originating server 106a-106n. In brief overview, at step 410, an object, such as a dynamically generated object served from an originating server 106a-106n is stored in the cache. At step 412, the cache manager 232 forms a group of previously served objects stored in the cache. In one embodiment, the group may be associated with or identified by one or more object determinants as will be described in further detail below. At step 414, the cache manager 232 maintains a record of the group of objects. At step 422, the cache manager 232 receives an invalidation command to invalidate the group of objects. At step 432, the cache manager 232 marks the group of objects as invalid in response to the invalidation command.

[00124] Step 410 is the same as in FIG. 4A, wherein an object is stored in the cache of the appliance 104, such as dynamically generated objects previously served from an originating server 106a-106n. In some embodiments, one or more of the objects may not be identified as cacheable, or otherwise may not have any cache or cache control information. For example, the server 106a-106n may assume the dynamically generated objects will not be cached.

[00125] According to step 412, the appliance 104 forms a group out of a set of the objects previously served from the originating server 106a-106n and stored in the cache. Any suitable or desired set of objects may be associated with each other to form a group. For example, any dynamically generated objects generated for, or associated with, serving a web page may form a group. In some embodiments, an object may be associated with multiple groups. In other embodiments, one group of objects may form a subset of another groups of objects. In some embodiments, the formed group of objects have objects served from the same server 106a-106n, while in other embodiments, the formed group of objects have objects served from different servers 106a-106n. In further embodiments, the formed group of objects may comprise objects from a client 102a-102n, objects from a server 106a-106n, or objects generated by or served from both clients 102a-102n and servers 106a-106n. In one embodiment, one object in the group is static while another object in the group is

dynamically generated. In some cases, one object in the group is not identified as cacheable while another object in the group is identified as cacheable. In other cases, the objects in the group may be logically related in accordance with functionality or application provided by a server 106a-106n. In another case, the objects in the group may be related as associated with the same client 102a-102n or the same user.

[00126] In step 414, a record of the group of objects is maintained. Various techniques for recording and maintaining a record of a group of objects, or otherwise associating objects, may be used in practicing the operations of the present invention described herein. In one embodiment, the record may be maintained directly in, for example, a look-up table. In another embodiments, the records could be represented in a hash-table format. In some embodiments, the cache manager 232 maintains the association of objects in a database, or a data structure or object in memory. In further embodiments, a flag, property or attribute of each object in the group is assigned or set to a value identifying the group, such as a value equal to, identifying, or referencing the name or identifier of the group, such as a group's object determinant that will be described in more detail below. In some embodiments, a group of objects is arranged, placed or located in a portion of cache memory identified as holding the group

[00127] In step 422, an invalidation command is received at the appliance 104 or cache manager 232. According to the embodiment described in FIG. 4B, the invalidation command identifies that one or more objects are invalid, or otherwise are stale. In some embodiments, the invalidation command references, identifies or specifies a name or identifier of the group of objects. In one embodiment, the invalidation command comprises a single invalidation request to invalidate all the objects in the group. In another embodiment, the invalidation command identifies one object in the group to invalidate. In other embodiments, the invalidation command comprises a plurality of invalidation request to invalidate a plurality of objects in the group

[00128] According to step 432, the group of previously served objects is marked as invalid if the invalidation command references, identifies, or specifies an object of the group as invalid, each object in the group as invalid, or the group as invalid. In some embodiments, if the invalidation command identifies an object in the group as invalid, the cache manager 232 marks the object as invalid. In other embodiments, if the invalidation command identifies an object in the group as invalid, the cache manager

232 marks the group of objects as invalid or each object in the group as invalid. In yet further embodiments, the cache manager 232 may only invalidate the group of objects when a plurality of objects are identified as invalid via one or more invalidation commands. In another embodiment, the invalidation command may specify a name or identifier of the group, and the cache manager 232 marks the group as invalid, or each object in the group as invalid.

[00129] In one embodiment, the appliance 104 or cache manager 232 flushes from the cache memory a group of objects that has been marked as invalid. In some embodiments, the objects in the group may be flushed from cache memory only when each object in the group is marked as invalid. In other embodiments, if one object of the group has been marked as invalid then the entire group is flushed. In another embodiment, the group of objects, or any object in the group, marked as invalid may be flushed upon receipt of a request for the group of objects, or any object in group, by a client 102a-102n. In other embodiments, the group of objects, or any object in the group, marked as invalid may be flushed upon receipt of a response from a server 106a-106n provide one or more new objects in the group.

[00130] An example of the above described embodiments follows. Customer resource management ("CRM") applications are used by many businesses to track and evaluate all aspects of resource management. Often, CRM applications are implemented and accessed over private and public networks including the Internet. These applications, which provide access to large amounts of data that is frequently being accessed, thus benefit from caching the data generated by such applications. For example, sales reports are frequently generated and served to remotely connected users. These sales reports are built by the relevant application through compiling data from sales information that is posted to such application servers and/or their underlying databases. As many users request the same document (i.e., a certain sales report), without caching, the application server must re-generate the object for each request. If, however, such objects can be stored in the cache, then application and database processing is conserved, including potentially valuable bandwidth, as the cache is placed closer to the requesting clients.

[00131] The challenge for caching such objects arises because each time a new sale is posted to the application running at the originating server (or to its underlying database), the information in the sales report needs to be updated. As a result, all sales reports that may have been stored in any caches supporting these application

servers must be invalidated and the content flushed out of cache memory. The traditional approach to caching, however, has no way of accurately determining when the change to the underlying database or application is going to occur and therefore cannot reasonably evaluate the freshness of dynamic content. Every time a change occurs in database or application or originating server, the cache has to be able to identify that the change has been made, and which group of objects should be invalidated as a consequence of such change. Generation of invalidation commands that contain object determinants linked to groups of previously served objects, as described above, can meet this need.

[00132] Multiple groups of related objects may be formed at a single hierarchical level. Alternatively, sub-groups of objects may be formed to create multiple hierarchical levels. In an embodiment, the groups or sub-groups of objects may be pre-designated by a user. In another embodiment, a user may establish rules by which the appliance 104 automatically forms groups of related objects, and associates object determinants therewith.

5. Identification of Object Determinants in a Client Request or Response

[00133] An embodiment of the present invention also addresses the need to be able to identify all objects affected by a state change at the originating application server 106a-106n (and/or underlying database) by generating groupings of objects and implementing parameterized invalidation. In this embodiment, any object or pre-defined group of objects can be invalidated by an intercepted HTTP request, for example from a client, that the cache parses in order to identify an object determinant. The term "object determinant" refers to any information, data, data structure, parameter, value, data pattern, request, reply, or command that references, identifies or specifies one object or a set of objects, uniquely or otherwise. In some embodiments, an object determination is a pattern of bytes or characters in a communication that may be associated with an object or used to uniquely identify that the communication is associated with, or referencing, the object. In one embodiment, an object determinant indicates whether change has occurred or will occur, in the originating server, to a group of previously served objects stored in the cache manager 232 with which the object determinant is associated. In some embodiments, the objects in a group of objects are related in that they are associated with at least one

object determinant. Specific, non-limiting examples of object determinants and further illustrations of their use are described more fully below.

[00134] In some embodiments of the present embodiment, object determinants are certain pre-defined parameters or data structures included or embedded in a client request or response. In other embodiments, the client 102a-102n, server 106a-106n or appliance 104 embeds in a communication one or more object determinants, such as pre-defined strings or sets of characters representing the object determinant. The object determinants indicate whether such request will have the effect of causing a change in the state of objects stored in the originating server 106a-106n or databases linked thereto. In one embodiment, the existence of the object determinant in a request indicates a change has or will occur to an object. In another embodiment, the syntax, structure, parameter, or value of the object determinant indicates a change has or will occur to an object. In an embodiment, the cache receives an object request from a client 102a-102n. The request may include certain parameters or values (object determinants) that the cache recognizes will change the state of the originating server or application server which will, as a consequence, make stale certain related objects stored by the cache manager 232 that had been previously generated by such originating server or application server 106a-106n. Depending on the invalidation policy set by the user, the parameters (object determinants) may require invalidation of one or more previously served objects or group of objects, by the originating server, that have been stored by the cache. The cache is configured to identify the relevant objects that will be effected by this state change (i.e., those objects or groups of objects linked to the object determinant), and invalidate these objects via the method marking each of the objects as invalid and/or flushing such objects from the cache memory.

[00135] The above described technique is illustrated in FIG. 4C. As with other embodiments described herein, step 410 comprises storing, in the cache, objects, such as dynamically generated objects previously served from an originating server. The objects could be generated by an application running on the originating server 106a-106n, or could be drawn, for example, from a database accessed by the originating server 106a-106n. In some embodiments, the dynamically generated objects are identified as not cacheable or otherwise not identified as cacheable.

[00136] According to step 421, the cache intercepts or otherwise receives a communication between the client and the server, such as a request from a client or a

response from a server. In some embodiment, the request is for a specified object, the object having been previously served and stored in the cache. In another embodiment, the communication includes a response from a server having a requested object. In one embodiment, such receipt or interception occurs according to established caching protocol and communications standards. Although the cache manager 232 or appliance 104 may be generally described as receiving a request, response or communication, in receiving such request, response or communication, the cache 232 or appliance 104 may intercept or obtain by any suitable means and/or mechanisms the request, response or communication even though not communicated directly or explicitly to the cache.

[00137] In step 423, an object determinant is identified in the intercepted communication. The cache manager 232 may extract, interpret, parse, access, read, or otherwise process the intercepted communication to determine or identify one or more objects determinants in the communications. Any parameter, value, syntax, data, structure or set of one or more characters of the communication may be used to identify an object determinant. In one embodiment, the cache manager 232 may identify the name or identifier of an object in a request from the client 102a-102n to the server 106a-106n, in which the client requests the object. In another embodiment, the cache manager 232 may identify the name or identifier of a first object in the request of the client 102a-102n or response from the server 106a-106n that indicates a change has occurred or will occur to a second object stored in the cache. In other embodiments, the cache manager 232 determines if any patterns of characters in the request match any object determinants associated with an object or group of objects in the cache. In some embodiments, an object determinant may be determined for an object not currently stored in cache. In other embodiments, an object determinant may be determined for an object currently marked as invalid. In other embodiments, an object determinant for a requested object is determined to be associated with an object determinant of a cached object. In yet another embodiment, upon the first reference, request, or response for an object in a communication, the cache manager 232 establishes the identified object determinant as the object determinant for the object.

[00138] By receiving and parsing the communication, such as a client request or server response, to identify an object determinant, the cache manager 232 or appliance 104 may effectively determine whether to mark as invalid a cached object that has been

associated with the identified object determinant. Thus, according to step 425, a determination is made as to whether the object determinant indicates a change to the cached object. In some embodiments, the identified object determinant may be part of a communication that does not alter, modify or generate an object. In other embodiments, the identified object determinant is a part of a communication that indicates a change has occurred or will occur to the object associated with the object determinant. For example, the communication may be a get request for a dynamically generated object or a submit request that will change the data used for one or more dynamically generated objects. In some embodiments, the existence of the object determinant in the communication indicates a change has or will occur on one or more objects. In another embodiment, the type or name of a command, directive or instruction in the communication along with the object determinant indicates a change has or will occur on one or more objects. In yet a further embodiment, the existence, value or setting of a parameter or variable of a command, directive or instruction indicates a change has or will occur on one or more objects associated with an object determinant.

[00139] In other embodiments, the cache manager 232 performs a hash function, algorithm, or operation on the intercepted communication or object determinant to determine if a change has occurred in the object. In some embodiments, the hash value is compared with a previous stored hash value for the object and if different then the cache manager 232 recognizes the object has changed. In yet another embodiment, a hash value for the object may be included in the communication or object determinant. In one embodiment, the communication indicates the object has changed by the value or setting of a parameter, such as with a Boolean flag. In other embodiments, an entity tag control and validation mechanism as will be described in more detail below may be used to identify the object and determine if the object has changed.

[00140] If a change is indicated, then at step 431, then the object associated with or identified by the object determinant is marked as invalid. In some embodiments, an object requested by the intercepted communication is marked as invalid in accordance with step 431, and retrieved from the originating server 106a-106n in accordance with step 440. Otherwise, in other embodiments, the requested object is retrieved from the cache in accordance with step 450. In one embodiment, any object marked as invalid will be flushed from the cache.

6. Invalidation of Groups of Objects Based on Object Determinants

[00141] The above embodiment of the present invention describes the case of invalidating a previously served object in the cache manager 232 based on identification of an object determinant in the client request. This general concept may also be used, in another embodiment, to identify and invalidate a group of objects with which one or more object determinants have been associated. This embodiment is illustrated in FIG. 4D.

[00142] The method described in FIG. 4D begins in the same fashion as the method of FIG. 4C. Step 410 comprises storing, in the cache, objects, such as dynamically generated objects previously served from an originating server. In some embodiments, one or more of the objects are not identified as cacheable. According to step 412 and similar to FIG. 4B, previously served objects are formed into groups. In one embodiment and in accordance with the object determinant technique of the present invention, a group of objects is associated with or identified by at least one object determinant. As described more fully below, in some embodiments, the association of groups with object determinants depends on the nature and details of the users caching policy, such as a policy defined, controlled or used by the policy engine 236. In other embodiment, the one or more object determinant of the group comprises the one or more object determinants of the objects in the group. In another embodiment, the object determinant of the group comprises a combination of object determinants of objects in the group.

[00143] According to step 414, a record is maintained of the group, along with its associated object determinants, if applicable. This step is similar to step 414, illustrated in FIG. 4B. In one embodiment, the record and/or any object determinants of the group is maintained in a look-up table. In other embodiments, the record and/or any object determinants of the group may be maintained in a hash-table format. The hash-table may be designed to efficiently store non-contiguous keys that may have wide gaps in their alphabetic and numeric sequences. In another embodiment, an indexing system can be built on top of a hash-table. In some embodiments, the cache manager 232 maintains the association of objects as a group with one or more object determinants in a database, or a data structure or object in memory. In further embodiments, a flag, property or attribute of each object in the group is assigned or

set to a value identifying the group, such as a value equal to, identifying, or referencing the name or identifier of the group, or a group's object determinant. In some embodiments, a group of objects is arranged, placed or located in a portion of cache memory identified as holding the group. In another embodiment, the one or more object determinants are stored in association with the group of objects.

[00144] Steps 421 and 423 are similar to steps 421 and 423 as illustrated in FIG. 4C. According to step 421, the cache manager 232 or appliance 104 intercepts or otherwise receives a communication between the client 102a-102n and server 106a-106n, such as a request from a client for an object previously served and stored in the cache. In one embodiment, the cache manager 232 intercepts a request from the client 102a-102n to the server 106a-106n. In some embodiments, the request is for an object stored in cache. In other embodiments, the request is an instruction, command or directive to the server 106a-106n that will cause a change to an object stored in cache, such as to cause an object to be dynamically generated. In another embodiment, the cache manager 232 intercepts a response from a server 106a-106n to the client 102a-102n comprising or identifying an object stored in cache.

[00145] In step 423, an object determinant is identified in the intercepted communication. As noted above, the object determinant indicates whether a change has occurred or will occur in the requested object, at the originating server 106a-106n. However, in the embodiment of FIG. 4D, the object determinant may be associated with a group of objects. This enables efficient invalidation of all objects stored in the cache that may be affected by a particular object determinant. In some embodiments, an object determinant of an object in the group is identified. In other embodiments, an object determinant, for example, a group object determinant, for the group of objects is identified. In another embodiment, a combination of object determinants of one or more objects in the group are identified.

[00146] Thus, according to step 427, a determination is made as to whether the object determinant indicates a change in the group of previously served objects. In some embodiments, the existence of the object determinant of the group in the intercepted communication indicates a change has occurred or will occur to one or more, or all of the objects in the group. In other embodiments, the name and type of a command, directive or instruction in the intercepted communication indicates such changes. In yet another embodiment, the existence, value or setting of any parameters or variables in the communication may also indicate such changes.

If at step 427, the object determinant indicates a change in the group, then the group of previously served objects is marked as invalid in the cache in accordance with step 435. In some embodiments, one or more, or all of the objects of the group are requested and retrieved from the originating server 106a-106n in accordance with step 440. If at step 427, the object determinant does not indicate a change in the group, then in some embodiments, any objects requested as part of intercepted communication and previously served and stored in the cache is retrieved from the cache manager 232 in accordance with step 450. In an embodiment, any object or group of objects marked as invalid may be flushed by the cache manager 232 from the cache.

7. Designation of Groups

[00148] The cache administrator may specifically designate which objects get included into a particular group. Whenever an object is stored in the cache, the administrator may make that object a member of one of the configured or implicit groups depending on the configuration. The configured groups can be based on configurations that an administrator has previously established or alternatively based on application behavior and other data related to object invalidation. An object may also be part of an implicit group if its configured group is dynamic. Objects in the implicit group are grouped by the value of the significant invalidation parameters.

[00149] By permitting very flexible grouping of objects, a cache can achieve a level of flexibility and coordination in invalidation that is necessary to effectively cache dynamically generated content. The cache can invalidate a very specific group of objects simultaneously, thereby making the cache more responsive to the frequent need to invalidate dynamically generated content. At the time the cache assigns an object to a group, the group determines a number of things relative to that object, including the invalidation parameters and the hit determinants, in order to associate one or more object determinants therewith.

[00150] In the customer resource management ("CRM") example, the cache administrator may pre-designate each of the groupings. For example, the administrator configures the cache to group each of the sales departments by name. Thus the administrator can designate an auto department, a motorcycle department, etc., and each time an object determinant is recognized in a request coming to the

cache, the cache can then invalidate all objects stored in a designated group linked to an appropriate department via the object determinant.

8. Ruled-Based Grouping

[00151] Alternatively, the cache administrator may establish rules that allow the cache appliance to determine on the run which objects to include in a particular group or groups. Such rules-based groupings may rely on the designation of groups by virtue of established rules that link the object to significant object determinants that the cache utilizes to create the relevant groups. An example of this approach may involve configuring the cache with rules that the cache uses to recognize what objects to put in each group.

[00152] Again turning to the CRM example, a rule may state that each subdivision of the Sales Department that is set up on the application should be recognized by the cache as its own grouping. In this way the groupings can be created without the cache administrator having to specifically identify each grouping but allows the cache to determine based on the relevant rules. This technique creates a more flexible and often less work intensive way to designate groupings. The cache administrator could configure a rule that states that every subdivision department of Sales (i.e., sales\auto, sales\motorcycle etc.) should generated a new grouping by the cache. As a request from the Auto Sales Department is processed and returned by the application via the cache, the cache can recognize each subgrouping of sales and automatically create a grouping for it, based on the pre-configured rule.

[00153] The rule may be implemented by the cache each time it sees a new request for an object of the type report/sales/auto or report/sales/motorcycle, etc. This process can then be repeated when a Motorcycle Sales Department request showing that it is a sub-grouping of the Sales Department, then the Bicycle Sales Department and so forth, as the cache recognizes these subgroups and establishes an object grouping for each of them. When a known invalidation request comes to the cache linked to one of these groupings, or if a relevant object determinant is identified in a client request (for example a post of a sales report to the Motorcycle Sales Department sales/motorcycle found in the parsing the request), the cache knows to invalidate all the cached objects in the Motorcycle Sales Department Grouping.

[00154] In this way, when a cache recognizes that a change has occurred or will occur to data served by the application (either because the cache recognizes that contents of a request received by the cache will trigger a change at the application or because of the occurrence of some outside change), the above technique enables the cache to quickly and simply identify which objects require invalidation through the process of grouping. In this way, the cache is able to invalidate large numbers of dynamically generated objects that are no longer fresh because of changes in the application or database state.

[00155] The ability of the cache to successfully store and serve out of its cache memory dynamically generated content can also be enhanced with an intelligent statistical engine that examines the pattern of request and response traffic in order to determine, over a period of time, the set of objects that would provide the most caching benefit. The engine can either be integrated into the cache appliance itself, or run in a separate computer as a heuristic to select some subset of objects for further investigation to determine suitability for dynamic caching.

9. Further Use of Object Determinants

[00156] As described above, object determinants may be any data structure that indicates whether a change has occurred or will occur, in the originating server, to the group of previously served objects stored in the cache with which the object determinant is associated. Object determinants could be set up on the basis of predefined string values embedded in the request. For example, when a request comes in with a certain USERID, the USERID can be linked to a group of objects in the cache memory that should be invalidated each time a post or other request comes from that certain USERID. Potential candidates for object determinants could also include using service identifiers of the server that originally served the object. The service identifier contains service IP address, TCP port and service identifier present in the HTTP request.

[00157] Another potential object determinant present in the request the request uniform resource locator ("URL"). In the case of caching of static objects, the request URL is typically sufficient to uniquely identify the object. For requests for dynamically generated content, however, the information present in the URL may not be sufficient to identify the cached object. The cache must therefore inspect other information in

the request to find object determinants including in HTTP headers, cookie header or in other custom HTTP headers. The cache can additionally look for a subset of relevant parameter information in a variety of other places in the client request, including, without limitation: in the URL query string, in the POST body, in a cookie header, or in any other request or response headers.

[00158] The problem in parsing a URL for object determinants is that the URL and other headers may contain a lot of information in addition to what is relevant for the cache's decision. The cache must therefore be able to parse through quite a lot of information to be able to identify the appropriate object determinants. In addition, the data in the header is often arbitrarily ordered, meaning there are no standardized ways that such data is placed into the HTTP header and therefore a simple comparison is often insufficient to locate the relevant object determinants in such string.

[00159] If there is no pre-configured policy to match a particular object determinant to a relevant object or group of objects stored in cache memory, the cache may still, in another embodiment, make such a determination. For example, the cache may examine and parse various aspects of the request to discover whether any other object determinants may be found in such request and used to link such request to particular objects stored in the cache memory that should be invalidated. Alternatively, one could also enable the cache to examine a request for certain object determinants that the cache determines, based on certain pre-defined heuristics, may meaningfully linked to particular objects or group of objects. For example, when the request comes into the cache for an update of a calendar associated with a particular USERID, an embodiment of the present invention could be set up to recognize that all cached objects with USERID equal to the USERID of the request updating the calendar, and that contains the user's calendar for any one particular day, will need to be invalidated.

[00160] The cache may also assume that the object determinants are present as a group of name=value or similar pairs in a non-specified order in the URL Stem, in the queries present in the URL, in the POST body or in a Cookie header. In an embodiment, it is assumed that the query is formatted as a list of name=value pairs. The user can therefore configure which parameter names are significant. Every cached object is keyed using first its access URL. The URL may look like /site/application/special/file.ext?p1=v1&p2=v2&p3=v3. The /site/application/special/file.ext part is the URL stem. The p1=v1&p2=v2&p3=v3

part is the URL query and contains parameter-value pairs. These parameter-value pairs may also be present in the POST body or in the Cookie headers.

[00161] In an embodiment, the user or administrator establishes that p1 and p2 shall be the invalidation parameters or object determinants. The cache will thereafter automatically group objects that have matching p1 and p2 values. One way of implementing this grouping is to map p1 and p2 to primary keys in database tables, i.e., to uniquely identifiable objects in the table that the cache will know how to reference in order to determine validation status. To update something in those database tables, in order to reflect the fact that data stored in the cache is no longer valid, the cache will specify new values for p1 and p2 and when the cache recognizes such new values the next time it goes to serve such content, it will know to invalidate the linked objects stored in its memory. The cache, when it encounters such a request, on seeing the update request knows that it has to invalidate the group with matching p1 and p2 values – because the cache understands that data in the origin will change, thereby affecting all objects that are related to those p1 and p2 object determinants.

[00162] To address the more complex case where the administrator has not pre-configured specific parameters embedded in the request as object determinants, the cache can deploy user-configured policies to extract the relevant object determinants from the request to assist in identifying when to invalidate groupings of objects. The determinant string is then used to locate the group of objects stored in the cache and invalidate such objects. These object determinants can be used to configure the cache to generate lists of significant parameter values. If an incoming write-request has matching values for the significant parameters then the objects tied to those parameter names should be invalidated. Alternatively, a user could specify the policy framework action that can extract the object determinant string from the request. The object determinant string is extracted from the write-request and all objects with matching determinant strings are invalidated. In this alternative approach, a request arrives at the cache, the cache makes a determination whether the request string matches an invalidation policy. The invalidation policy specifies objects in which content group should be invalidated.

[00163] Alternatively, the cache could use any other user information that may be present in the client request. As noted above, the authentication and authorization integration allows the cache access to the user information. The USERID or the GROUPID could be one of the determinants in the event the relevant grouping of

cached objects are linked to a user or a group of users. Although user information is often an important object determinant, the user information often may not be present in the HTTP request. In a further embodiment of the present invention, the dynamic caching aspects of the invention can be combined with another of Applicant's patent applications. To accomplish this, reference is made to Applicant's above referenced copending Patent Application No. 11/169,002 ("the Integrated Caching patent"). That application describes a system and method for integrating the cache with a variety of other networking elements including the ability to perform certain kinds of authentication, access control and audit (AAA) infrastructure. Thus, the level of security accorded to data that is generated by the applications is applied to data that is instead served from a cache. This technique allows the applications to cache sensitive, access controlled information that could not otherwise be cached.

[00164] This approach allows the cache to identify users that do not include identifiable user information in the HTTP request but that may be identifiable via the AAA approach described in the Integrated Caching patent. Such an approach enables the cache to identify the relevant user to a particular request through examining the authorization state information that can be shared from the AAA processing. In a further embodiment, the integration enables the application of security policies to information stored in the cache to prevent unauthorized users from accessing information stored at the cache.

[00165] This approach also address the challenge posed by the fact that a significant portion of dynamically generated data requires that the client requesting such data be authorized and authenticated before the cache can respond to the relevant request from the client. The cache must have the ability to authorize requests made by authenticated users so that applications can cache access-controlled objects and by integrating such dynamic caching technology with authentication and authorization information, this security can be achieved. The USERID or the GROUPID will be one of the object determinants if the objects are personalized to a user or a group of users. Thus, the level of security accorded to data that is generated by the applications is applied to cached information as well. This technique allows the applications to cache sensitive, access controlled information that could not otherwise be cached.

[00166] Finally, other information like time of day, state of the database at the origin, etc., may be parsed from the request and used as object determinants to determine

whether objects stored in the cache are still valid. The cache may take care of this situation by configuring appropriate expiration behavior in groupings of objects that are configured to be sensitive to such external variables.

[00167] To further address the challenge presented by the fact that requests for dynamic content must be parsed and interpreted by the cache, the cache in accordance with an embodiment of the present invention can limit which parameters are deemed to be relevant object determinants for the cache. In this way, the success rate for serving objects from the cache rather than forwarding such requests to the applicable application server can be enhanced. By way of example, a request query from a client may contain both a city and a state parameter. However, the cache may be configured to comply with the requirements of the application for which the cache is storing content to recognize that the response can be served to requests coming from clients that the query shows come from all clients in a given state without regard to the city value. For this purpose, the city parameter is not relevant and the cache could recognize this fact. An alternate embodiment involves configuring the cache so that a response can be served from the cache if just the city parameter makes a match regardless of what is specified for the state parameter.

[00168] In summary, the cache implements generalized parameterized object matching. In this approach, the cache is configured to recognize the subset of information in the request that will be useful as object determinants, and that are linked to a particular object so that when such object determinants are recognized, the cache can utilize the presence (or conversely the absence of such determinants) in evaluating whether the object or group of objects remains fresh and capable of being served from the cache. The cache maintains a table that it consults each time a request comes in to check against the configured parameters to determine if the requested data remains fresh, and which also allows the cache to match the relevant data to the proper object stored in the cache memory.

10. Incarnation Numbers

[00169] In yet another embodiment, the cache can utilize incarnation numbers to invalidate a group of objects. Where a cache needs to change the state of each of a group of objects at one time because of a change in the state at the origin, incarnation numbers provides a simple technique for effecting this invalidation. Whereas

identifying each object and changing the state individually is an inefficient approach to assuring freshness of data stored in a cache, use of incarnation numbers enables a much more simple and effective approach to invalidating groups of objects. The present embodiment describes how each object points to a data structure that represents the group and therefore the server need only send a command that changes the state in the data structure for the group. When a subsequent request for a cached object arrives from a client, the cache must first figure out whether the state has changed. To do so it looks up the data structure to reference whether the state has changed for the group.

[00170] In order to implement the data structure effectively, the cache must be able to determine whether to look up for a state change. Therefore, the cache must be able to determine whether it has already looked at the state change in the group or not. This is where the incarnation numbers are helpful. The cache associates dynamically generated objects into content groups. Each of these content groups may be represented through a hash table look-up process with a particular index value or "incarnation number" contained in a data structure. Thereafter, whenever the cache receives a client request that the cache recognizes as causing a state change, the client parses the client request for the relevant parameters, performs the hash look-up based on the recognized object determinants, and increments the index or incarnation number in the data structure. Each time an object stored within a designated grouping is requested by a client, the cache performs the hash algorithm on the object, and compares it to the original stored value in the data structure for such content group. If the stored value is the same as the number calculated by the cache for such object, then the cache knows the content remains fresh and can be served to the requestor. In the event the cache detects a discrepancy between the current incarnation number calculated for such object in and the number stored for such content group in the data structure, the cache knows that the stored object is no longer fresh. The cache then invalidates the stored object and sends the request along to the application server. When the response comes back the cache appliance will store the new response in the cache memory and link such response again to the new data structure. Thereafter, each time the cache receives a request for an object in that grouping, the cache can make the comparison and assuming no further changes have been made to the data structure, the cache can serve the newly stored object.

By utilizing invalidation of a group of objects in this fashion, the cache is able to invalidate very quickly – and the time taken is constant regardless of the number of objects invalidated. Through this faster and more efficient process of invalidation, the techniques of the present invention enables the cache to more effectively handle dynamically generated objects. The approach allows cache appliances that sit in front of applications to more aggressively store and serve dynamically generated objects without serving invalid or stale content because of rapid changes in such data. The embodiment enables the cache to serve data that frequently or unpredictably changes thereby improving the performance of the cache. The cache is also able to invalidate objects and group of objects stored in the cache memory using user commands and also by examining and grouping various kinds of web traffic.

11. Flash Cache and Flash Crowds.

[00172] Another embodiment of the present invention may also include a technique that is able to increase cache hit rates for extremely fast changing dynamically generated objects that would not otherwise be cacheable. When a cache manager 232 receives a first request from a client for a particular object, the cache manager 232 forwards that request to the originating server 106a-106n for processing because, as the first such request, the cache manager 232 cannot yet have such object stored. As the response object is generated and then returned to the requesting client via the cache manager 232, the cache manager 232 automatically stores a copy of the object. With objects that change extremely quickly, traditional caches, in contrast to the present invention, just pass all responses along to the originating server 106a-106n for processing as such content is always assumed to be not fresh, and therefore not valid for cache storage.

[00173] In one aspect, the present invention is directed towards a “flash cache” technique for handling additional requests for an object received by the cache manager 232 during the time the cache manager 232 or appliance 104 is in the processing of transmitting or waiting to transmit a response for a first requestor of the object. Referring now to FIG. 5, method 500 of the present invention depicts the flash cache technique of the present invention. In brief overview and in view of FIGs. 1 and 2, at step 510, the cache manager 232 of the present invention receives a response from an originating server 106a-106n for an object requested by a first

client, for example client 102a illustrated in FIG. 1. The object may comprise a dynamically generated object which is generated by one of the originating servers 106a-106n, and the response from the originating server 106a-106n may include the dynamically generated object. At step 515, the cache manager 232 requests transmission of the response to the first client, such as via a network stack, for example, a TCP/IP stack of the device 104. At step 520, the response may be stored or held in a buffer while waiting to be transmitted. For example, the device 104 may throttle communications to a slow connected or low bandwidth client 102a, and therefore, the device 104 queues network packets, which represent the response.

[00174] At step 525, while the response of the first client 102a is waiting to be transmitted in the buffer or otherwise in the process of being transmitted and/or prior to completing transmission of the response to the first client 102a, the cache manager 232 receives a second request for the object from a second client 102B. At step 530, the cache manager 232 determines the object is currently in the buffer and provides the object from the buffer to response to the second request to the second client 102B. In one embodiment, as the device uses a single TCP/IP stack, the same object can be provided for transmission to the first client 102a and the second client 102B. At step 535, the first response is transmitted to the first client 102a, and a response is transmitted to the second client 102B. At step 540, the response of the first client 102B is removed from the buffer.

[00175] In further detail, at step 510, the device 104, such as by the cache manager 232, intercepts or otherwise receives a response from an originating server 106a-106n to a request from a first client 102a for an object, such as dynamically generated object. In one embodiment, the dynamically generated object is identified as non-cacheable, such as by the originating server 106a-106n, or is otherwise not identified as cacheable. In some embodiments, the cache manager 232 receives the request from the first client 102a and forwards the request to the originating server 106a-106n. In one embodiment, the cache manager 232 checks the cache manager 232 for the requested object prior to forwarding the request to the originating server 106a-106n. In some cases, the cache manager 232 determines the object is invalid or has expired. In other cases, the cache manager 232 determines the object is not stored or otherwise available in the cache manager 232. For example, the object may have been flushed for example, or this may be the first time the object has been requested from an originating server 106a-106n. In some embodiments, the response comprises

the object, and in other embodiments, the response indicates a status, such as a failure or error message regarding the object.

[00176] At step 515, the response for the request by the first client 102a is provided for transmission to the first client 102a. In some embodiments, the device 104 receives the response, and requests the packet processing engine 240 to transmit the response to the client 102a. In other embodiments, the cache manager 232 receives the response and requests the packet processing engine 240 to transmit the response to the client 102a. In one embodiment, the device 104 comprises one network stack for receiving and transmitting network packets to the clients 102a-102n. In some embodiments, the network stacks comprises a TCP/IP stack. In other embodiments, the device 104 may include multiple network stacks, each associated with or used by one or more clients 102a-102n. The network stacks or stacks are associated with the kernel 204, network ports 226 and/or the packet processing engine 240 as those ordinarily skilled in the art will recognize and appreciate.

[00177] At step 520, during the course of transmitting the response to the first client 102a requested at step 515, the device 104 queues, holds, or otherwise stores the response in a buffer. In some embodiments, the buffer is a part of the network stack, such as a TCP/IP stack. For example, the buffer may comprise a data structure used in or by the TCP/IP stack or by a network driver, filter or other network related software processing or handling any portion of the network stack. In other embodiments, the buffer is part of the network port 226. In additional embodiments, the buffer is part of the packet processing engine 240, or yet in further embodiments, the kernel comprises the buffer. The buffer may be any memory or storage element located in any portion of the device 105. In many embodiments, the response is queued, held or stored in the buffer as one or more network packets, such as TCP/IP protocol based packets. In one embodiment, the response is held in one or more data structures providing the buffer for use by the present invention. As those ordinarily skilled in the art will recognize and appreciate, the response may be queued, held or stored in the buffer in any suitable form.

[00178] Furthermore, the response may be stored or queued in the buffer for any duration of time period, either arbitrary, predetermined, implicit, explicit or otherwise, and using any suitable means and/or mechanisms. In one embodiment, the response is stored in the buffer as the network packet awaits to be transmitted by the device 104 based on a rate of transmission, packet queuing rate, network traffic and congestion,

or any other network characteristics and as may be determined by the device 104, such as via the packet processing engine 240 and/or the network ports 226. In another embodiment, the response is held in a buffer for a determined time period. In some embodiments, the device 104 queues, manages and transmits network packets to a client 102a based on characteristics, such as bandwidth, type of connection, etc. of the client's network connection(s), or characteristics of the client 102a. For example, for a client 102a with a slower connection to the device 104 than the device's connection to an originating server 106a-106n, the device 104 may throttle or otherwise manage the transmission of the network packets to the client 102a to provide a desired performance of or behavior for the client 102a, or the network connection of the client 102a. As such, the device 104 may queue or hold the response in a buffer for a desired time period in accordance with any queuing or buffer management logic, function, rule or operation.

[00179] At step 525, while the response of the first client 102a is held in the buffer, or is being transmitted or otherwise prior to completing transmission of the response to the first client 102a, such as a buffer of a TCP/IP stack, the device 104, such as via the cache manager 232, intercepts or otherwise receives a second request from a second client 102B for the object obtained by the cache manager 232 for the first request. That is, in one embodiment, the object requested by the second client 102B is stored in the buffer waiting to be transmitted. In some embodiments, the device 104 receives multiple requests from multiple clients 102b-102n for the same object requested by the first client 102a and is currently being held in the buffer for transmission to the first client 102a or is other currently being transmitted to the first client 102a. In some embodiments, the cache manager 232 determines the second request is requesting the same object as the first request using any suitable means and/or mechanisms, such as any of the techniques described herein for example, by using object determinants. In some embodiments, the appliance 104 or the cache manager 232 queues the second request, and any additional requests in a queue using any type and form of queuing mechanism. As such, in one embodiment, the clients 102a-102n do not need to resubmit the requests as the appliance 104 or the cache manager 232 queues and responds to the requests on behalf of the servers 106a-106n. In another embodiment, the appliance 104 or cache manager 232 queues and responds the requests transparently to the client 102a-102n without submitting the request to the server 106a-106n. The appliance 104 or cache manager 232 may use any

mechanism such as a queue, object, or data structure in memory to queue the requests from the client 102a-102n to be responded to using the techniques of the present invention.

[00180] At step 530, the device 104 determines the object requested by the second client 102b is in the response stored in the buffer or is in the response currently being transmitted to the first client 102a. In some embodiments, the cache manager 232 determines the object is stored in the buffer. In some cases, the cache manager 232 first checks cache memory or storage for the object, and then checks the buffer for the object. In other cases, the buffer may be considered a storage location for cached objects to be searched or managed by the cache manager 232. In other embodiments, the cache manager 232 requests the packet processing engine 240 to check if the object is in the buffer. In additional embodiments, the cache manager 232 determines if the object is in the buffer via any application programming interface (API) of the kernel 204. In some cases, the cache manager 232 interfaces to any drivers, network handlers, filters, or other software handling or processing any portion of the network stack to determine if the object is in the buffer.

[00181] Further at step 530, the present invention provides the object stored in the buffer or being transmitted as the response to the first client 102a for use in the response to the second request from the second client 102B at step 525. In some embodiments, the device 104, such as via the packet processing engine 104, forms a response to the second request by using any portion of the buffer and related data structures of the network stack. In one embodiment, the object is stored in the buffer once, such as in a single network stack configuration, and used to form the network packets to be transmitted to the first client 102a and second client 102B respectively. In other embodiments, a copy of the object from the buffer is used to form the response to the second request. In further embodiments, the object is provided from a first network stack to a second network stack to form the response to the second request for the second client 102B. The response of the first client 102a may be maintained in the buffer, and any portion or all of the response in the buffer may be used to provide the response for the second client 102B. In one embodiment, the response of the first client stored in the buffer may be used without modification to respond to the second client 102B. In another embodiment, the data representing the object stored in the response of the first client 102a in the buffer is modified to provide the response for the second client 102B.

At step 535 of method 500 of the present invention, the first response to the first client 102a is transmitted from the device 104, and the second response in response to the second request of the second client 102B is transmitted from the device 104. The response to the first client 102a and second client 102B may be transmitted in any order from the device 104, with any time interval between each other, using the same network stack 267 or different network stacks. In some embodiments, the response for the second client 102B is also queued, held or stored in a buffer waiting to be transmitted. As such, any additional requests for the same object may also be provided via the first response or second response stored in the buffer and not yet transmitted, flushed, removed or otherwise unavailable.

[00183] At step 540, the response of the first client 102a is flushed or otherwise removed from the buffer. In some embodiments, if the response of the second client 102B was stored to the buffer, the response of the second client is also removed. In other embodiments, the response of the second client may remain in the buffer and steps 525 to 535 of method 500 are practiced again using the second client response while receiving a third request for the object from a third client.

[00184] In another aspect, the present invention is directed towards a “flash crowd” technique for handling situations where the appliance 104 or cache manager 232 receives additional requests, e.g., nearly simultaneous requests, for the same object during the time the server is processing and returning the response object for a first requestor. Once all such nearly simultaneous requests are responded to by the cache, the object is immediately flushed from the cache memory, with no additional expiry time or invalidation action by application or administrator. This technique of the present invention enables data to be cached and served for very small amounts of time for objects that would otherwise be considered non-cacheable. This approach yields a significant improvement in applications that serve fast changing data to a large volume of concurrent users, such, for example, as real time stock quotes, or a fast evolving news story.

[00185] FIG. 6 describes a method 600 depicting the steps taken by the cache of the present invention to serve the same content to each of the client requestors as it receives such content in response to the first request. In brief overview of method 600, at step 610, the cache manager 232 or appliance 104 receives a first request from a first client 102a for an object, such as a dynamically generated object, from an originating server 106a-106n. At step 615, the cache forwards the first request to the

originating server 106a-106n. While waiting for a response to the first request or prior to responding to the first request of the first client 102a, at step 620, the cache receives a second request from a second client 102b for the object. In some embodiments, the cache does not forward the second request to the originating server but instead queues the second request. At step 625, the cache receives the response to the first request from the originating server, and at step 630, transmits the received response and/or object from the response to the first client 102a in response to the first request, and to the second client 102b in response to the second request. At step 635, the cache may receive an additional third request from a third client 102c for the object prior to completing transmission of the received response and/or object from the response to either the first client 102a or second client 102b. At step 640, the cache also transmits the received response and/or object from the response to the third client 102c. Upon responding to the requests, the cache, at step 645, flushes the object from the cache memory, for example, without any expiry time or invalidation action by the application, originating server, or administrator.

[00186] In further details and in view of FIGs.1 and 2, step 610 of method 600 of the present invention, the cache manager 232 of appliance 104, in one embodiment, may intercept or otherwise receive a request from a client 102a-102n, for example client 102a in FIG. 1 for an object, such as a dynamically generated object, at any point in time during communications between the client 102a and the originating server, for example 106a. In one embodiment, the request from the first client 102a may be the first time the first client 102a has requested the object. In one example, the client 102a may request the object for the first time upon connection to the origination server 106a. In another example, the cache manager 232 may have not previously cached or requested the object from the originating server 106a on behalf of the client 102a or any other client 102a-102n. In other embodiments, the request from the client 102a may be a subsequent request for a previously requested object. In yet other embodiments, the cache manager 232 may have previously cached the object requested by the client 102a at step 610. As such, the techniques of the present invention as illustrated in FIG. 6 may be practiced at any point during the course of communications between the clients 102a-102n and the originating servers 106a-106B, regardless if the object is currently cached, has been previously cached or has never been cached.

At step 615, the cache manager 232 forwards the request from the client 102a to the originating server 106a. In one embodiment, the cache manager 232 forwards the request to the originating server without checking if the object is stored in the cache manager 232. In some embodiments, the cache manager 232 first checks if the object requested by the client 102a is available from the cache manager 232, and if not, then forwards the request to the originating server 106a. In other embodiments, the cache manager 232 checks if the object is stored in the cache but is either invalid or otherwise needs to be updated, or is otherwise about to be invalid or shortly need to be updated. In another embodiment, the cache manager 232 determines a dynamically generated object stored in the cache should be dynamically regenerated and served from the originating server 106a-160N. As such, if the cache manager 232 determines the object should be requested from an originating server 106a-106n, the cache manager 232 forwards the request from the client 102a to the originating server 106a-106n. In one embodiment, the cache manager 232 forwards the same or original request received from the first client 102a to the originating server 106a. In another embodiment, the cache manager 232 sends a request formed by the cache manager 232 for the object requested or identified in the request by the first client 102a.

[00188] While waiting for a response to the request of the first client 102a from the originating server 106a-160N and/or prior to responding to the request of the first client 102a, the cache manager 232, at step 620, may intercept or otherwise receive one or more additional requests for the object identified, associated with or requested by the first request of the first client 102a at step 610. For example, in one embodiment, the cache manager 232 receives a second request from a second client 102B for the object, which may be dynamically generated by the originating server 106a in response to the first request. In one embodiment, the appliance 104 or the cache manager 232 queues the second request, and any additional requests in a queue using any type and form of queuing mechanism. In some embodiments, the second and/or additional requests for the object may occur or be received by the cache manager 232 nearly simultaneously. In other embodiments, the second and/or additional requests may be intercepted or received by the cache manager 232 at any point between receiving the first request and receiving a response to the first request from the originating server 106a. In yet another embodiment, the second and/or additional requests may be received by the cache manager 232 at any point between

receiving the first request and providing a response to the first request to the client 102a.

[00189] In one embodiment, prior to responding to the first request, the cache manager 232 does not forward the second request from the second client 102b for the object and/or any other additional requests for the objects from any other clients 102c-102n. In one embodiment, the cache manager 232, at step 625, queues or otherwise holds the second and additional requests, for example, at the appliance 104, until receiving the response to the first request. As such, the cache manager 232 transparently to the second client 102B does not forward the second client's request. However, in some embodiments, it appears to the second client 102B that the request is being processed, for example, by an originating server 106a-106n

[00190] In some embodiments, the second request from the second client 102a may request

the same object as the first request at step 610 but also include an additional request for another object or other information, static, dynamic or otherwise. For example, the request of the second client 102B may include multiple requests. In these embodiments, the cache manager 232 may identify the portion of the request related to the object already requested in connection with the first request of the first client 102a, and not process the identified portion until the response to the first request is received. For the other portions of the request from the second client 102B, the cache manager 232 may process these other requests in any suitable manner, such as by any techniques of the present invention described herein. For example, the second request may include a request for a second object, for which the cache manager 232 requests the second object from an originating server 102a-20B or obtains the second object from the cache manager 232.

[00191] At step 625, the method 600 of the present invention receives a response to the first request forwarded to the originating server 106a at step 615. In some embodiments, the response includes the object, such as an object dynamically generated by any of the servers 106a-106n. In other embodiments, the response includes a pointer to a location of the object to be fetched or otherwise obtained by the cache manager 232. In further embodiments, the response may indicate some error message or otherwise indicate the object requested cannot be provided. In one embodiment, the cache manager 232 compares the received response to the corresponding request to

determine if the response is suitable or appropriate for the request. In further embodiments, the cache manager 232 obtains contents of the response, such as the dynamically generated object, and creates, modifies, or otherwise provides the response desired to be communicated to the first client 102a.

[00192] At step 630, the cache manager 232 of the present invention uses the response and/or object from the response received for the first request to respond to the first request and the second or additional requests for the same object. In one embodiment, the cache manager 232 transmits the received response to the first client 102a in response to the first request, and transmits the received response to the second client 102B in response to the second request. In another embodiment, the cache manager 232 transmits the received response to the first client 102a, but modifies the received response and transmits the modified response to the second client 102B in response to the second request. For example, the cache manager 232 may modify the response to the second client 102B to include session information or other information corresponding to the second request. In further embodiments, the cache manager 232 may obtain the object requested from the received response, and create, modify or otherwise provide a response corresponding to the first request and the second request and including the object, such as a first response and a second response, and transmit the corresponding response to the first client 102a and second client 102B. As such, the cache manager 232 may use all or any portion of the response received for the first request from an originating server 106a-106n in responding to the requests from the clients 102a-102b for the object.

[00193] In some embodiments, the cache manager 232 may intercept or otherwise receive, at step 635, additional requests for the object from other clients 102c-102n after receiving the response to the first request from the server 106a and before completing transmission of a response to the first client 102a, second client 102b or any other client 102a-102c requesting the object at step 625. For example, the cache manager 232 may receive a third request for the object from a third client 102C. In these embodiments, the cache manager 232 may use all or any portion of the response received for the first request from an originating server 106a-106n and communicated to clients 102a and/or 102b at step 630 to respond to these additional requests. In another embodiment, the cache manager 232 may have completed transmission of responses to the first client 102a and second client 102b but may still have any of the responses in memory or otherwise available to respond to the third request of the third

client 102c with the object provided via the first request. As such, at step 640, the cache manager 232 of the present invention can respond to the third request or further additional requests without requesting the object again from the originating server and/or without storing the requested object in cache.

[00194] In some embodiments, the appliance 104 or the cache manager 232 queues the second request, third request or any additional requests in a queue. As such, in one embodiment, the clients 102a-102n do not need to resubmit the requests as the appliance 104 or the cache manager 232 queues and responds to the requests on behalf of the servers 106a-106n. In another embodiment, the appliance 104 or cache manager 232 queues and responds the requests transparently to the client 102a-102n without submitting the request to the server 106a-106n. The appliance 104 or cache manager 232 may use any mechanism such as a queue, object, or data structure in memory to queue the requests from the client 102a-102n to be responded to using the techniques of the present invention. Additionally, the appliance 102 or cache manager 232 may respond to the client requests in any order, and not necessarily in the order received or in a first-in first-out manner from the queue.

[00195] At step 645, the cache manager 232 flushes the received response to the first request from the originating server 106a, any responses to the client's requests, and/or the object requested by the requests from the cache manager 232. In one embodiment, the cache manager 232 immediately flushes the response and/or object from cache upon completing the transmission of the last response, such as at step 630 in one embodiment or step 640, in another embodiment. In some embodiments, the cache manager 232 maintains the object for a predetermined time period, such as in accordance with an expiry time. In other embodiments, the cache manager 232 maintains the object in cache until the originating server 106a-160N invalidates the object. In yet another embodiment, the object remains in memory or storage until the memory or storage is written over or otherwise used by the cache manager 232 for other purposes.

[00196] Although the techniques of methods 500 and 600 of the present invention are generally described above in regards to a request for a single object from multiple clients, those ordinarily skilled in the art will recognize and appreciate that the flash cache and flash crowd handling techniques of the present invention may be used to handle multiple object requests either in a single request or a series of requests from multiple clients that may occur either sequentially, concurrently, or nearly

concurrently or simultaneously. As such, multiple instances of practicing methods 500 and 600 may be executing at various frequencies and at various times during the course of operation of the cache manager 232 or appliance 104 in a network environment 200. Furthermore, each request or series of requests from each client may request different sets of objects, with a common set of one or more objects requested among one or more clients.

[00197] For example in view of the flash crowd techniques of the present invention, a first client may request a first object and a second object, while a second client may request the second object and a third object, and the third client may request the first object and the third object. In some embodiments, the cache of the present invention uses multiplexing and/or demultiplexing techniques to request an object once for each object of multiple requests outstanding between multiple clients, and provide the corresponding response to each client accordingly. Further to the above example, the cache may obtain the first object from the originating server once for the first client and the third client, the second object once for the first client and second client, and the third object once for the second client and third client. In some embodiments, when all the objects for one or more requests for a client are received by the cache, the cache transmits the corresponding response or responses to the client. In other embodiments, the cache may respond to the client with each object upon receipt of the object by the cache.

[00198] Likewise, in view of the flash cache technique of method 500 of the present invention, the buffer of the appliance may comprise multiple objects waiting to be transmitted while additional requests for those objects are received by the appliance. In one embodiment, the appliance of the present invention uses a single TCP/IP stack from which multiple buffers are associated with and temporarily store multiple objects during the course of transmission. As such, a first object in the buffer may be used to respond to requests from a first client and a second client for the first object, and a second object in the buffer may be used to respond to requests from a third client and fourth client for the second object, or to respond to a request from either or both of the first client or second client for the second object.

[00199] Furthermore, the flash cache and flash crowd techniques of the present invention may be practiced in conjunction or in any combination with each other as those ordinarily skilled in the art will recognize and appreciate. For example, the appliance of the present invention may be providing responses to additional requests

for an object from the buffer via the flash cache technique, but determines the object in the buffer is either invalid or otherwise needs to be requested from the originating server. At that point, the appliance may switch to the flash crowd technique by requesting the object from the originating server. The appliance then uses the object requested from the originating server to respond to additional requests for the object received and queued at the appliance while waiting for the originating server to respond.

12. Etag Insertion

[00200] In another aspect, the present invention is directed towards using techniques to increase cache hit rates by inserting information, such as entity tag and cache control information for an object, in a response to a client to enable the cache to check for a hit in a subsequent request. The present invention provides a technique for identifying a dynamically generated object as cacheable to a client when the originating server did not identify the object as cacheable. In some embodiments, such as an embodiment handling HTTP requests and responses for objects, the techniques of the present invention insert an entity tag, or “etag” into the response to provide cache control for objects provided without entity tags and/or cache control information from an originating server. In brief reference to the HTTP protocol as known to those ordinarily skilled in the art, an “etag” is the ETag header field of an HTTP based message which provides the current value of the entity tag for a requested variant, such as an object, and may be used for comparison with other entities from the same resource. The etag provides a mechanism for validation in the HTTP protocol. Also, an HTTP based message may include a Cache-Control header field to specify directives to caching mechanisms along the request/response message or communication chain. The cache-control directives specify desired caching behavior.

[00201] Referring now to FIGs. 7A and 7B and in view of FIGS. 1 and 2, steps are depicted for an embodiment of practicing the “etag” techniques of the present invention. In brief overview of FIG. 7A, method 700 of the present invention, step 710, receives a response having an object requested by a client 102a-102n, such as object served from an originating server 106a-106n and received by a cache manager 232 of the appliance 104. At step 715, the cache manager 232 determines if the object from the response is under tag entity control or otherwise has an entity tag. If the

object does not have an entity tag, then at step 720, the cache manager 232 generates an entity tag for the object. At step 725, the cache manager 232 modifies the response to provide entity control information including the cache generated entity tag. At step 730, the cache manager 232 stores the object with the entity tag information to the cache manager 232. At step 735, the cache manager 232 responds to the client 102a-102n requesting the object with the modified response. As such, the client 102a-102n receives an entity tag controlled object from the cache manager 232 in a manner transparent to the client 102a-102n, such that the client 102a-102n is not aware that the originating server 106a-106n did not provide the object with the entity tag information.

[00202] In further detail, at step 710 of the present invention, the cache manager 232 receives a response to a request of a client 102a-102n for the object from an originating server 106a-106n. In some embodiments, the response may be received from a request sent from the client 102a-102n to the server 106a-106n, or from the cache manager 232 on behalf of the client 102a-102n to the server 106a-106n. In one embodiment, the cache manager 232 forwards the request from the client 102a-102n to the server 106a-106n after checking the cache manager 232 for a matching and/or valid object. In another embodiment, the cache manager 232 requests the object from the server 106a-106n during practicing the flashing crowd technique of the present invention described in conjunction with FIG. 6.

[00203] At step 715 of method 700, the cache manager 232 determines if the object received in the response at step 710 is under entity tag control or has an entity tag. In some embodiments, the cache manager 232 inspects, examines, checks or otherwise reads any field, header, data, or other information of the one or more packets providing the response. As those ordinarily skilled in the art will recognize and appreciate, the portion of the packet or packets to inspect or read depends on the type of one or more protocols used for the response. In one embodiment, the cache manager 232 determines the entity tag portion of the response packet is missing, empty or otherwise not provided. In another embodiment, the cache manager 232 determines the entity tag portion of the response packet is corrupted, invalid or otherwise not useable by the cache manager 232 and/or the client 102a-102B. In some embodiments, the cache manager 232 determines the response has an entity tag or is under entity control but not in the type, form or manner desired by the cache manager 232. In addition to or instead of the entity tag, the cache manager 232, in a

similar manner to the entity tag, determines if there exists desired cache-control information in the response. In further embodiments, the cache manager 232 may determine via the policy engine 236 if the object should be under entity tag and/or cache control. As such, in some embodiments, the cache manager 232 may not continue performing the etag techniques of the present invention, such as steps 720 to 725, if the appliance 104 determines via the policy engine 236 the object should not be under entity tag or cache control.

[00204] At step 720, the cache manager 232 generates desired entity tag and cache control information for the object of the received response. The etag may comprise any type and/or form of entity tag representation and may comprise any numeric, alphanumeric, letters in one or more characters. In the embodiment of HTTP, the etag may comprise a string or quoted string, and may further include a prefix for indicating the etag is weak or strong in accordance with the HTTP protocol. In some embodiments, the cache manager 232 uses an entity header, or etag, counter which increments an etag counter sequentially, or in any desired increments. In some embodiments, the etag counter is global to all objects for which the cache manager 232 generates an etag. In other embodiments, the cache manager 232 may have multiple etag counters, each associated with or specific to an object or groups of objects.

[00205] In one embodiment, the etag may comprise a universal identifier (UID) or global universal identifier (GUID) associated with the object. In another embodiment, the etag may comprise an incarnation number as previously described above. In some embodiments, the cache manager 232 obtains the etag from an application, program, or other form of executable instructions running on the appliance 104 or in other embodiments, from another system or device accessible by the appliance 104 via the network. In further embodiments, the cache manager 232 may use any suitable algorithm, business rule or logic to generate a desired unique etag for the object. In one embodiment, the appliance 104 comprises a database, file, or other organized storage element for generating and maintaining etags for objects for the cache manager 232.

[00206] At step 725 of method 700, the cache manager 232 modifies the response from the originating server 106a-106n to include desired entity tag and/or cache control information, such as the entity tag generated by the cache at step 720. In one embodiment, the cache manager 232 inserts the entity tag information via an entity

tag header of the protocol, such as the HTTP protocol. In another embodiment, the cache manager 232 modifies any entity tag header fields in the response to include the desired entity tag. In yet another embodiment, the cache manager 232 replaces the entity tag fields in the response with the entity tag fields desired to be used by the cache manager 232. In further embodiments, the entity tag information is inserted, modified or placed into any portion of the network packet or packets for communicating the response to the client 102a-102n.

[00207] Additionally, step 725, the cache manager 232 may insert, modify or otherwise place any desired cache-control information into any header, fields of a protocol, or any other portion of network packets for communicating the response. The cache-control information may be generated by the cache manager 232 by any suitable means and/or mechanisms. In one embodiment, the cache-control information is configured into and obtained via the policy engine 236 of the device 104. In another embodiment, the cache manager 232 determines the desired cache-control information from any algorithms, business rules, or logic of the cache manager 232. In some embodiments, the cache-control information may be based on any history of the network and caching performance between the client 102a-102n and the appliance 104 or cache manager 232.

[00208] At step 730 of the method 700 of the present invention, the cache manager 232 stores the generated entity tag with the object in any suitable or desired manner in the cache manager 232. In some embodiments, the entity tag is included in or is made a part of the stored object. In other embodiments, the entity tag is stored separately but associated with the object. In another embodiment, the entity tag and object are associated via a database, linked list, lookup table, or any other mechanism for associating one entity with another entity as known to those skilled in the art. Additionally, the cache-control information provided with the object and/or entity tag, in some embodiments, may also be stored in association with the object in the cache manager 232.

[00209] At step 735, the cache manager 232 causes the appliance 104 to transmit the modified response to the client 106a-106n. For example, the cache manager 232 requests the packet processing engine 240 to transmit the modified response to the client 102a-102n. As such, the client 102a-102n receives the object with entity tag and cache control information, even though unbeknownst to the client 102a-102n the originating server 106a-106n did not generate such information in one embodiment,

or the cache manager 232 changed such information generated by the server 106a-106n. The etag technique of the present invention can ensure that every object has a validator, i.e., an etag, or has an etag controlled or otherwise desired by the cache manager 232. If an object doesn't have a validator or a desired validator, then the cache manager 232 inserts its own etag. These etag validators can avoid serving full responses on repeat requests for an object as will be discussed in connection with FIG. 7B.

[00210] FIG. 7B depicts a method 750 of the present invention for using the entity tag to check by a client 102a-102n if an object has been modified. As such, the client 102a-102n can cache the object on the client 102a-102n and check with the cache manager 232 or originating server 106a-106n if the object has changed. In brief overview of method 750, at step 775, the cache manager 232 receives a request from the client 102a-102n to check the entity tag of an object. At step 760, the cache compares the entity tag received from the request of the client 102a-102n with the entity tag of object stored in cache. At step 765, the cache manager 232 determines if the entity tag from the client matched the entity of the object in cache. If the entity tags do not match, then at step 770, the cache manager 232 responds to the client 102a-102n with a response providing the current entity tag for the object. If the entity tags do match, then at step 775, the cache manager 232 responds to the client 102a-102n with a response indicating the object has not been modified.

[00211] In further detail, at step 755, the cache manager 232 may receive a request from the client 102a-102n to check or validate the entity tag of an object by any type and/or form of communication. In one embodiment, the request comprises an HTTP message with an If-Match or If-None-Match request header field, which includes an entity tag string value provided by the client 102a-102n. For example, the client 102a-102n may have a copy of the object in a client cache or otherwise available on the client 102a-102n. The client 102a-102n uses the entity tag of the client's copy of the object in the request to determine or validate if the object is current.

[00212] At step 760, the cache manager 232 compares the entity tag value received from the client 102a-102B with the entity tag for the object to determine if the cache manager 232 has a different version of the object than the client 102a-102n. The cache manager 232 may use any suitable means and/or mechanisms to compare the entity tag value from the client 102a-102n with the entity tag value for the cached object. In one embodiment, the cache manager 232 performs a byte-compare on the

etag value from the request with the etag stored with the object in the cache. In another embodiment, the cache manager 232 performs and uses hash algorithms on the entity tag and/or the cache object to determine if the object has been modified. In further embodiments, the cache manager 232 requests the object from the server 106a-106n, and compared the received object with the cached object to determine if the object has been modified.

[00213] If, at step 765, the entity tag from the client 102a-102n does not match the entity tag stored with the object in the cache manager 232 then at step 770, the cache manager 232 sends a response to the client indicating the entity tags do not match or otherwise the object has been modified. In one embodiment, the cache manager 232 sends a response to the client 102a-102n providing a new entity tag value for the object. In another embodiment, the cache manager 232 sends a response to the client 102a-102n provide a new version of the object, or the modified object, such as a modified object with a new entity tag stored in the cache manager 232. If, at step 765, the entity tag from the client 102a-102n does match the entity tag stored with the object in the cache manager 232 then at step 775, the cache manager 232 sends a response to the client indicating the entity tags do match or otherwise the object has not been modified.

[00214] Although the etag techniques of the present invention has generally been described in reference to the HTTP protocol, those ordinarily skilled in the art will recognize and appreciate the present invention may be practiced with any type and/or form of protocol to provide entity tag and cache information. In some embodiments, the protocol may not support entity tag and cache information fields, while in other embodiments, the protocol may support entity tag and cache information fields. In any of these embodiments, the techniques of the present invention provide entity tag and cache control information when the server does not provide this information in a response or when the cache decides to change or control the entity and cache control information provided by the server in a manner controlled or desired by the appliance 104.

[00215] In order to address the challenge to dynamically caching content dynamically generated by origin servers and applications the cache appliance of the present invention, can incorporate pre-packaged policies for certain commonly used applications such as Outlook® Web Access, Oracle® Configurator, and Serena® TeamTrack®. These packages are designed with an understanding of the way such

applications process data thereby enabling the cache to make more intelligent decisions about which object and how long to keep objects generated by such applications stored in the cache memory.

D. Exemplary Computer System-Based Implementation

[00216] The functions of the present invention may be implemented using hardware, software, or a combination thereof and may be implemented in one or more computing devices or other processing systems. For example, FIGS. 8A AND 8B depict an example computing device 800 that may be utilized to implement any of the techniques, methods and functions of the present invention. As such, the appliance 104 of the present invention may be a computing device 800 or a specific purpose device designed and constructed to provide any of the techniques, functions and operations of the present invention described herein.

[00217] FIGs. 8A and 8B depict block diagrams of a computing device 800, and in some embodiments, also referred to as a network device 800, useful for practicing an embodiment of the present invention. As shown in FIGs. 8A and 8B, each computing device 800 includes a central processing unit 802, and a main memory unit 822. As shown in FIG. 8A, a typical computing device 800 may include a visual display device 824, a keyboard 826 and/or a pointing device 827, such as a mouse. Each computing device 800 may also include additional optional elements, such as one or more input/output devices 830a-830b (generally referred to using reference numeral 830), and a cache memory 840 in communication with the central processing unit 802.

[00218] The central processing unit 802 is any logic circuitry that responds to and processes instructions fetched from the main memory unit 822. In many embodiments, the central processing unit is provided by a microprocessor unit, such as: those manufactured by Intel Corporation of Mountain View, California; those manufactured by Motorola Corporation of Schaumburg, Illinois; those manufactured by Transmeta Corporation of Santa Clara, California; those manufactured by International Business Machines of White Plains, New York; or those manufactured by Advanced Micro Devices of Sunnyvale, California. The computing device 800 may be based on any of these processors, or any other processor capable of operating as described herein.

[00219] Main memory unit 822 may be one or more memory chips capable of storing data and allowing any storage location to be directly accessed by the microprocessor 802, such as Static random access memory (SRAM), Burst SRAM or SynchBurst SRAM (BSRAM), Dynamic random access memory (DRAM), Fast Page Mode DRAM (FPM DRAM), Enhanced DRAM (EDRAM), Extended Data Output RAM (EDO RAM), Extended Data Output DRAM (EDO DRAM), Burst Extended Data Output DRAM (BEDO DRAM), Enhanced DRAM (EDRAM), synchronous DRAM (SDRAM), JEDEC SRAM, PC100 SDRAM, Double Data Rate SDRAM (DDR SDRAM), Enhanced SDRAM (ESDRAM), SyncLink DRAM (SLDRAM), Direct Rambus DRAM (DRDRAM), or Ferroelectric RAM (FRAM). The main memory 822 may be based on any of the above described memory chips, or any other available memory chips capable of operating as described herein. In the embodiment shown in FIG. 8A, the processor 802 communicates with main memory 822 via a system bus 850 (described in more detail below). FIG. 8A depicts an embodiment of a computing device 800 in which the processor communicates directly with main memory 822 via a memory port 803. For example, in FIG. 8B the main memory 822 may be DRDRAM.

[00220] FIG. 8B depicts an embodiment in which the main processor 802 communicates directly with cache memory 840 via a secondary bus, sometimes referred to as a backside bus. In other embodiments, the main processor 802 communicates with cache memory 840 using the system bus 850. Cache memory 840 typically has a faster response time than main memory 822 and is typically provided by SRAM, BSRAM, or EDRAM.

[00221] In the embodiment shown in FIG. 8A, the processor 802 communicates with various I/O devices 830 via a local system bus 850. Various busses may be used to connect the central processing unit 802 to any of the I/O devices 830, including a VESA VL bus, an ISA bus, an EISA bus, a MicroChannel Architecture (MCA) bus, a PCI bus, a PCI-X bus, a PCI-Express bus, or a NuBus. For embodiments in which the I/O device is a video display 824, the processor 802 may use an Advanced Graphics Port (AGP) to communicate with the display 824. FIG. 8B depicts an embodiment of a computer 800 in which the main processor 802 communicates directly with I/O device 830b via HyperTransport, Rapid I/O, or InfiniBand. FIG. 8B also depicts an embodiment in which local busses and direct communication are mixed: the processor

802 communicates with I/O device 830a using a local interconnected bus while communicating with I/O device 830b directly.

[00222] The computing device 800 may support any suitable installation device 816, such as a floppy disk drive for receiving floppy disks such as 3.5-inch, 5.25-inch disks or ZIP disks, a CD-ROM drive, a CD-R/RW drive, a DVD-ROM drive, tape drives of various formats, USB device, hard-drive or any other device suitable for installing software and programs such as any software 820, or portion thereof, related to the present invention or otherwise providing any of the techniques of the present invention. The computing device 800 may further comprise a storage device 828, such as one or more hard disk drives or redundant arrays of independent disks, for storing an operating system and other related software, and for storing application software programs such as any program related to the software 820 of the present invention. Optionally, any of the installation devices 816 could also be used as the storage device 828.

[00223] Furthermore, the computing device 800 may include a network interface 818 to interface to a Local Area Network (LAN), Wide Area Network (WAN) or the Internet through a variety of connections including, but not limited to, standard telephone lines, LAN or WAN links (*e.g.*, 802.11, T1, T3, 56kb, X.25), broadband connections (*e.g.*, ISDN, Frame Relay, ATM), wireless connections, or some combination of any or all of the above. The network interface 818 may comprise a built-in network adapter, network interface card, PCMCIA network card, card bus network adapter, wireless network adapter, USB network adapter, modem or any other device suitable for interfacing the computing device 800 to any type of network capable of communication and performing the operations described herein.

[00224] A wide variety of I/O devices 830a-830n may be present in the computing device 800. Input devices include keyboards, mice, trackpads, trackballs, microphones, and drawing tablets. Output devices include video displays, speakers, inkjet printers, laser printers, and dye-sublimation printers. The I/O devices may be controlled by an I/O controller 823 as shown in FIG. 8A. The I/O controller may control one or more I/O devices such as a keyboard 826 and a pointing device 827, *e.g.*, a mouse or optical pen. Furthermore, an I/O device may also provide storage 828 and/or an installation medium 816 for the computing device 800. In still other embodiments, the computing device 800 may provide USB connections to receive

handheld USB storage devices such as the USB Flash Drive line of devices manufactured by Twintech Industry, Inc. of Los Alamitos, California.

[00225] In further embodiments, an I/O device 830 may be a bridge 870 between the system bus 850 and an external communication bus, such as a USB bus, an Apple Desktop Bus, an RS-232 serial connection, a SCSI bus, a FireWire bus, a FireWire 800 bus, an Ethernet bus, an AppleTalk bus, a Gigabit Ethernet bus, an Asynchronous Transfer Mode bus, a HIPPI bus, a Super HIPPI bus, a SerialPlus bus, a SCI/LAMP bus, a FibreChannel bus, or a Serial Attached small computer system interface bus.

[00226] A computing device 800 of the sort depicted in FIGs. 8A and 8B typically operate under the control of operating systems, which control scheduling of tasks and access to system resources. The computing device 800 can be running any operating system such as any of the versions of the Microsoft® Windows operating systems, the different releases of the Unix and Linux operating systems, any version of the Mac OS® for Macintosh computers, any embedded operating system, any network operating system, any real-time operating system, any open source operating system, any proprietary operating system, any operating systems for mobile computing devices or network devices, or any other operating system capable of running on the computing device and performing the operations described herein. Typical operating systems include: WINDOWS 3.x, WINDOWS 95, WINDOWS 98, WINDOWS 2000, WINDOWS NT 3.51, WINDOWS NT 4.0, WINDOWS CE, and WINDOWS XP, all of which are manufactured by Microsoft Corporation of Redmond, Washington; MacOS, manufactured by Apple Computer of Cupertino, California; OS/2, manufactured by International Business Machines of Armonk, New York; and Linux, a freely-available operating system distributed by Caldera Corp. of Salt Lake City, Utah, or any type and/or form of a Unix operating system, among others.

[00227] In other embodiments, the computing device 800 may have different processors, operating systems, and input devices consistent with the device. The computing device 800 can be any workstation, desktop computer, laptop or notebook computer, server, handheld computer, mobile telephone, any other computer, or other form of computing or telecommunications device that is capable of communication and that has sufficient processor power and memory capacity to perform the operations of the present invention described herein. Moreover, the computing device 800 can be any type and/form of network device, such as a remote access device, a Virtual Private Network (VPN) device, a Secure Socket Layer (SSL) VPN device,

router, switch, bridge, or other network device in any form capable of performing the operations of the present invention described herein

[00228] Although the appliance of the present invention is generally discussing using or communicating via TCP/IP, the appliance may use or communicate with any modified transport control protocol, such as Transaction TCP (T/TCP), TCP with selection acknowledgements (TCP-SACK), TCP with large windows (TCP-LW), congestion prediction such as in the TCP-Vegas protocol, and TCP spoofing. Additionally, the appliance of the present invention may use any internode or high-performance protocol. Moreover, the appliance of the present invention may use or operate with any other transport and network protocols, such as Sequenced packet exchange (SPX) protocol over the Internetwork Packet Exchange (IPX) protocol.

Conclusion

[00229] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, it will be understood by those skilled in the relevant art(s) that various changes in form and details may be made therein without departing from the spirit and scope of the invention as defined in the appended claims. Accordingly, the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

WE CLAIM:

1. A method for caching a dynamically generated object not identified as cacheable, the method comprising the steps of:
 - (a) storing in a cache a dynamically generated object not identified as cacheable, the dynamically generated object served from an originating server;
 - (b) receiving, by the cache, a request to invalidate the cached dynamically generated object; and
 - (c) marking, by the cache in response to the request, the cached dynamically generated object as invalid.
2. The method of claim 1, comprising requesting, by the originating server to invalidate the cached dynamically generated object.
3. The method of claim 2, comprising automatically requesting, by the originating server, to invalidate the cached dynamically generated object in response to a change to the dynamically generated object in the originating server.
4. The method of claim 1, comprising requesting, by a client, to invalidate the cached dynamically generated object, the client in communication with the originating server to receive the dynamically generated object.
5. The method of claim 1, comprising requesting, by an external administrative control, to invalidate the cached dynamically generated object.
6. The method of claim 1, comprising flushing from the cache the cached dynamically generated object marked as invalid.
7. The method of claim 1, comprising receiving, by the cache, the request to invalidate within one of ten milliseconds or less of caching the dynamically generated object.

8. The method of claim 1, comprising receiving, by the cache, the request to invalidate responsive to expiration of an expiry of the cached dynamically generated object of one of ten milliseconds or less.

9. The method of claim 1, comprising operating the cache on one of an appliance, a network device or a computing device in communications between the originating server and a client.

10. A method for caching a group of dynamically generated objects having at least one object not identified as cacheable, the method comprising the steps of:
 - (a) identifying, in a cache, a group of dynamically generated objects previously served from an originating server, at least one of the dynamically generated objects not identified as cacheable;
 - (b) associating, by the cache, the group with an object determinant; and
 - (c) intercepting, by the cache, a communication identifying the object determinant of the group and indicating a change is one of about to occur or has occurred on the originating server to one of the objects of the group.

11. The method of claim 9, comprising marking, by the cache, the group of dynamically generated objects as invalid in response to receiving the communication.

12. The method of claim 10, comprising flushing from the cache the group of objects marked as invalid.

13. The method of claim 9, wherein the group of dynamically generated objects is pre-designated.

14. The method of claim 9, further comprising automatically identifying the group of dynamically generated objects and associating the object determinant with the group according to a rule.

15. A method for maintaining a cache of dynamically generated objects, the method comprising the steps of:

(a) intercepting, by a cache, a communication between a client and an originating server;

(b) identifying, by the cache, an object determinant in the communication indicating one of a change has occurred or will occur in a dynamically generated object at the originating server; and

(c) marking, by the cache, the dynamically generated object stored in the cache as invalid.

16. The method of claim 15, comprising obtaining the dynamically generated object from the originating server.

17. The method of claim 15, comprising flushing from the cache the invalid dynamically generated object.

18. The method of claim 15, wherein the dynamically generated object is not identified as cacheable.

19. The method of claim 15, comprising:

(e) associating, in the cache, the dynamically generated object with a group of dynamically generated objects previously served from the originating server;

(f) associating, by the cache, the group with the object determinant; and

(g) marking, by the cache, the group of dynamically generated objects as invalid in response to the communication.

20. The method of claim 19, comprising flushing from the cache the group of dynamically generated objects marked as invalid.

21. The method of claim 19, wherein the group of dynamically generated objects is pre-designated.

22. The method of claim 19, further comprising automatically identifying the group of dynamically generated objects and associating the object determinant with the group according to a rule.

23. The method of claim 15, comprising embedding by the client in the communication the object determinant as a pre-defined string.

24. The method of claim 15, comprising identifying the object determinant based on a pre-defined heuristic rule associated with the dynamically generated object.

25. The method of claim 15, comprising selecting, by the cache, the object determinant from one of the following elements of the communication: USERID, IP address, TCP port, HTTP header, custom HTTP header, client URL, cookie header, URL query string and POST body.
26. The method of claim 15, comprising extracting, by the cache, from the communication the object determinant based on a user-configured invalidation policy.
27. The method of claim 15, comprising maintaining a table in the cache to associate the object determinant with one or more objects or groups of objects stored in the cache.
28. The method of claim 15, comprising examining, by an intelligent statistical engine, communications from the client to identify a set of dynamically generated objects to store in the cache.
29. The method of claim 15, comprising associating, by the cache, objects stored in the cache into a content group, wherein the content group is represented by a hash table having an incarnation number as an index.
30. The method of claim 15, comprising performing, by the cache, a hash algorithm on the dynamically generated object identified via the communication to determine a change to the dynamically generated object.
31. A system for caching a dynamically generated object not identified as cacheable, the system comprising:
means for storing in a cache a dynamically generated object not identified as cacheable, the dynamically generated object served from an originating server;
means for receiving, by the cache, a request to invalidate the cached dynamically generated object; and
means for marking, by the cache in response to the request, the cached dynamically generated object as invalid.
32. The system of claim 31, comprising means for requesting, by the originating server to invalidate the cached dynamically generated object.
33. The system of claim 31, comprising means for automatically requesting, by the originating server, to invalidate the cached dynamically generated object in response to a change to the dynamically generated object in the originating server.

34. The system of claim 31, comprising means for requesting, by a client, to invalidate the cached dynamically generated object, the client in communication with the originating server to receive the dynamically generated object.
35. The system of claim 31, comprising means for requesting, by an external administrative control, to invalidate the cached dynamically generated object.
36. The system of claim 31, comprising means for flushing from the cache the cached dynamically generated object marked as invalid.
37. The system of claim 30, comprising means for receiving the invalidation request within one of ten milliseconds or less of caching the dynamically generated object.
38. The system of claim 31, comprising means for receiving the invalidation request within one of ten milliseconds or less of caching the dynamically generated object.
39. The system of claim 31, comprising means for receiving, by the cache, the request to invalidate responsive to expiration of an expiry of the cached dynamically generated object of one of ten milliseconds or less.
40. The system of claim 31, wherein the cache operates on one of an appliance, a network device or a computing device in communications between the originating server and a client.
41. A system for caching a group of dynamically generated objects having at least one object not identified as cacheable, the system comprising:
means for identifying, in a cache, a group of dynamically generated objects previously served from an originating server, at least one of the dynamically generated objects not identified as cacheable;
means for associating, by the cache, the group with an object determinant;
and

means for intercepting, by the cache, a communication identifying the object determinant of the group and indicating a change is one of about to occur or has occurred on the originating server to one of the objects of the group.

42. The system of claim 41, comprising means for marking, by the cache, the group of dynamically generated objects as invalid in response to receiving the communication.

43. The system of claim 41, comprising means for flushing from the cache the group of objects marked as invalid.

44. The system of claim 41, wherein the group of dynamically generated objects is pre-designated.

45. The system of claim 41, comprising means for automatically identifying the group of dynamically generated objects and associating the object determinant with the group according to a rule.

46. A system for maintaining a cache of dynamically generated objects, the system comprising:
means for intercepting, by a cache, a communication between a client and originating server;
means for identifying, by the cache, an object determinant in the communication that indicates one of a change has occurred or will occur in a dynamically generated object at the originating server; and
means for marking, by the cache, the dynamically generated object stored in cache as invalid.

47. The system of claim 46, comprising means for obtaining the dynamically generated object from the originating server.

48. The system of claim 46, comprising means for flushing from the cache the invalid dynamically generated object.

49. The system of claim 46, comprising:
means for associating, in the cache, the dynamically generated object with a group of dynamically generated objects previously served from the originating server;
means for associating, by the cache, the group with the object determinant;
and
means for marking, by the cache, the group of dynamically generated objects as invalid in response to the communication.
50. The system of claim 46, wherein the dynamically generated object is identified as non-cacheable.
51. The system of claim 46, comprising means for flushing from the cache the group of dynamically generated objects marked as invalid.
52. The system of claim 46, wherein the group of dynamically generated objects is pre-designated.
53. The system of claim 46, comprising means for automatically identifying the group of dynamically generated objects and associating the object determinant with the group according to a rule.
54. The system of claim 46, comprising means for embedding by the client in the communication the object determinant as a pre-defined string.
55. The system of claim 46, comprising means for identifying the object determinant based on a pre-defined heuristic rule associated with the dynamically generated object.
56. The system of claim 46, comprising means for selecting, by the cache, the object determinant from one of the following elements of the communication: USERID, IP address, TCP port, HTTP header, custom HTTP header, client URL, cookie header, URL query string and POST body.
57. The system of claim 46, comprising means for extracting, by the cache, from the communication the object determinant based on a user-configured invalidation policy.
58. The system of claim 46, comprising means for maintaining a table in the cache to associate the object determinant with to one or more objects or groups of objects stored in the cache.

59. The system of claim 46, comprising means for examining, by an intelligent statistical engine, communications from the client to identify a set of dynamically generated objects to store in the cache.

60. The system of claim 46, comprising means for associating, by the cache, objects stored in cache into a content group, wherein the content group is represented by a hash table having an incarnation number as an index.

61. The system of claim 46, comprising means for performing, by the cache, a hash algorithm on the dynamically generated object identified via the communication to determine a change to the dynamically generated object.

62. In a packet processing engine of a device, the packet processing engine processing network packets comprising at least one dynamically generated object and storing the at least one dynamically generated object in a cache, a method for invalidating a cached dynamically generated object in the cache responsive to a signal of a packet processing timer, the method comprising the steps of:

- (a) receiving a first signal from a packet processing timer to process a network packet;
- (b) processing the network packet responsive to the signal; and
- (c) transmitting a second signal to a cache indicating a cached dynamically generated object is invalid.

63. The method of claim 62, comprising providing, by the second signal, an invalidation command to invalidate the cached dynamically generated object.

64. The method of claim 62, comprising providing, by the second signal, an expiration of an expiry of the cached dynamically generated object

65. The method of claim 62, wherein step (c) is performed during processing of the network packet

66. The method of claim 62, wherein step (c) is performed upon completion of processing of the network packet.

67. The method of claim 62, comprising operating the packet processing timer at a time increment of one or more milliseconds.

68. The method of claim 62, comprising operating the packet processing timer at a time increment to invalidate an object stored in cache at an expiry time of one of

equal to or less than 10 milliseconds.

69. The method of claim 62, comprising marking, by the cache, the cached dynamically generated object as invalid.

70. The method of claim 62, wherein the device comprises one of the following: a bridge, a router, a switch, and a SSL VPN device.

71. The method of claim 62, comprising operating one of the cache, the packet processing engine, or the packet processing timer in a kernel space of the device.

72. The method of claim 62, comprising providing, via the packet processing timer, a time interval for expiry of the cached dynamically generated object one of less than or equal to a duration of time to dynamically generate and serve the dynamically generated object from an originating server.

73. The method of claim 62, comprising transmitting from an originating server the network packet comprising dynamically generated object.

74. The method of claim 62, wherein the dynamically generated object is identified as not cacheable.

75. A device for invalidating a cached dynamically generated object responsive to a signal of a packet processing timer, the device processing network packets comprising at least one dynamically generated object and storing the at least one dynamically generated object in a cache, the device comprising:

a packet processing timer generating a first signal ;

a packet processing engine, in communication with the packet processing timer, processing the network packet responsive to receipt of the first signal of the packet processing timer; and

a cache, in communication with the packet processing engine, having a dynamically generated object cached in a storage element, the cache receiving a second signal from the packet processing engine indicating the cached dynamically generated object is invalid.

76. The device of claim 75, wherein the second signal comprises an invalidation command to invalidate the cached dynamically generated object.

77. The device of claim 75, wherein the second signal comprises an expiration of an expiry of the cached dynamically generated object.

78. The device of claim 75, wherein the packet processing engine transmits the second signal to the cache during processing of the network packet by the packet

processing engine.

79. The device of claim 75, wherein the packet processing engine transmits the second signal to the cache upon completion of processing of the network packet.

80. The device of claim 75, wherein the packet processing timer operates at a time interval of one or more milliseconds.

81. The device of claim 75, wherein the packet processing timer operates at a time interval to trigger the second signal to the cache to invalidate the cached object at a first time interval of one of equal to or less than 10 milliseconds.

82. The device of claim 75, wherein the device comprises one of the following: a bridge, a router, a switch, and a SSL VPN device.

83. The device of claim 75, wherein one of the cache, the packet processing engine, or the packet processing timer operates in a kernel space of the device.

84. The device of claim 75, wherein the packet processing timer provides a time interval for expiry of the cached dynamically generated object one of less than or equal to a duration of time to dynamically generate and serve the dynamically generated object from an originating server.

85. The device of claim 75, wherein an originating server transmits the network packet comprising the dynamically generated object.

86. The device of claim 75, wherein the dynamically generated object is not identified as cacheable.

87. A method for responding to a request for a dynamically generated object from a plurality of clients, the method comprising the steps of:

(d) receiving, by a device, a response from an originating server to a first request of a first client for a dynamically generated object, the response comprising the dynamically generated object;

(e) requesting, by the device, to a packet processing engine to transmit the response to the first client, the response held in a buffer while waiting to be transmitted;

(f) receiving, by the device, from a second client a second request for the dynamically generated object prior to completing transmission of the response to the first client; and

(g) requesting, by the device, to the packet processing engine to transmit the dynamically generated object of the response to the first client held in the buffer to the second client.

88. The method of claim 87, comprising transmitting, by the packet processing engine, the response to the first client.
89. The method of claim 87, comprising transmitting, by the packet processing engine, the dynamically generated object to the second client in response to the second request.
90. The method of claim 87, comprising removing, by the packet processing engine, the response from the buffer after transmitting the response to one of the first client or the second client.
91. The method of claim 87, wherein the buffer comprises one of a TCP/IP buffer or a portion of a TCP/IP stack of the device.
92. The method of claim 87, wherein the device comprises a single TCP/IP stack for transmitting to the first client and the second client, the buffer associated with the single TCP/IP stack.
93. The method of claim 87, wherein step (c) further comprises receiving, by the device, from a plurality of clients a plurality of requests for the dynamically generated object, and wherein step (d) further comprises requesting, by the device, to transmit the dynamically generated object of the response to the first client held in the buffer to the plurality of clients.
94. The method of claim 87, comprising holding, by the packet processing engine, the response to the first client based on one or more characteristics related to a connection of the first client to a network
95. The method of claim 87, comprising queuing, by the device, the second request.
96. The method of claim 87, comprising removing the dynamically generated object from the buffer after transmitting the dynamically generated object to the first client and the second client.
97. The method of claim 87, wherein the dynamically generated object is identified as non-cacheable.
98. The method of claim 87, comprising determining, by the device, the second request is for the dynamically generated object of the first request via an object determinant associated with the second request.
99. The method of claim 87, wherein the device comprises on one of an appliance, a network device or a computing device in communication between the first client and the originating server.

100. The method of claim 87, wherein the device comprise a cache in communication with the packet processing engine.

101. The method of claim 100, comprising requesting, by the cache, the packet processing engine to transmit one of the response to the first client or the dynamically generated object of the first client to the second client.

102. A system for responding to a request for a dynamically generated object from a plurality of clients, the system comprising:

a device receiving a response from an originating server to a first request of a first client for a dynamically generated object, the response comprising the dynamically generated object, the device comprising:

a buffer for holding one or more network packets waiting to be transmitted from the device;

a packet processing engine receiving a request from the device to transmit the response to the first client, the buffer holding one or more network packets comprising the response; and

wherein the device receives from a second client a second request for the dynamically generated object prior to completing transmission of the response to the first client, and requests the packet processing engine to transmit the dynamically generated object of the response held in the buffer to the second client.

103. The system of claim 102, wherein the packet processing engine transmits from the buffer the response to the first client.

104. The system of claim 102, wherein the packet processing engine transmits from the buffer the dynamically generated object to the second client in response to the second request.

105. The system of claim 102, wherein the packet processing engine removes the response from the buffer after transmitting the response to one of the first client or the second client.

106. The system of claim 102, wherein the buffer comprises one of a TCP/IP buffer or a portion of a TCP/IP stack of the device.

107. The system of claim 102, wherein the device comprises a single TCP/IP stack for transmitting to the first client and the second client the buffer associated with the single TCP/IP stack.

108. The system of claim 102, wherein the device receives from a plurality of

clients a plurality of requests for the dynamically generated object, and wherein the device requests the packet processing engine to transmit the dynamically generated object of the response to the first client held in the buffer to the plurality of clients.

109. The system of claim 102, wherein the packet processing engine holds the response to the first client based on one or more characteristics related to a connection of the first client to a network

110. The system of claim 102, wherein the device queues the second request.

111. The system of claim 102, wherein the packet processing engine remove the dynamically generated object from the buffer after transmitting the dynamically generated object to the first client and the second client.

112. The system of claim 102, wherein the dynamically generated object is identified as non-cacheable.

113. The system of claim 102, comprising determining, by the device, the second request is for the dynamically generated object of the first request via an object determinant associated with the second request.

114. The system of claim 102, wherein the device comprises on one of an appliance, a network device or a computing device in communication between the first client and the originating server.

115. The system of claim 102, wherein the device comprise a cache in communication with the packet processing engine.

116. The system of claim 115, wherein the cache requests the packet processing engine to transmit one of the response to the first client or the dynamically generated object of the first client to the second client.

117. In a network device, a method for responding to a request for a dynamically generated object from a plurality of clients, the method comprising the steps of:

(h) receiving from a first client a first request for a dynamically generated object from an originating server;

(i) transmitting the first request to the originating server;

(j) receiving from a second client a second request for the dynamically generated object prior to responding to the first request of the first client; and

(k) receiving the response to the first request from the originating server, the response comprising the dynamically generated object; and

(l) transmitting the dynamically generated object to the first client in

response to the first request, and the dynamically generated object to the second client in response to the second request.

118. The method of claim 117, wherein step (c) further comprises receiving from a plurality of clients a plurality of requests for the dynamically generated object, and wherein step (e) further comprises transmitting the dynamically generated object to the plurality of clients in response to each client's request.

119. The method of claim 117, comprising the steps of:

(f) receiving a third request for the dynamically generated object from a third client prior to completing transmission of the dynamically generated object to the second client; and

(g) transmitting the dynamically generated object to the third client in response to the third request.

120. The method of claim 117, comprising receiving the first request and the second request nearly simultaneously.

121. The method of claim 117, comprising queuing the second request.

122. The method of claim 117, comprising flushing the dynamically generated object from a cache after transmitting the dynamically generated object to the first client and the second client.

123. The method of claim 117, wherein the dynamically generated object is identified as non-cacheable.

124. The method of claim 117, comprising determining the second request is for the dynamically generated object of the first request via an object determinant associated with the second request.

125. The method of claim 117, comprising modifying the response to include one of an entity tag header or a cache-control header and transmitting the modified response to one of the first client or the second client

126. A network device for responding to a request for a dynamically generated object from a plurality of clients, the network device comprising:

means for receiving from a first client a first request for a dynamically generated object from an originating server;

means for transmitting the first request to the originating server;

means for receiving from a second client a second request for the dynamically generated object prior to responding to the first request of the first client; and

means for receiving the response to the first request from the originating server, the response comprising the dynamically generated object; and

means for transmitting the dynamically generated object to the first client in response to the first request, and the dynamically generated object to the second client in response to the second request.

127. The network device of claim 126, further comprising means for receiving from a plurality of clients a plurality of requests for the dynamically generated object, and transmitting the dynamically generated object to the plurality of clients in response to each client's request.

128. The network device of claim 126, comprising

means for receiving a third request for the dynamically generated object from a third client prior to completing transmission of the dynamically generated object to the second client; and

means for transmitting the dynamically generated object, to the third client in response to the third request.

129. The network device of claim 126, comprising means for receiving the first request and the second request nearly simultaneously.

130. The network device of claim 126, wherein the network device queues the second request.

131. The network device of claim 126, wherein the network device flushes the dynamically generated object from a cache after transmitting the dynamically generated object to the first client and the second client.

132. The network device of claim 126, wherein the dynamically generated object is not identified as cacheable.

133. The network device of claim 126, comprising means for determining the second request is for the dynamically generated object of the first request via an object determinant associated with the second request.

134. The network device of claim 126, wherein the network device comprises one of an appliance or a computing device in communication between the client and the originating server.

135. The network device of claim 126, comprising means for modifying the response to include one of an entity tag header or a cache-control header and transmitting the modified response to one of the first client or the second client.

136. A method of using an entity tag for caching a dynamically generated object not identified as cacheable, the method comprising the steps of:

(m) receiving, by a cache, a response to a request of a client for a dynamically generated object from an originating server, the dynamically generated object identified as not cacheable,

(n) generating, by the cache, an entity tag for the dynamically generated object stored in the cache;

(o) modifying, by the cache, the response to include the generated entity tag; and

(p) transmitting, by the cache, to the client the modified response identifying the dynamically generated object as cacheable.

137. The method of claim 136, wherein the response received from the originating server does not comprise an entity tag.

138. The method of claim 136, comprising modifying, by the cache, the response to identify a duration for the client to cache the dynamically generated object.

139. The method of claim 136, comprising modifying, by the cache, the response to indicate to the client to store the response and serve the response after checking if the dynamically generated object has changed with one of the cache or the originating server.

140. The method of claim 136, wherein the response from the originating server does not have cache control information for the dynamically generated object, and the cache modifies the response transmitted to the client to include cache control information.

141. The method of claim 136, comprising generating, by the cache, a number for the entity tag.

142. The method of claim 136, comprising storing, by the cache, the generated entity tag with the dynamically generated object.

143. The method of claim 142, comprising storing, by the cache, cache control information with the dynamically generated object.

144. The method of claim 143, comprising serving, by the cache, the cached dynamically generated object with the generated entity tag and cache control information.
145. The method of claim 136, comprising operating the cache on one of an appliance, network device or computing device.
146. The method of claim 136, wherein one of the request or the response comprises HyperText Markup Language (HTML).
147. The method of claim 146, comprising modifying, by the cache, the response to include an etag header identifying the generated entity tag.
148. The method of claim 146, comprising modifying, by the cache, the response to include a cache-control header.
149. The method of claim 146, comprising receiving, by the cache, from the client a second request having an If-None-Match header, and in response to the second request, comparing the etag value associated with the If-None-Match header to the entity tag stored in the cache with the dynamically generated object.
150. The method of claim 149, comprising transmitting, by the cache, to the client a not modified response to the second request if the etag value matches the entity tag.
151. The method of claim 149, comprising transmitting, by the cache, to the client an updated dynamically generated object with a second entity tag if the etag value does not match the entity tag.
152. A system for using of using an entity tag for caching a dynamically generated object identified as not cacheable, the system comprising:
- means for receiving, by a cache, a response to a request of a client for a dynamically generated object from an originating server, the dynamically generated object identified as not cacheable,
 - means for generating, by the cache, an entity tag for the dynamically generated object stored in the cache;
 - means for modifying, by the cache, the response to include the generated entity tag; and
 - means for transmitting, by the cache, to the client the modified response identifying the dynamically generated object as cacheable.

153. The system of claim 152, wherein the response received from the originating server does not comprise an entity tag.
154. The system of claim 152, comprising means for modifying, by the cache, the response to identify a duration for the client to cache the dynamically generated object.
155. The system of claim 152, comprising means for modifying, by the cache, the response to indicate to the client to store the response and serve the response after checking if the dynamically generated object has changed with one of the cache or the originating server.
156. The system of claim 152, wherein the response from the originating server does not have cache control information for the dynamically generated object, and the cache modifies the response transmitted to the client to include cache control information.
157. The system of claim 152, comprising means for generating, by the cache, a number for the entity tag.
158. The system of claim 152, comprising means for storing, by the cache, the generated entity tag with the dynamically generated object.
159. The system of claim 152, comprising means for storing, by the cache, cache control information with the dynamically generated object.
160. The system of claim 159, comprising means for serving, by the cache, the cached dynamically generated object with the generated entity tag and cache control information.
161. The system of claim 152, comprising means for operating the cache on one of an appliance, network device or computing device.
162. The system of claim 152, wherein one of the request or the response comprises Hypertext Markup Language (HTML).
163. The system of claim 162, comprising means for modifying, by the cache, the response to include an etag header identifying the generated entity tag.
164. The system of claim 162, comprising means for modifying, by the cache, the response to include a cache-control header.
165. The system of claim 164, comprising means for comprising receiving, by the cache, from the client a second request having a If-None-Match header, and in response to the second request, comparing the etag value associated with the If-None-Match header to the entity tag stored in the cache with the dynamically generated

object.

166. The system of claim 165, comprising means for transmitting, by the cache, to the client a not modified response to the second request if the etag value matches the entity tag.

167. The system of claim 165, comprising means for transmitting, by the cache, to the client an updated dynamically generated object with a second entity tag if the etag value does not match the entity tag.




FIG. 1

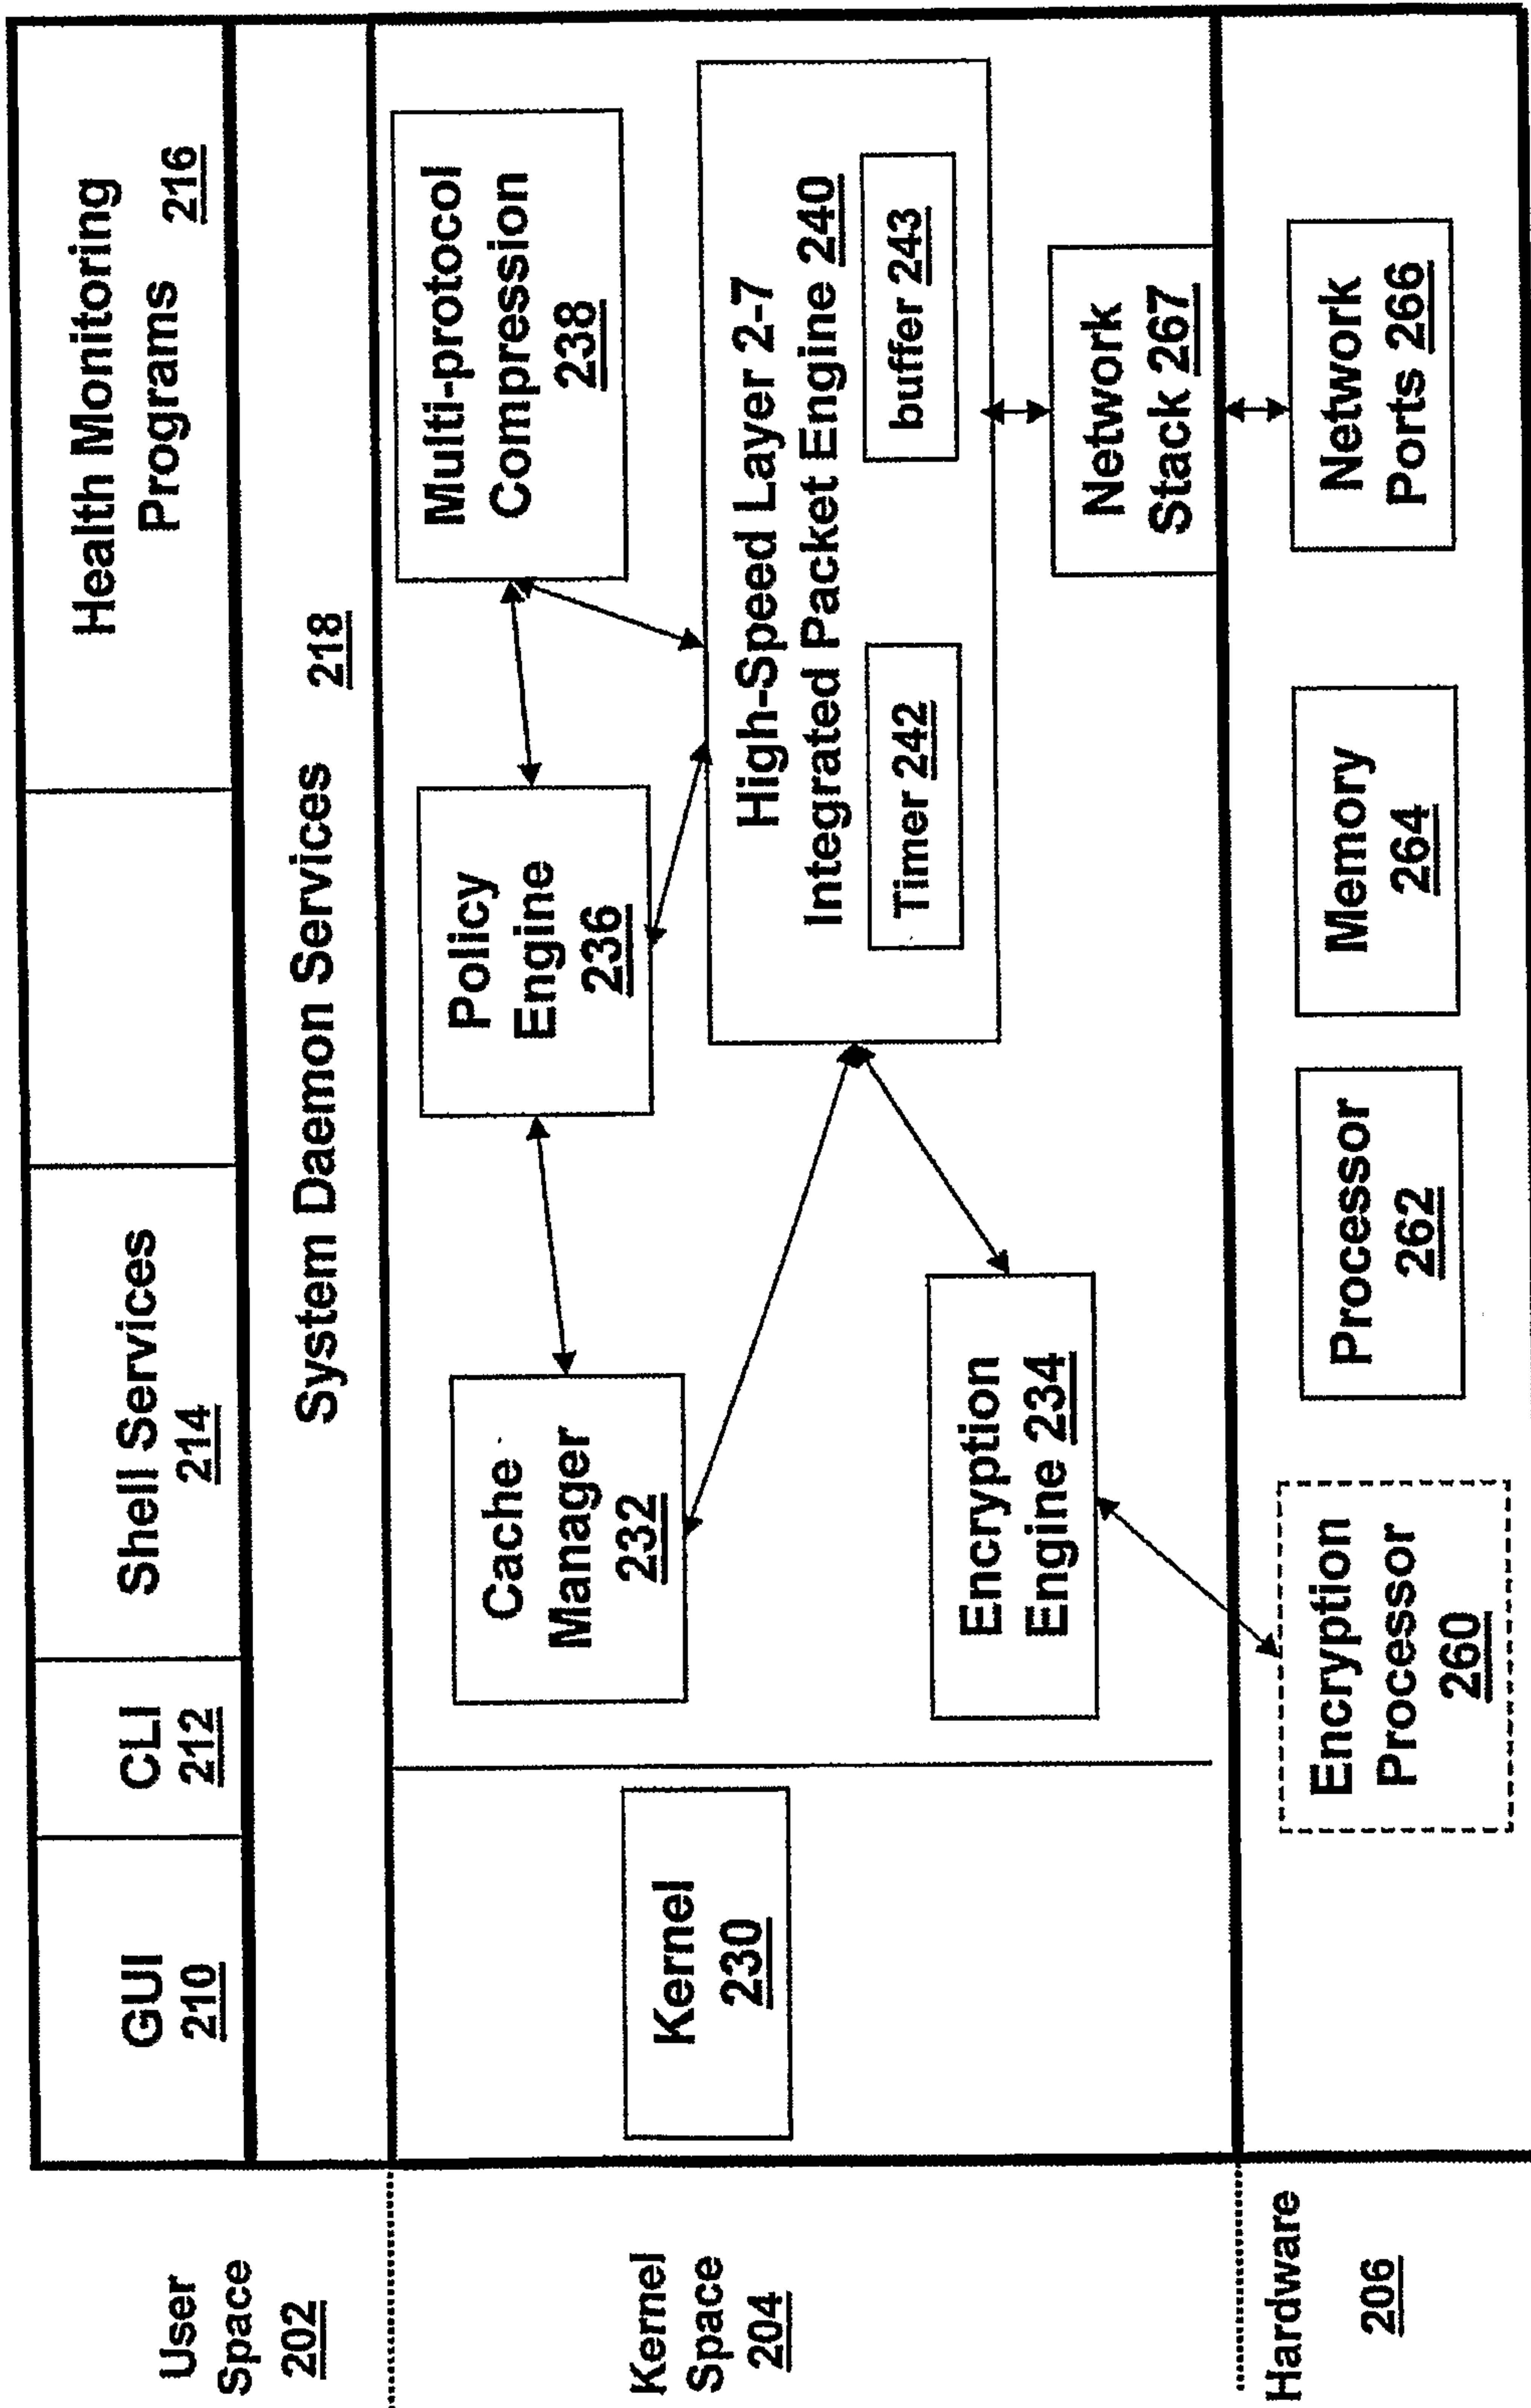


FIG. 2

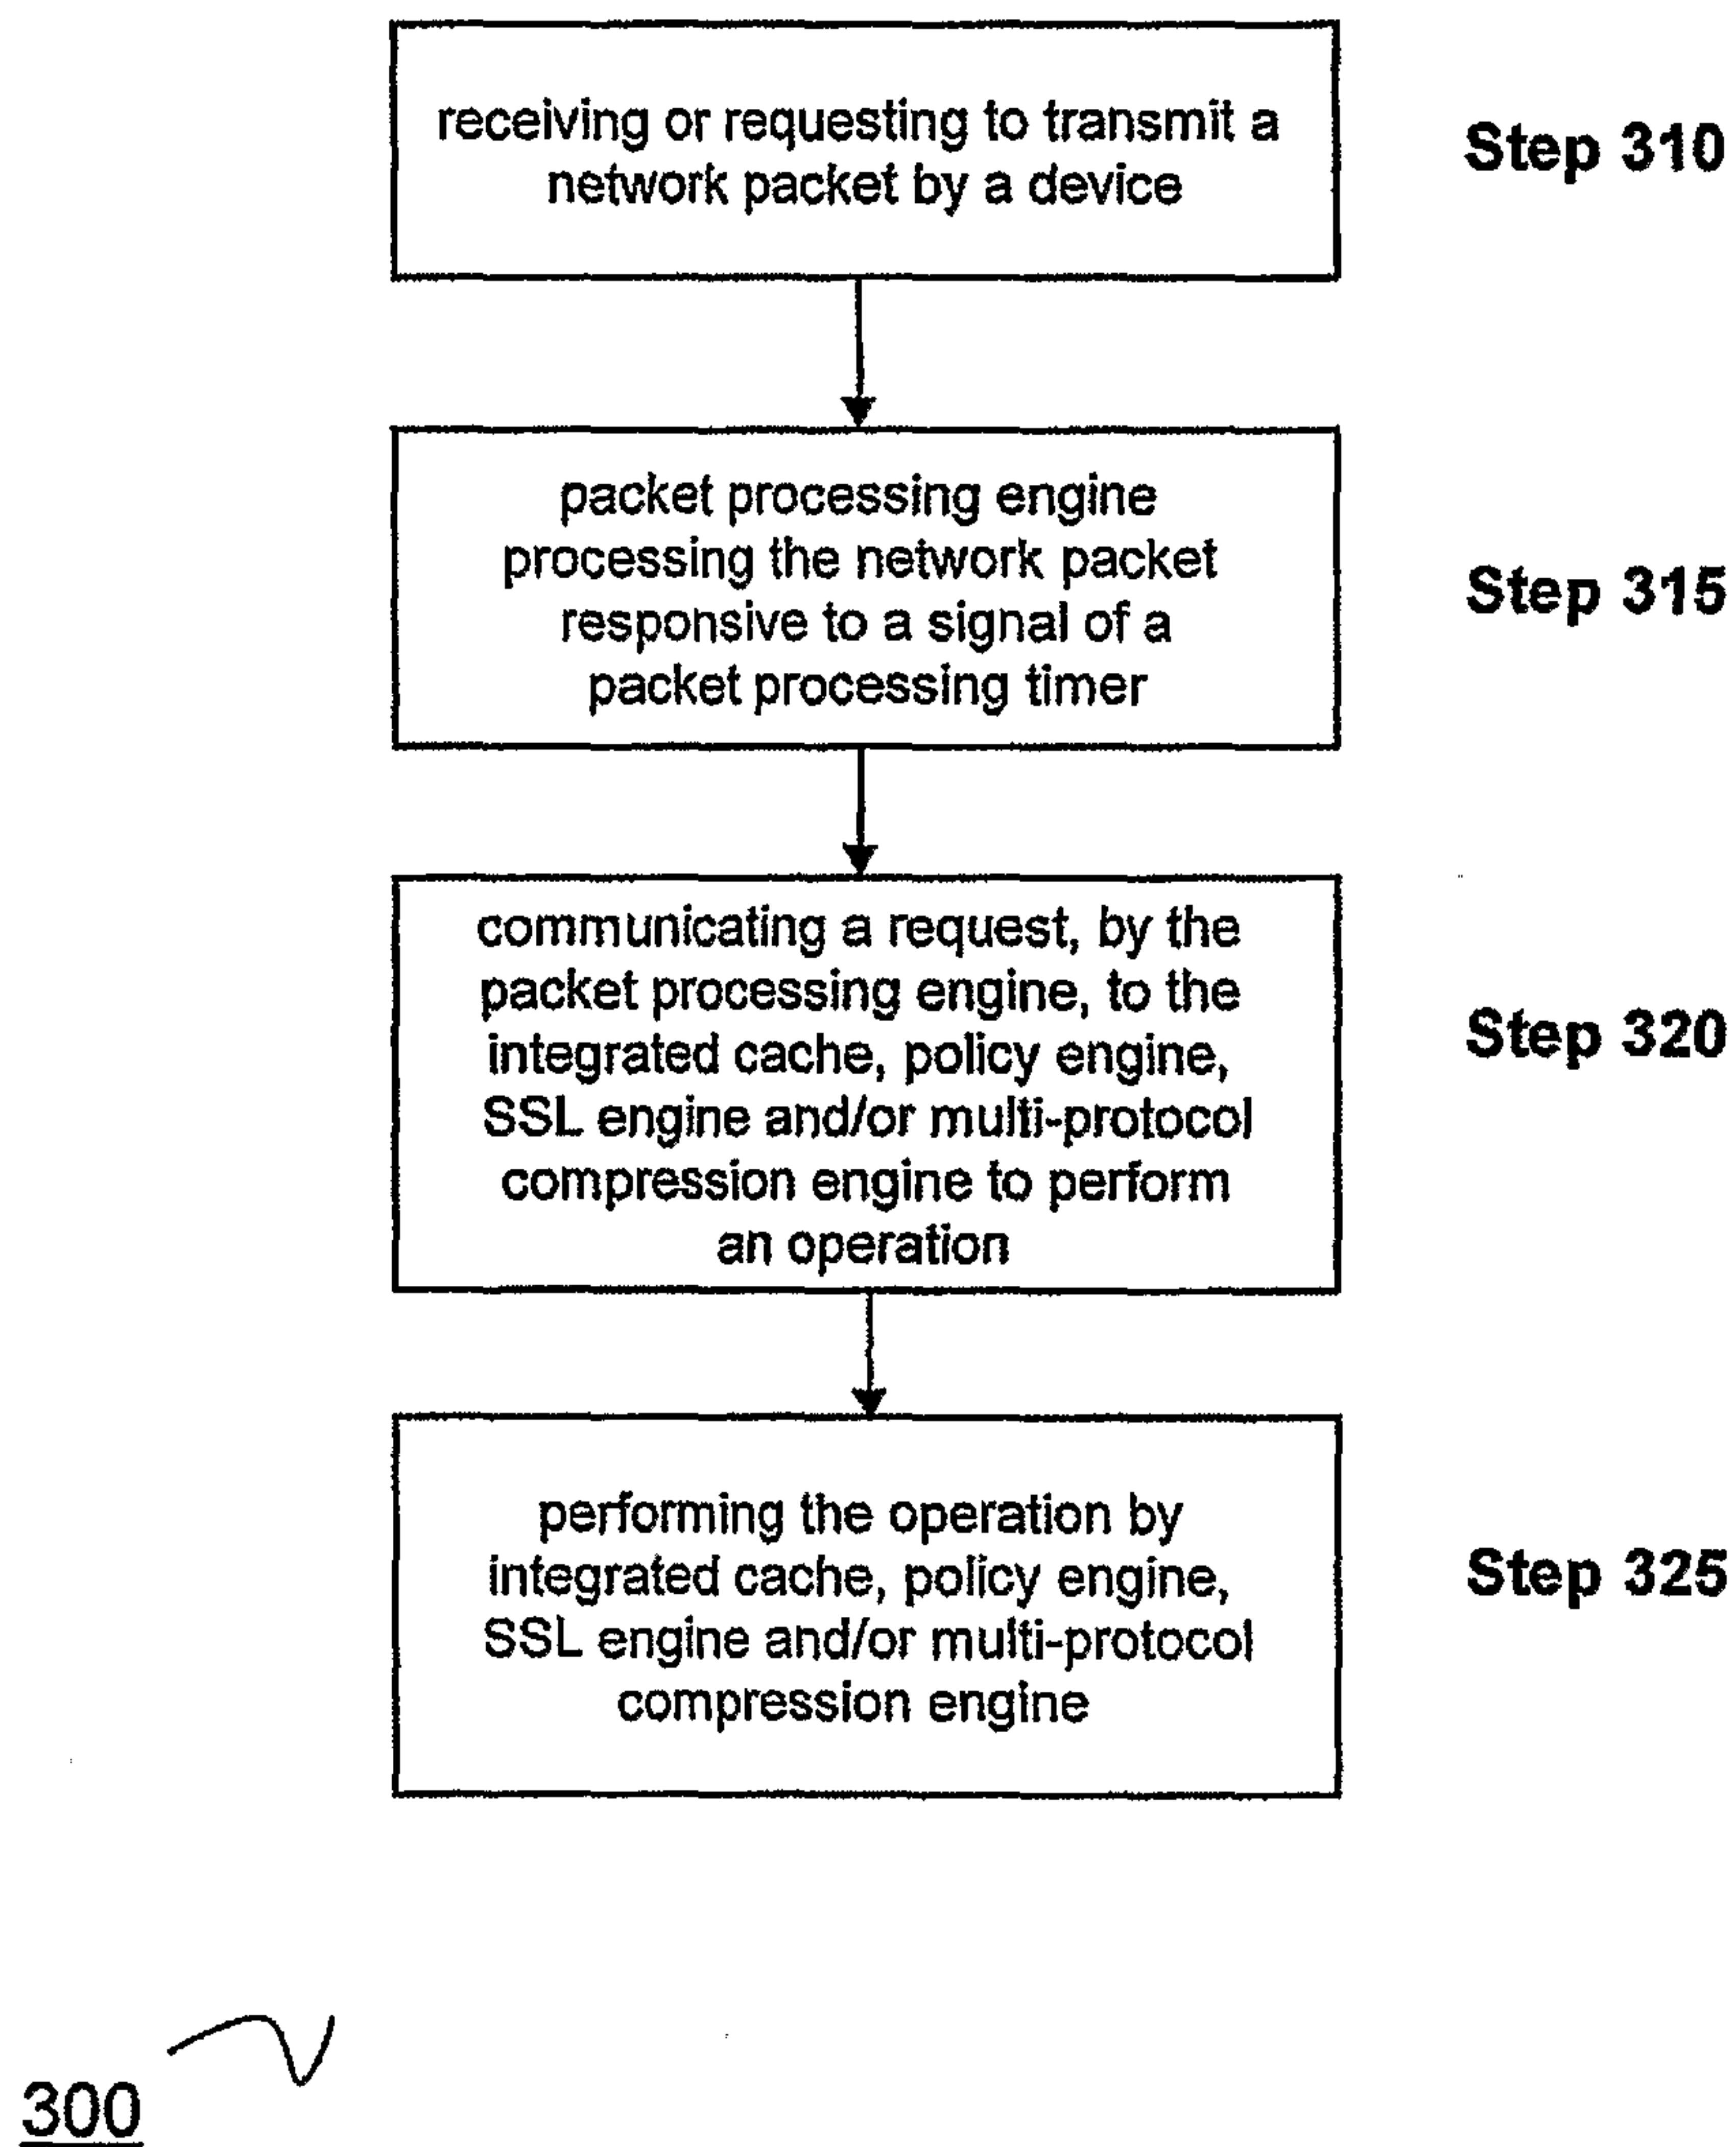
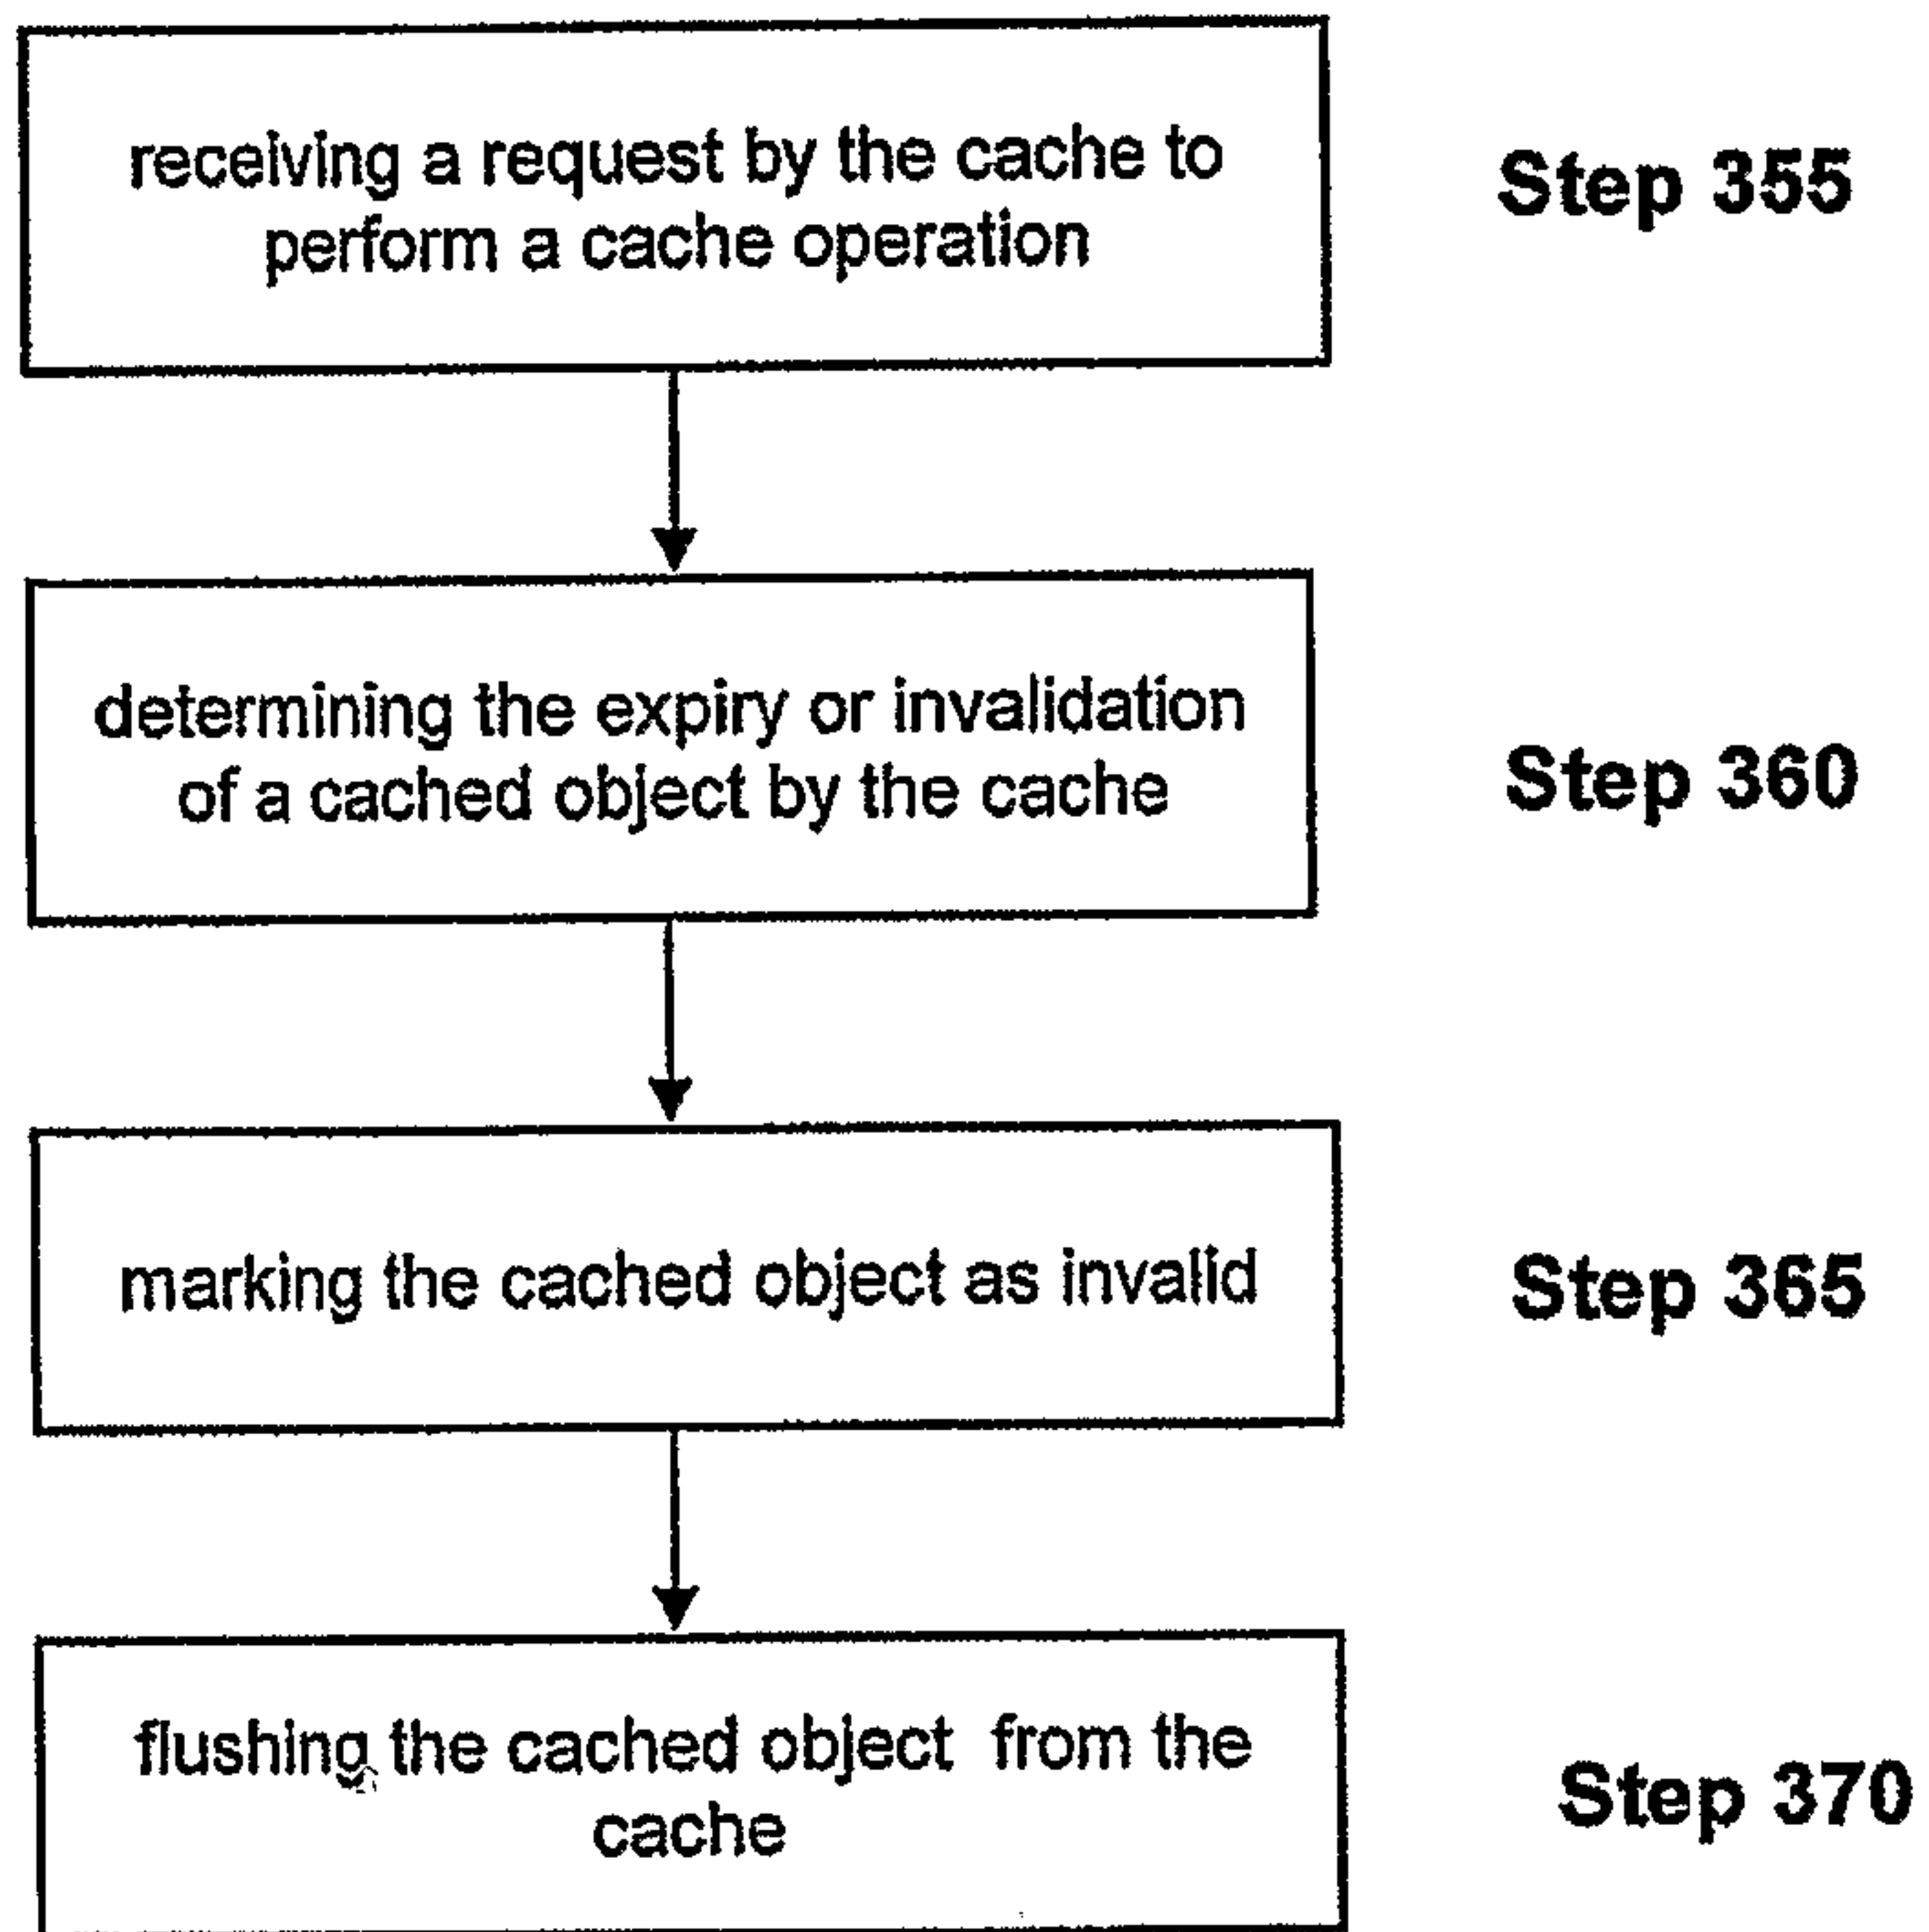


FIG. 3A



350 ~

FIG. 3B

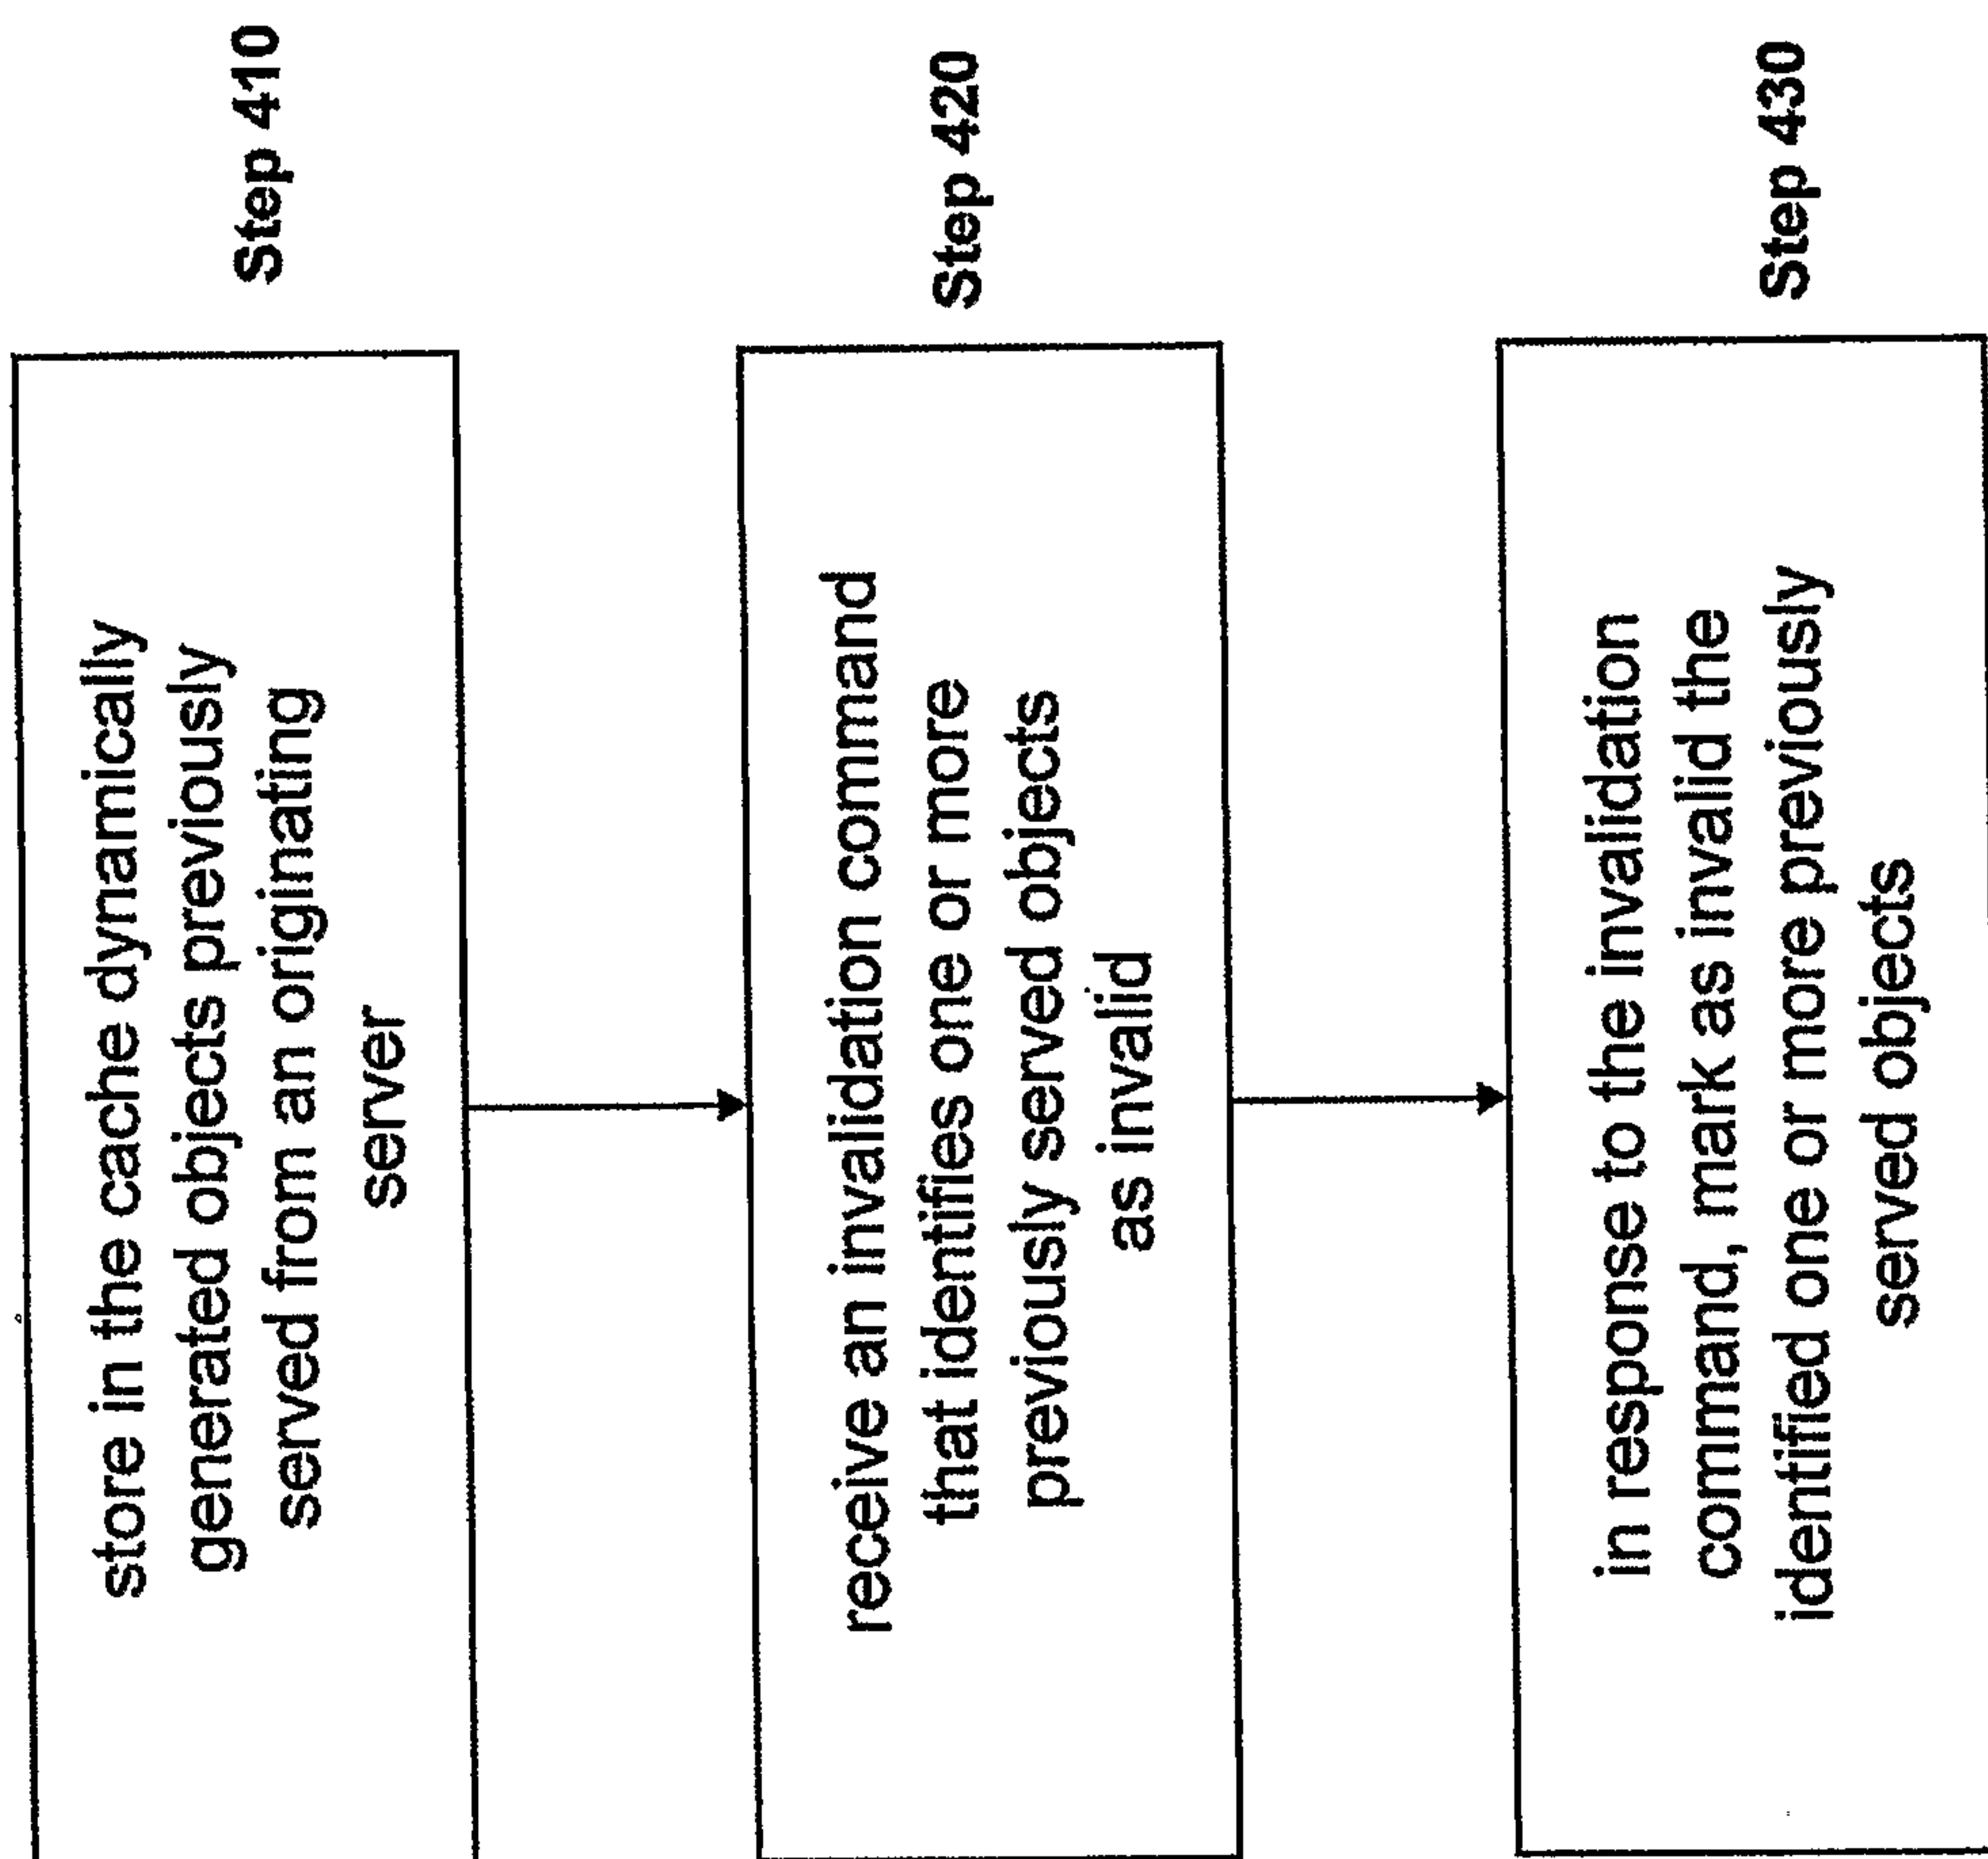


FIG. 4A

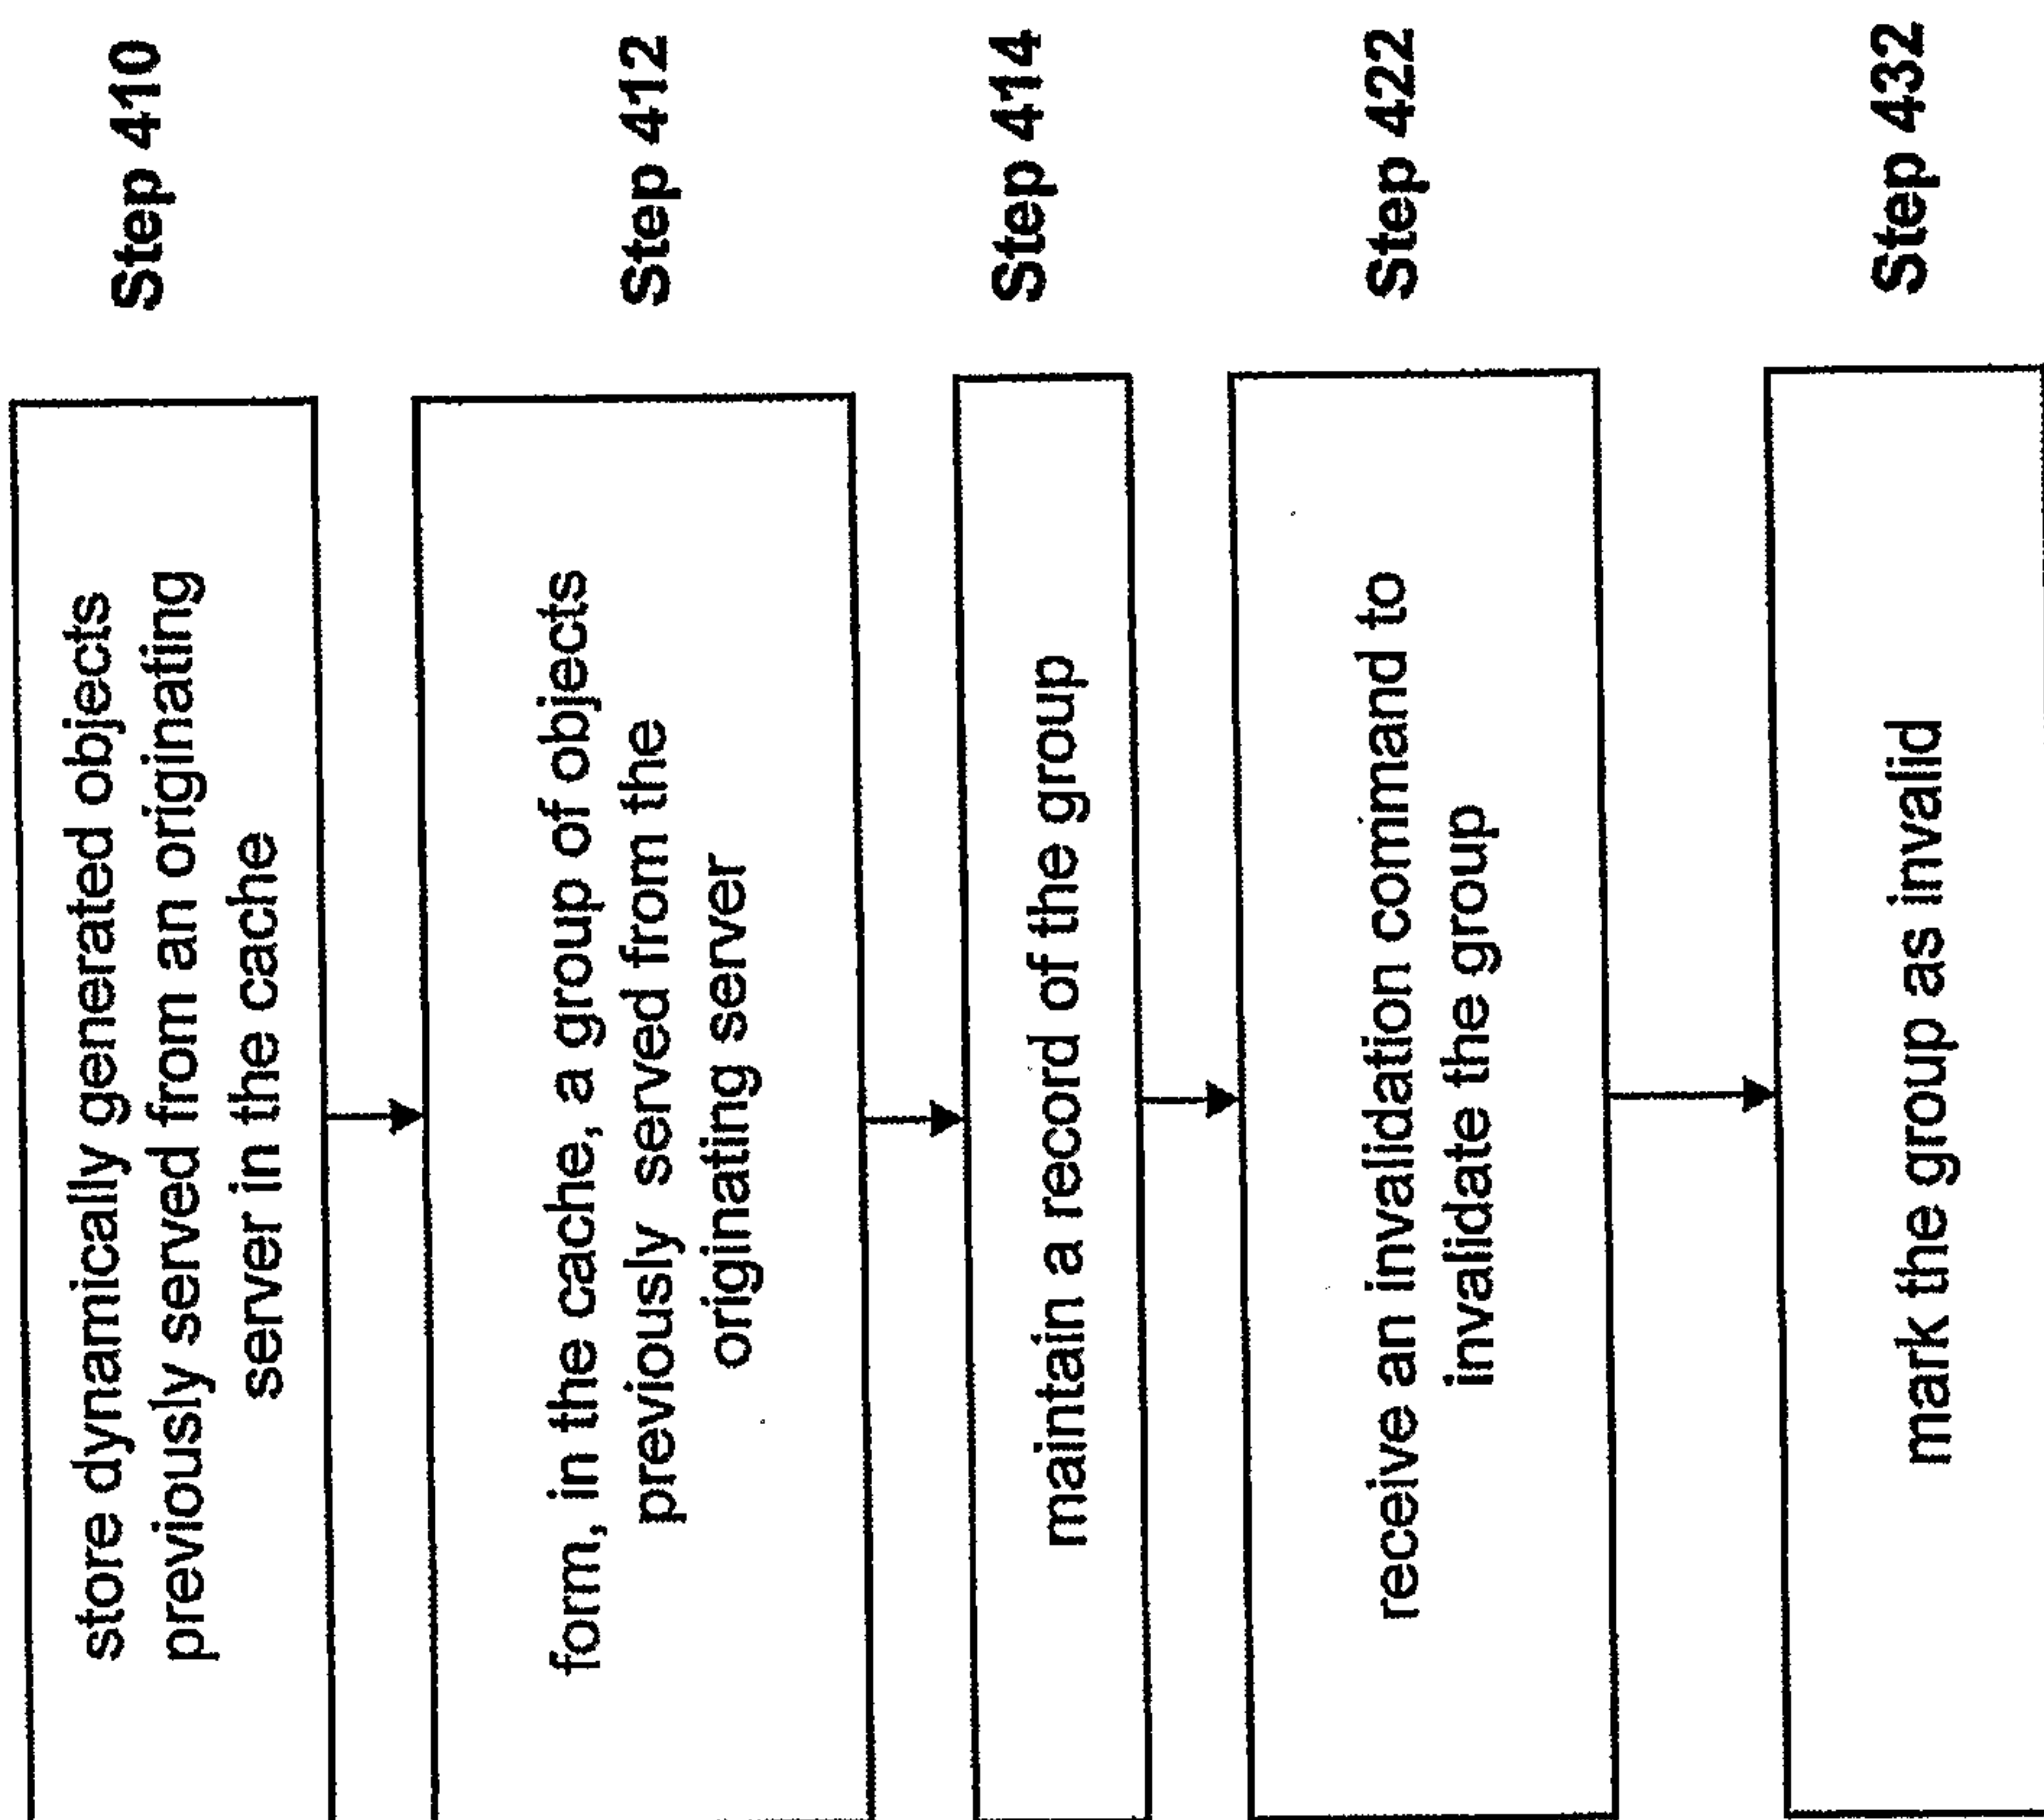
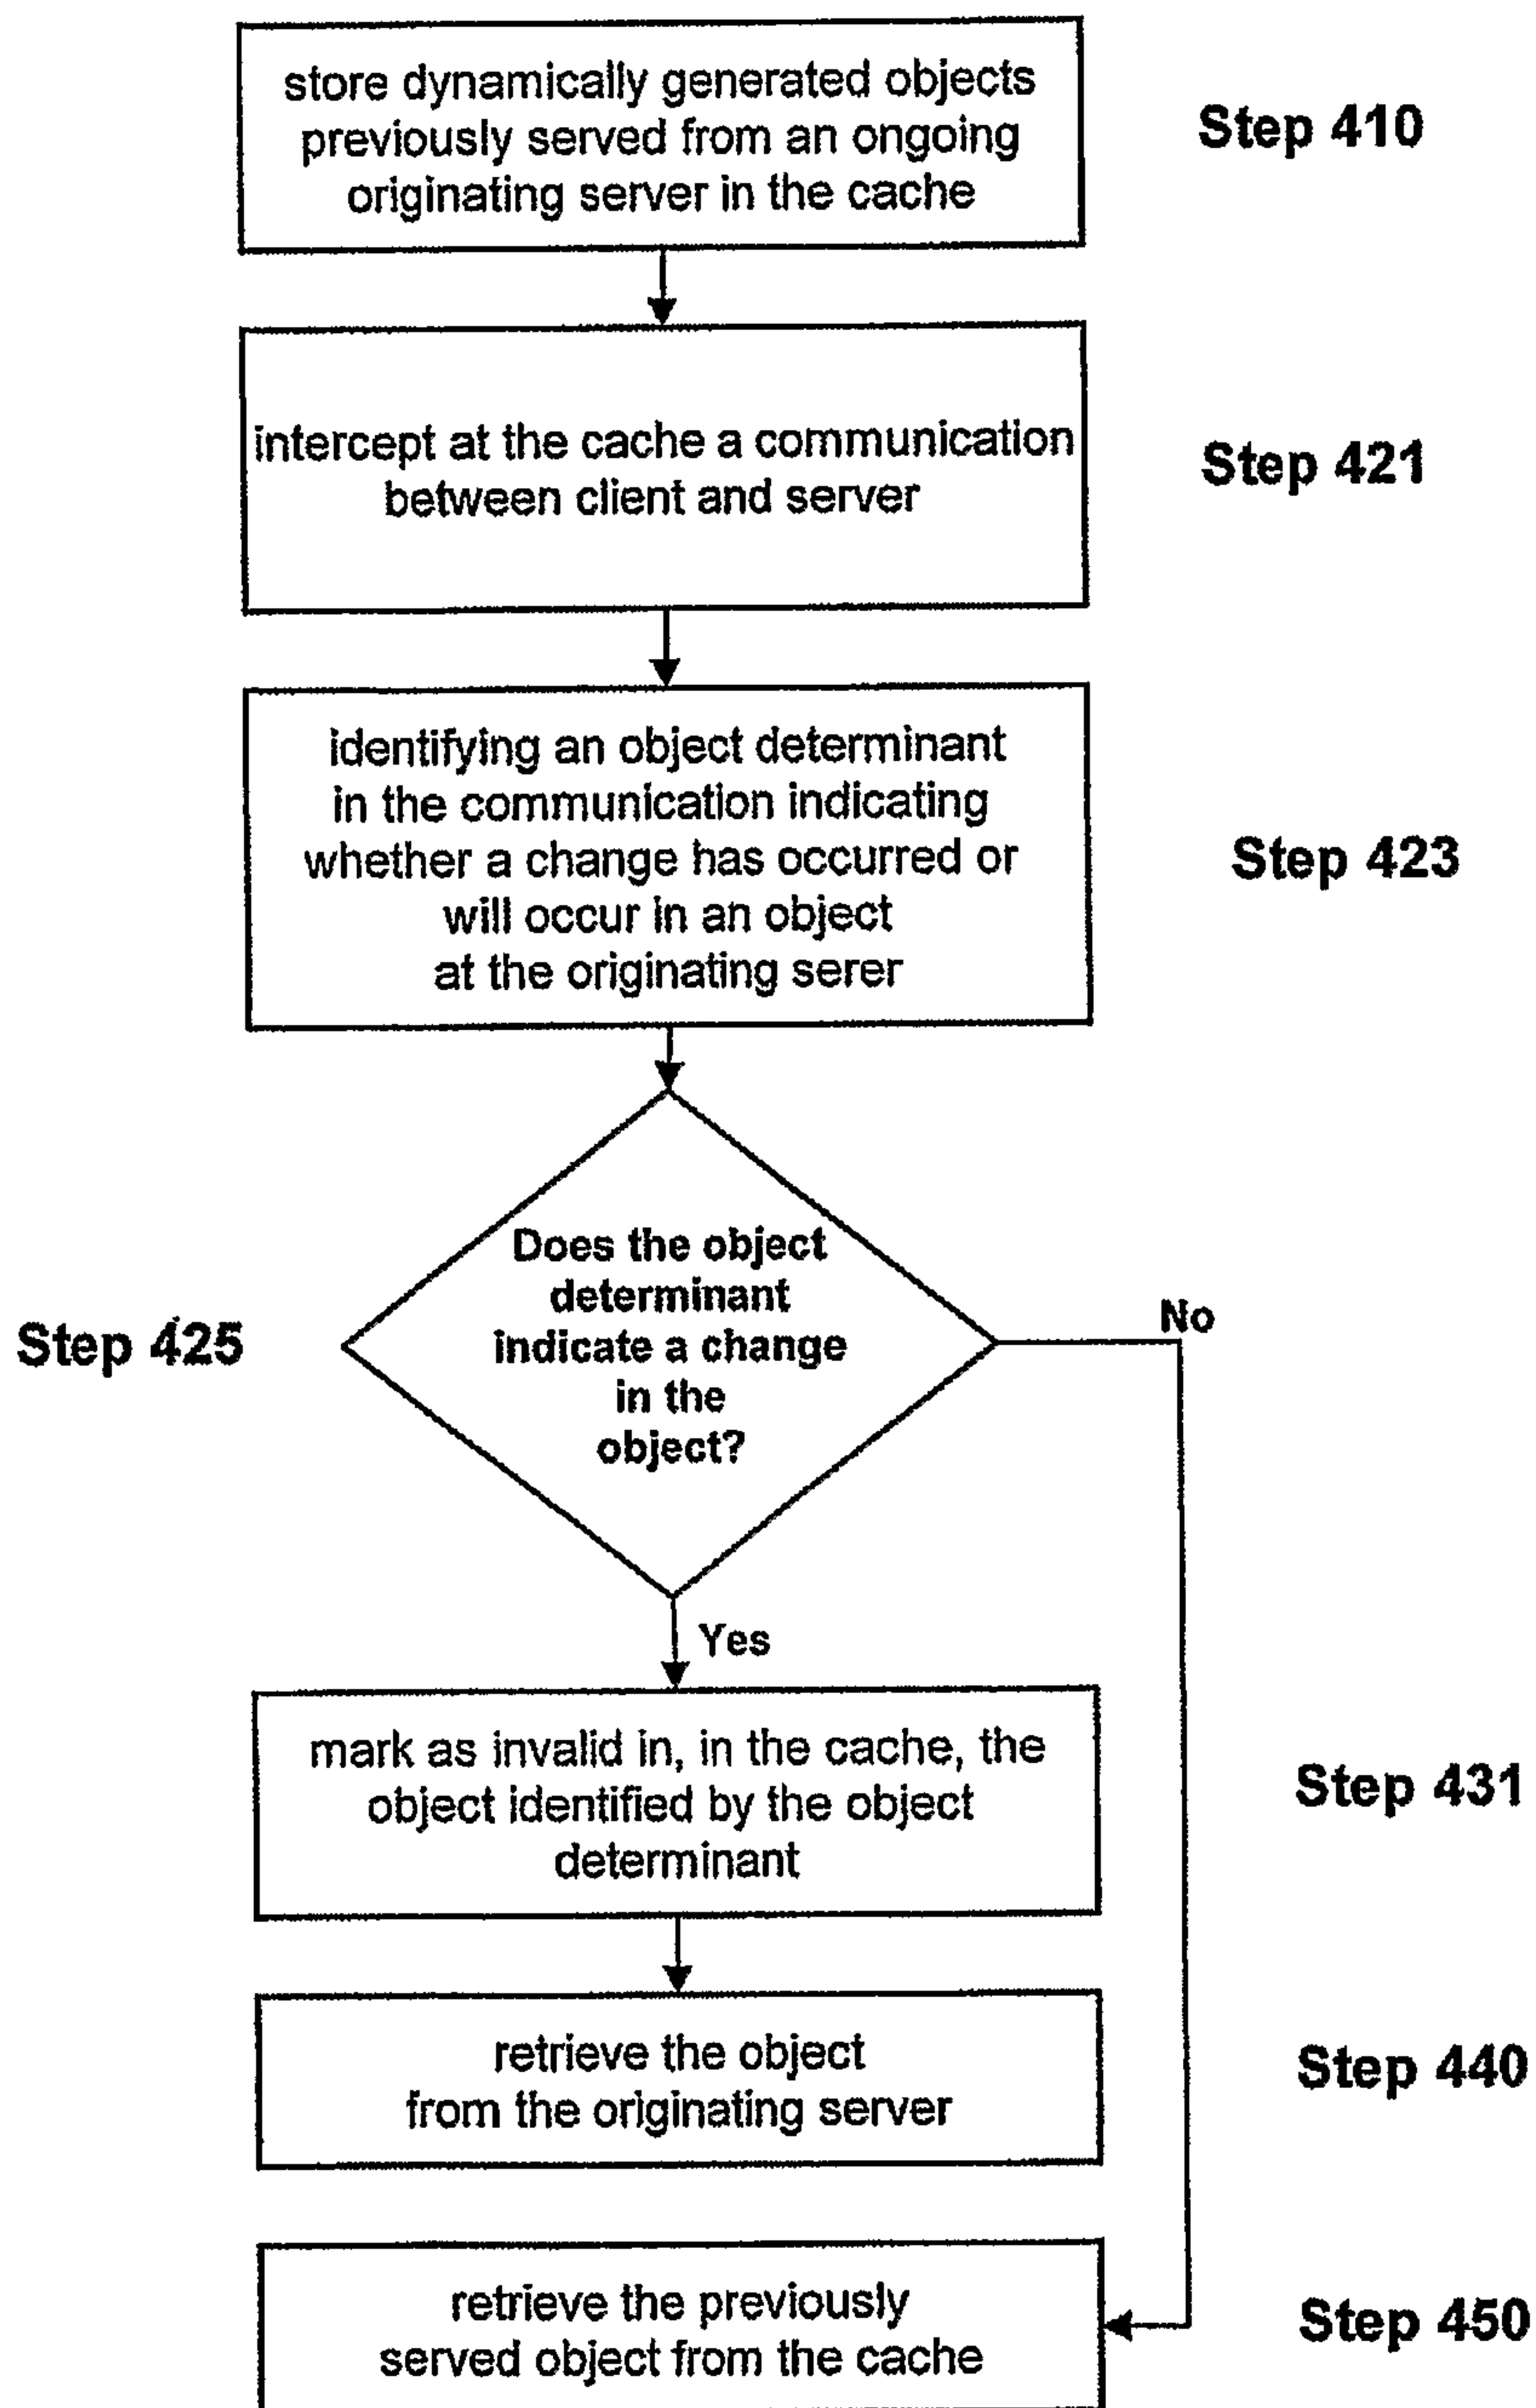


FIG. 4B

**FIG. 4C**

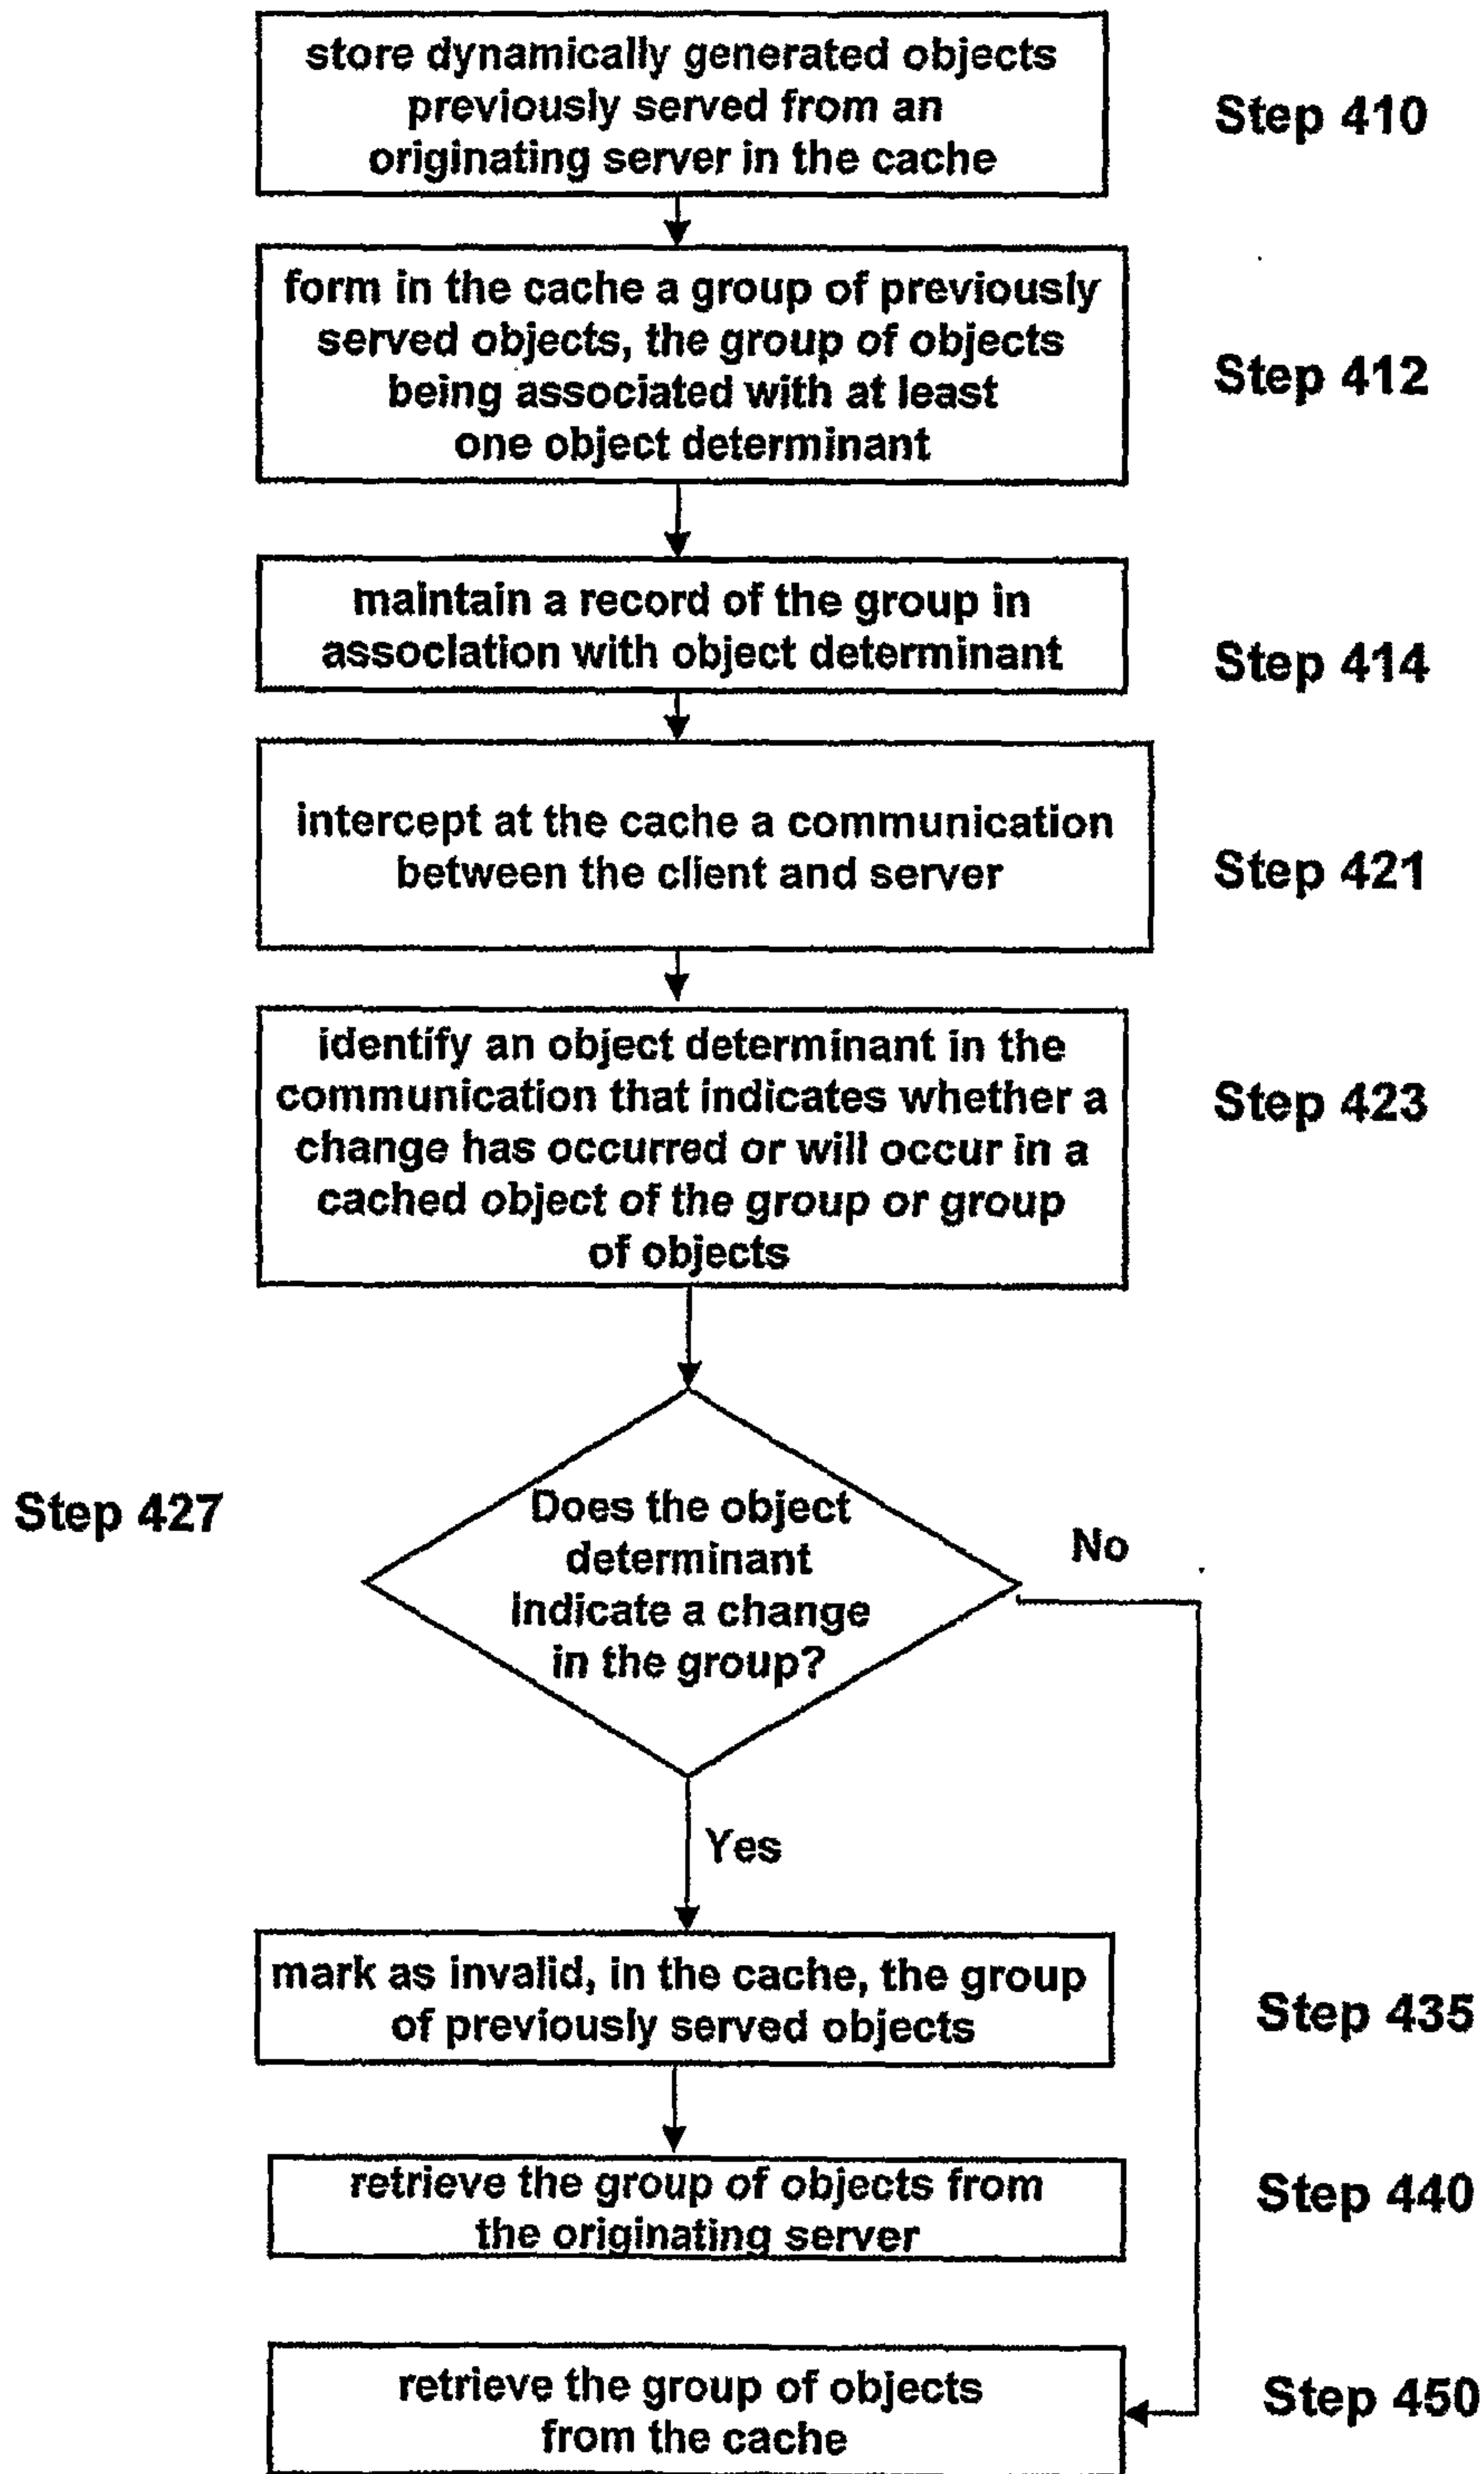
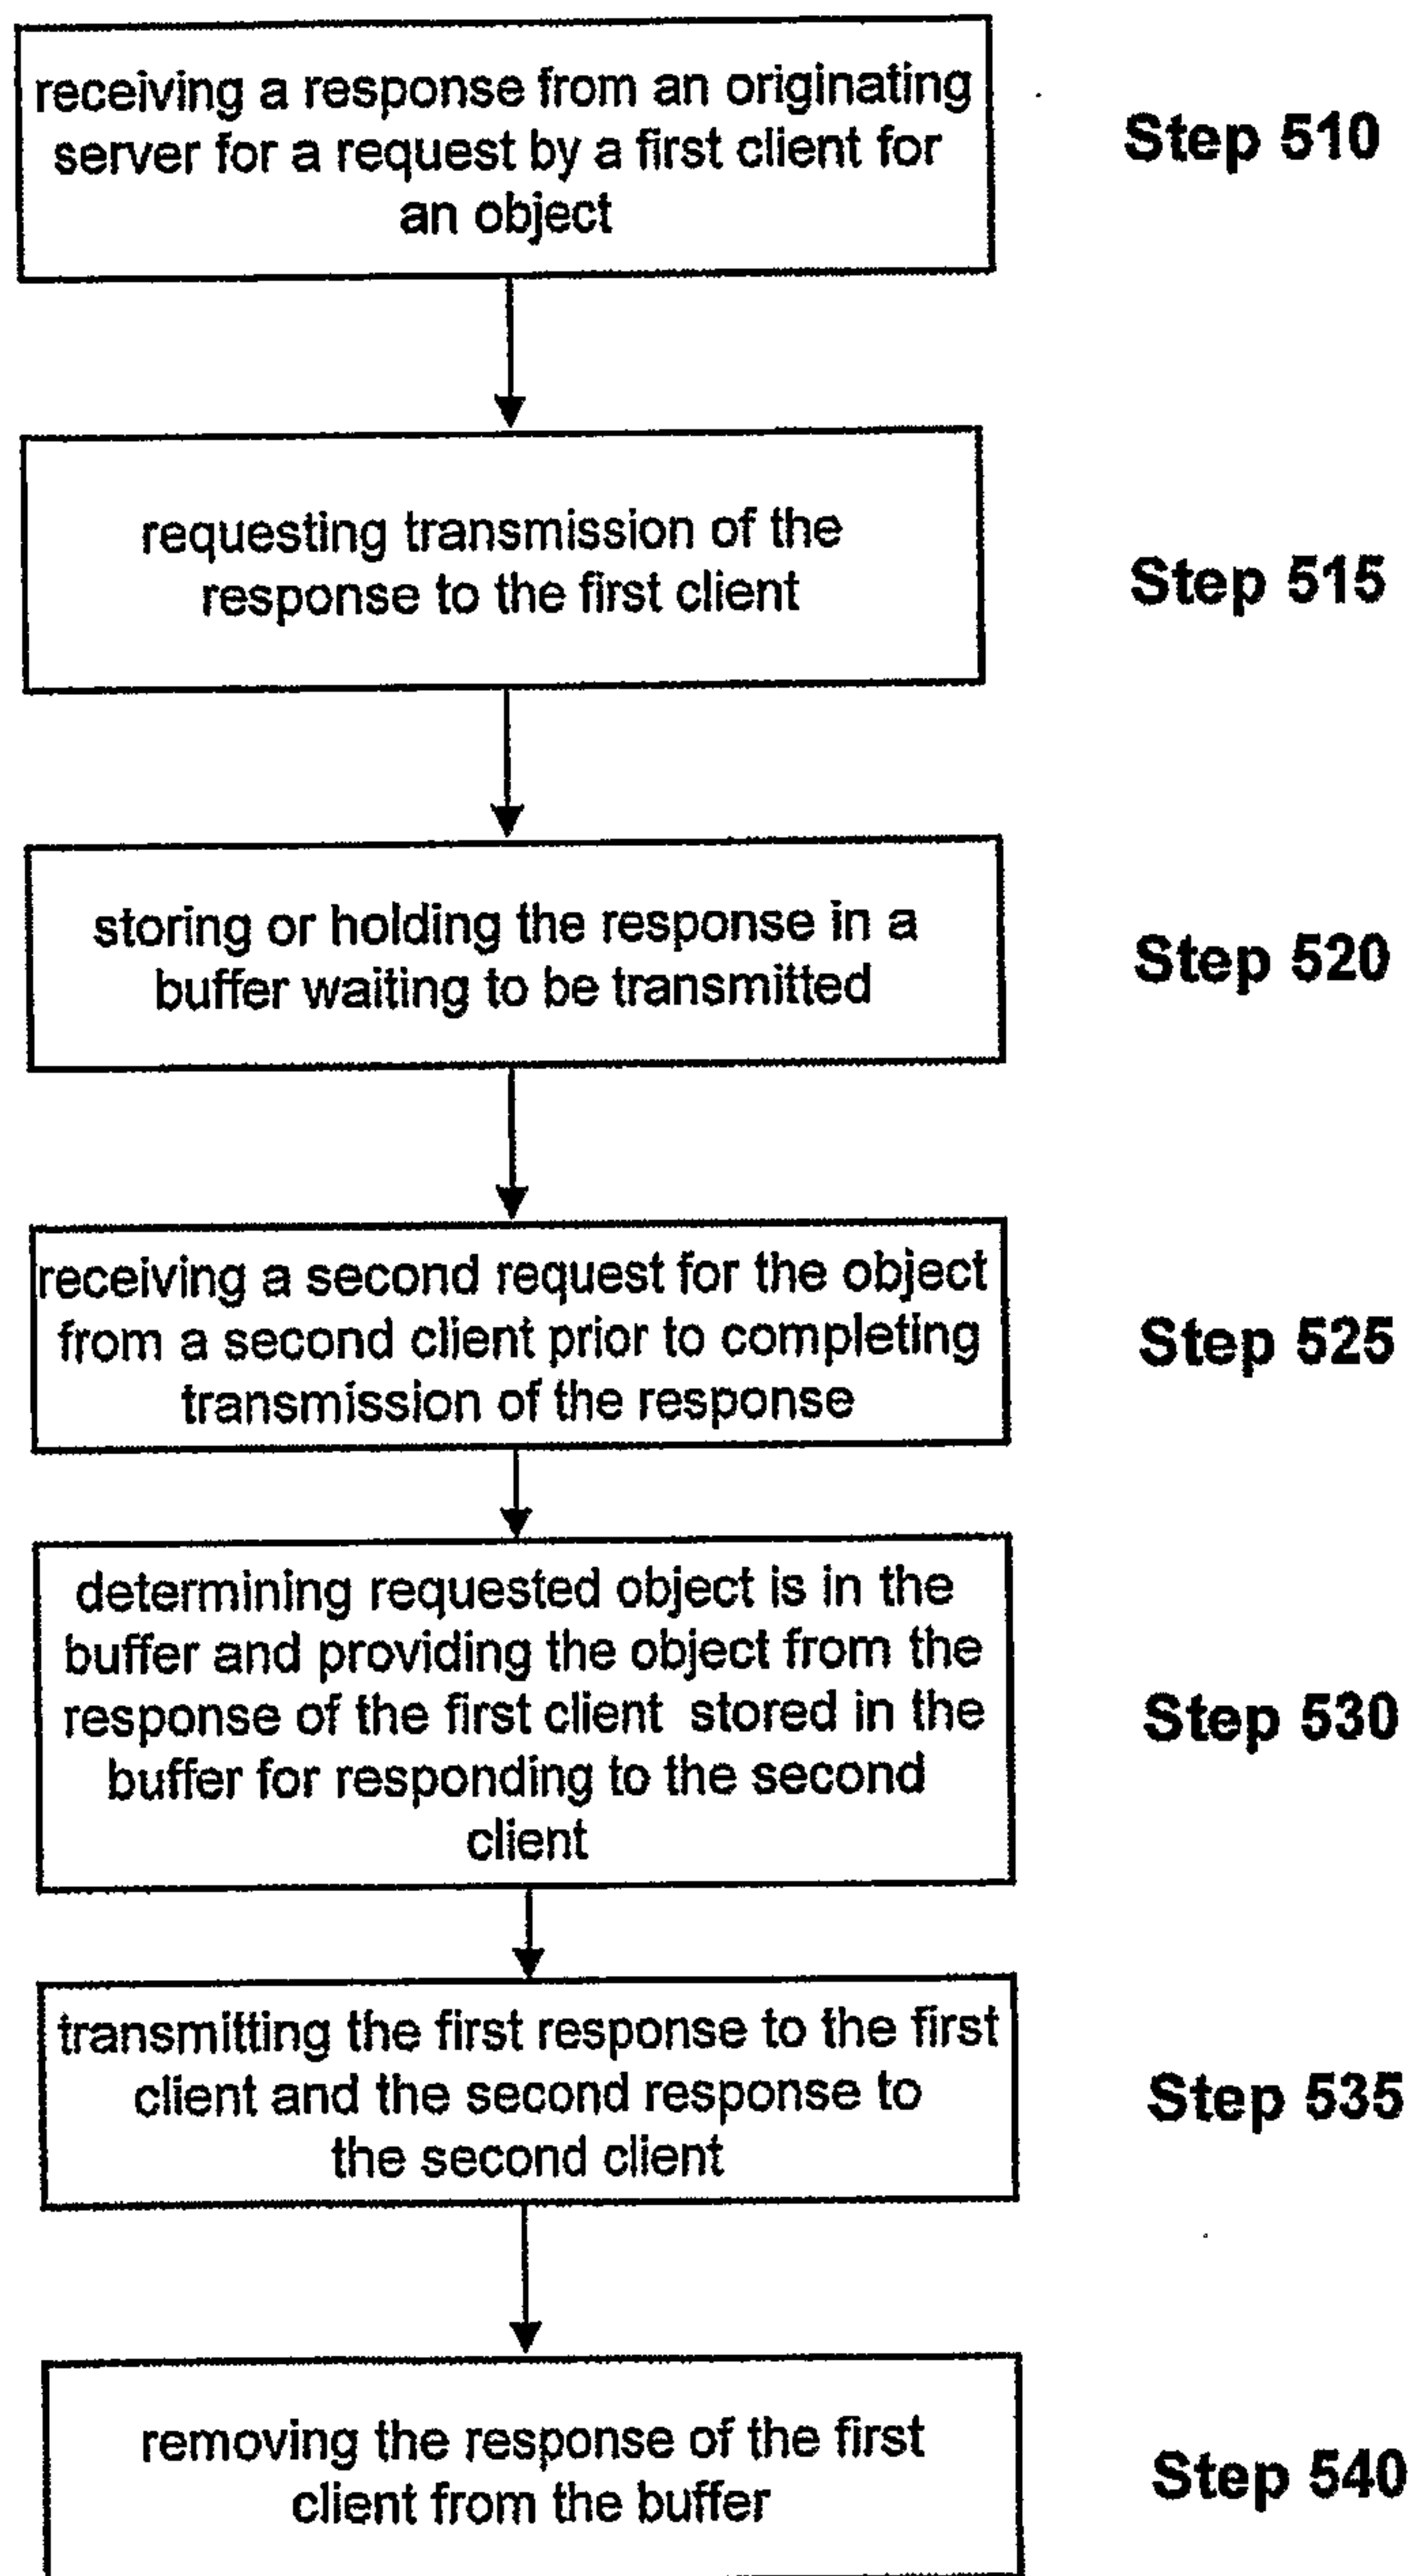
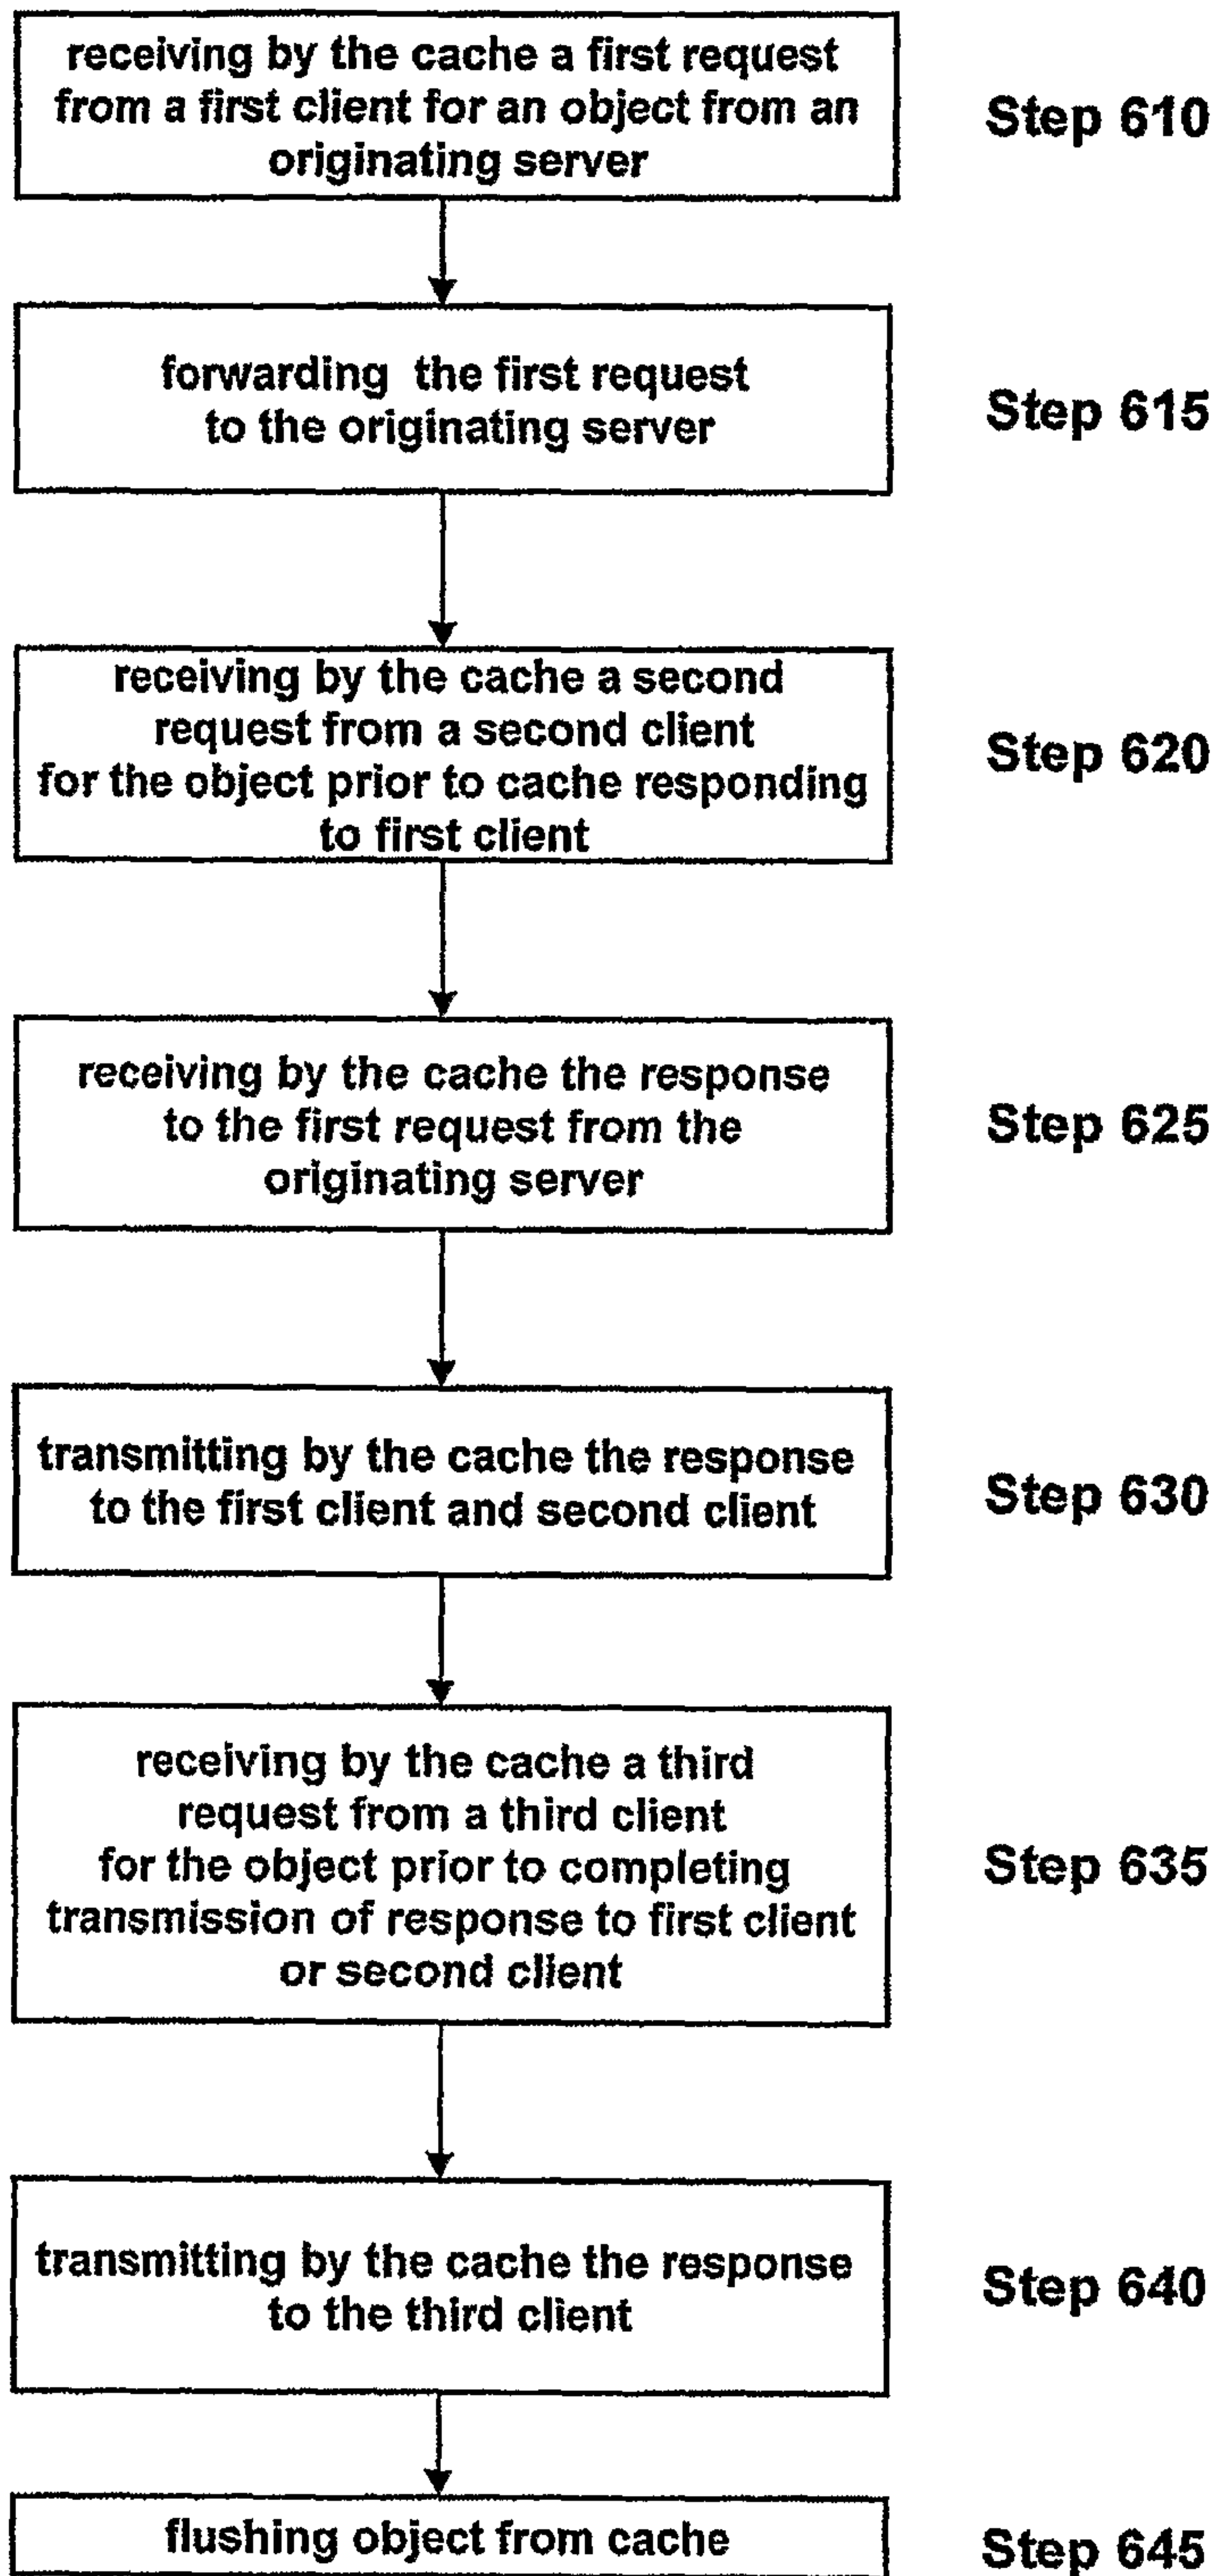


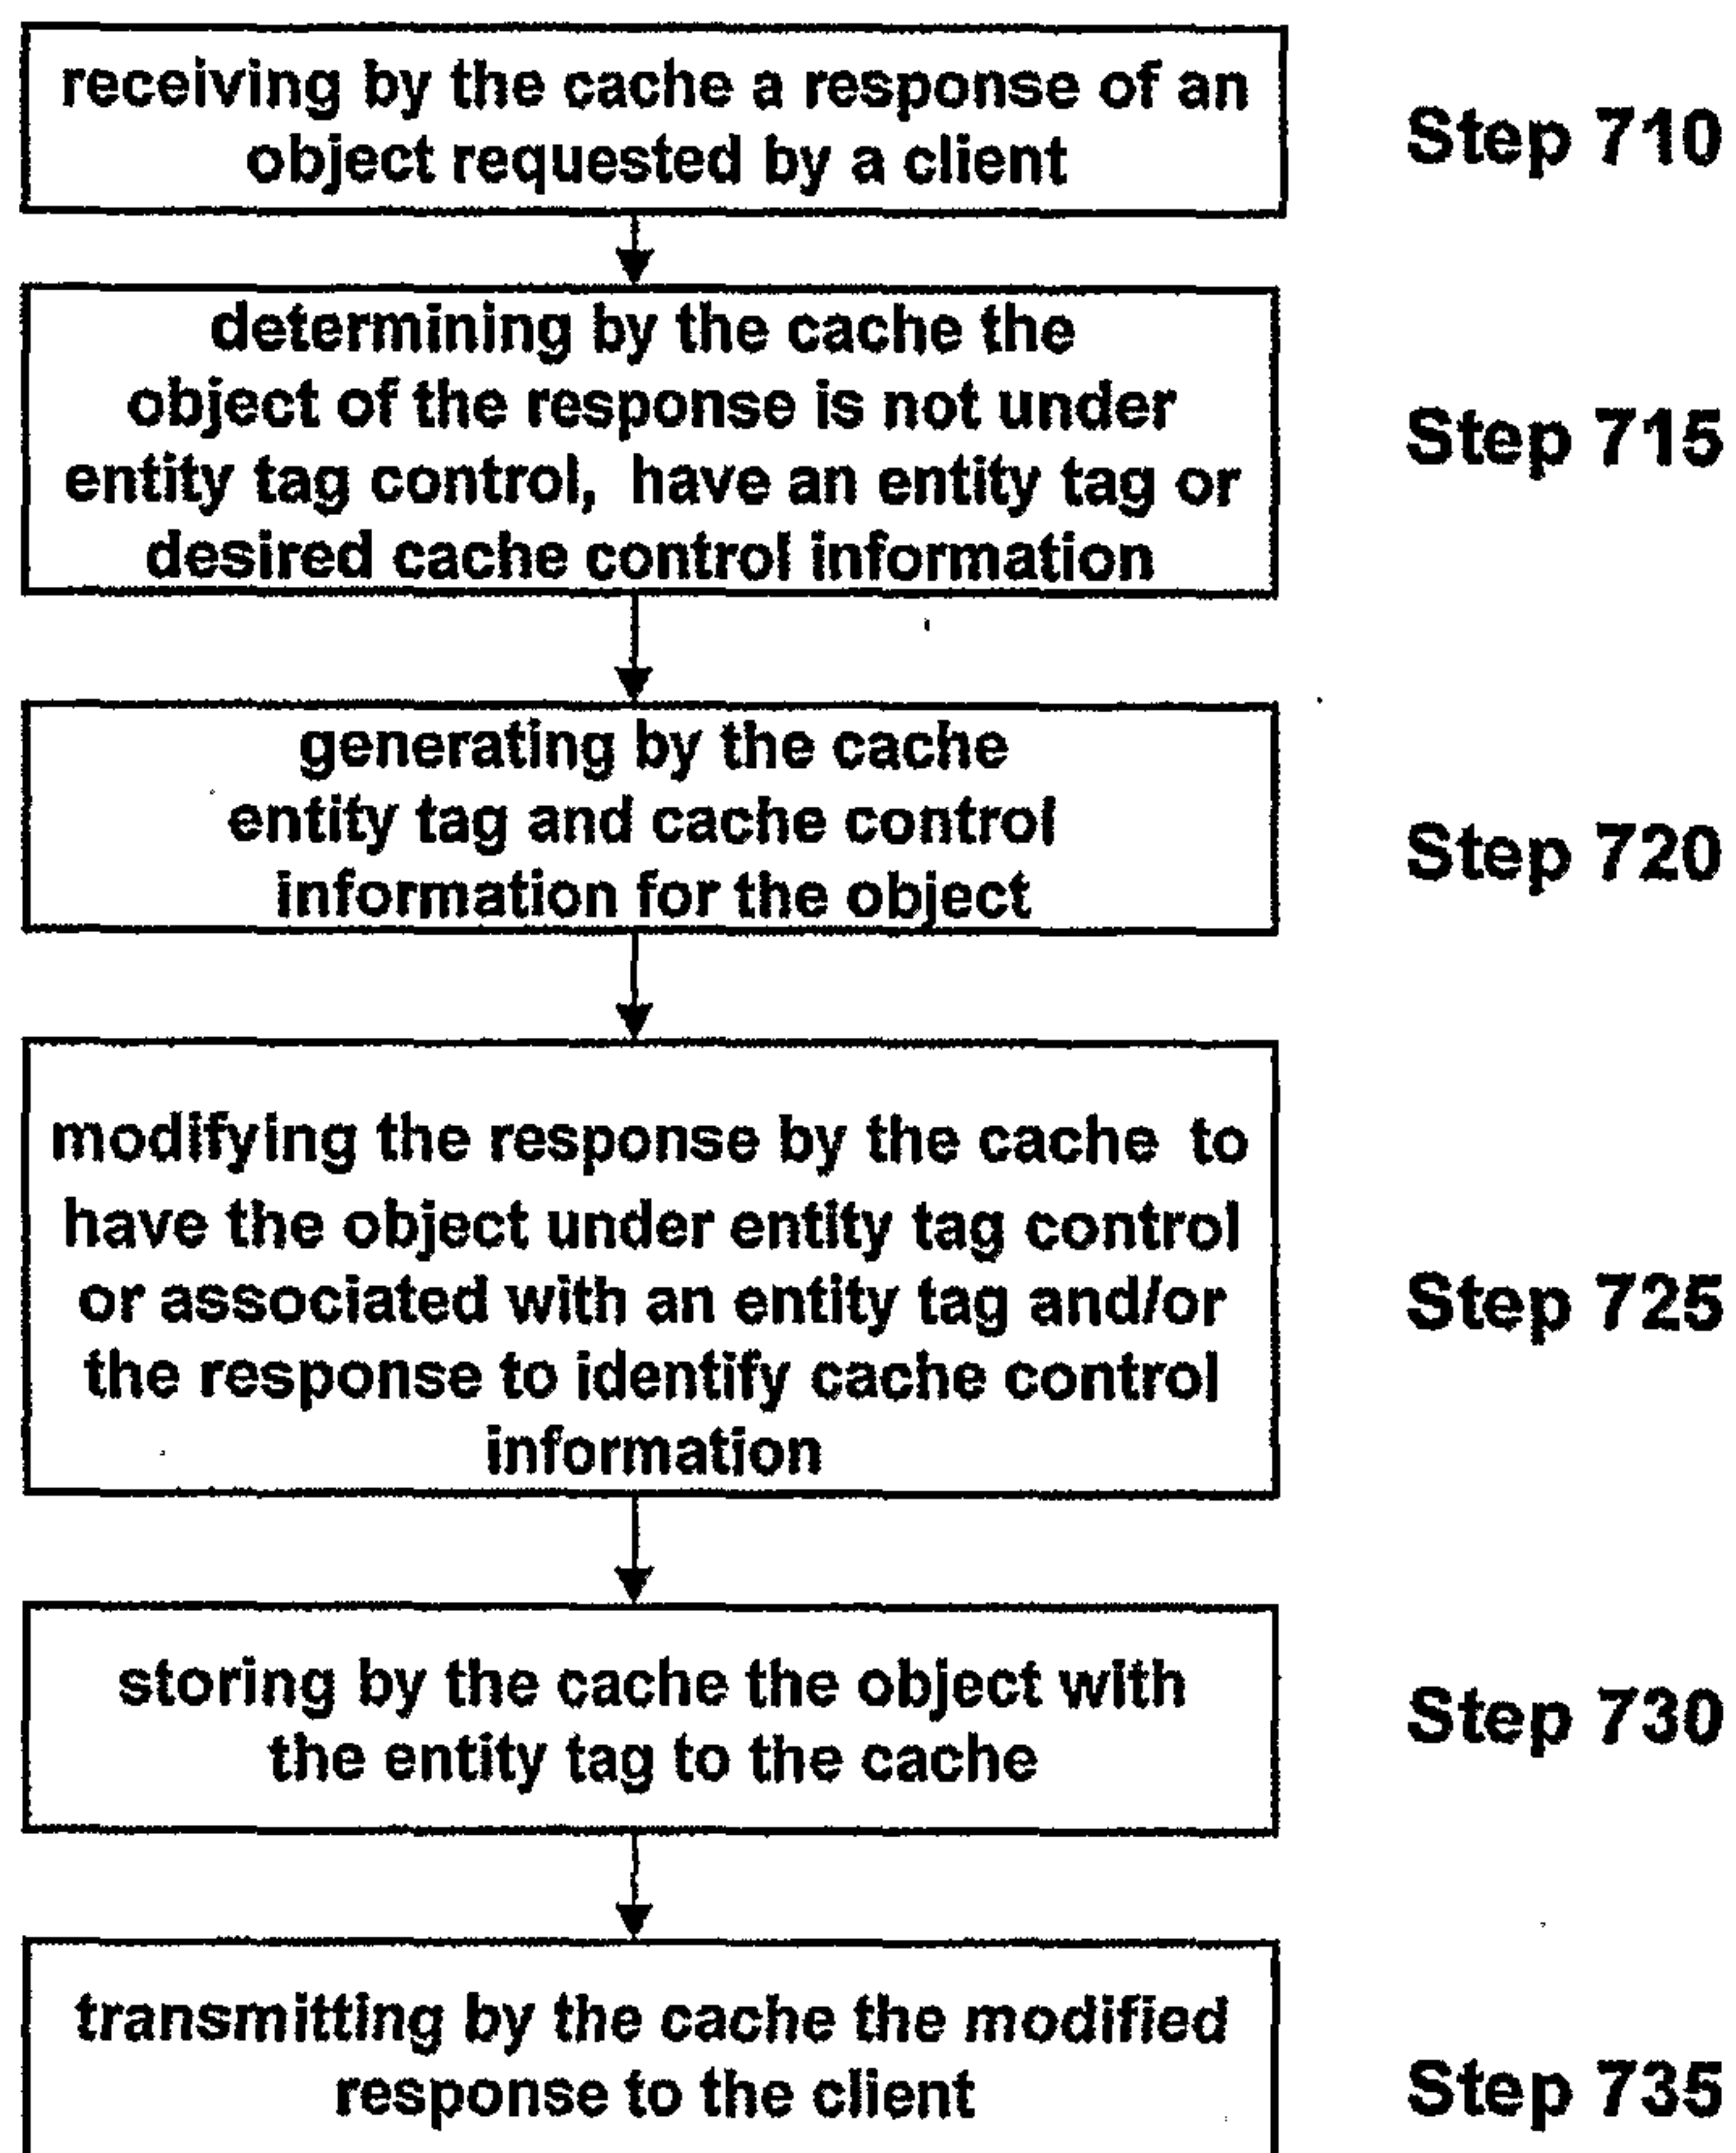
FIG. 4D

**FIG. 5**



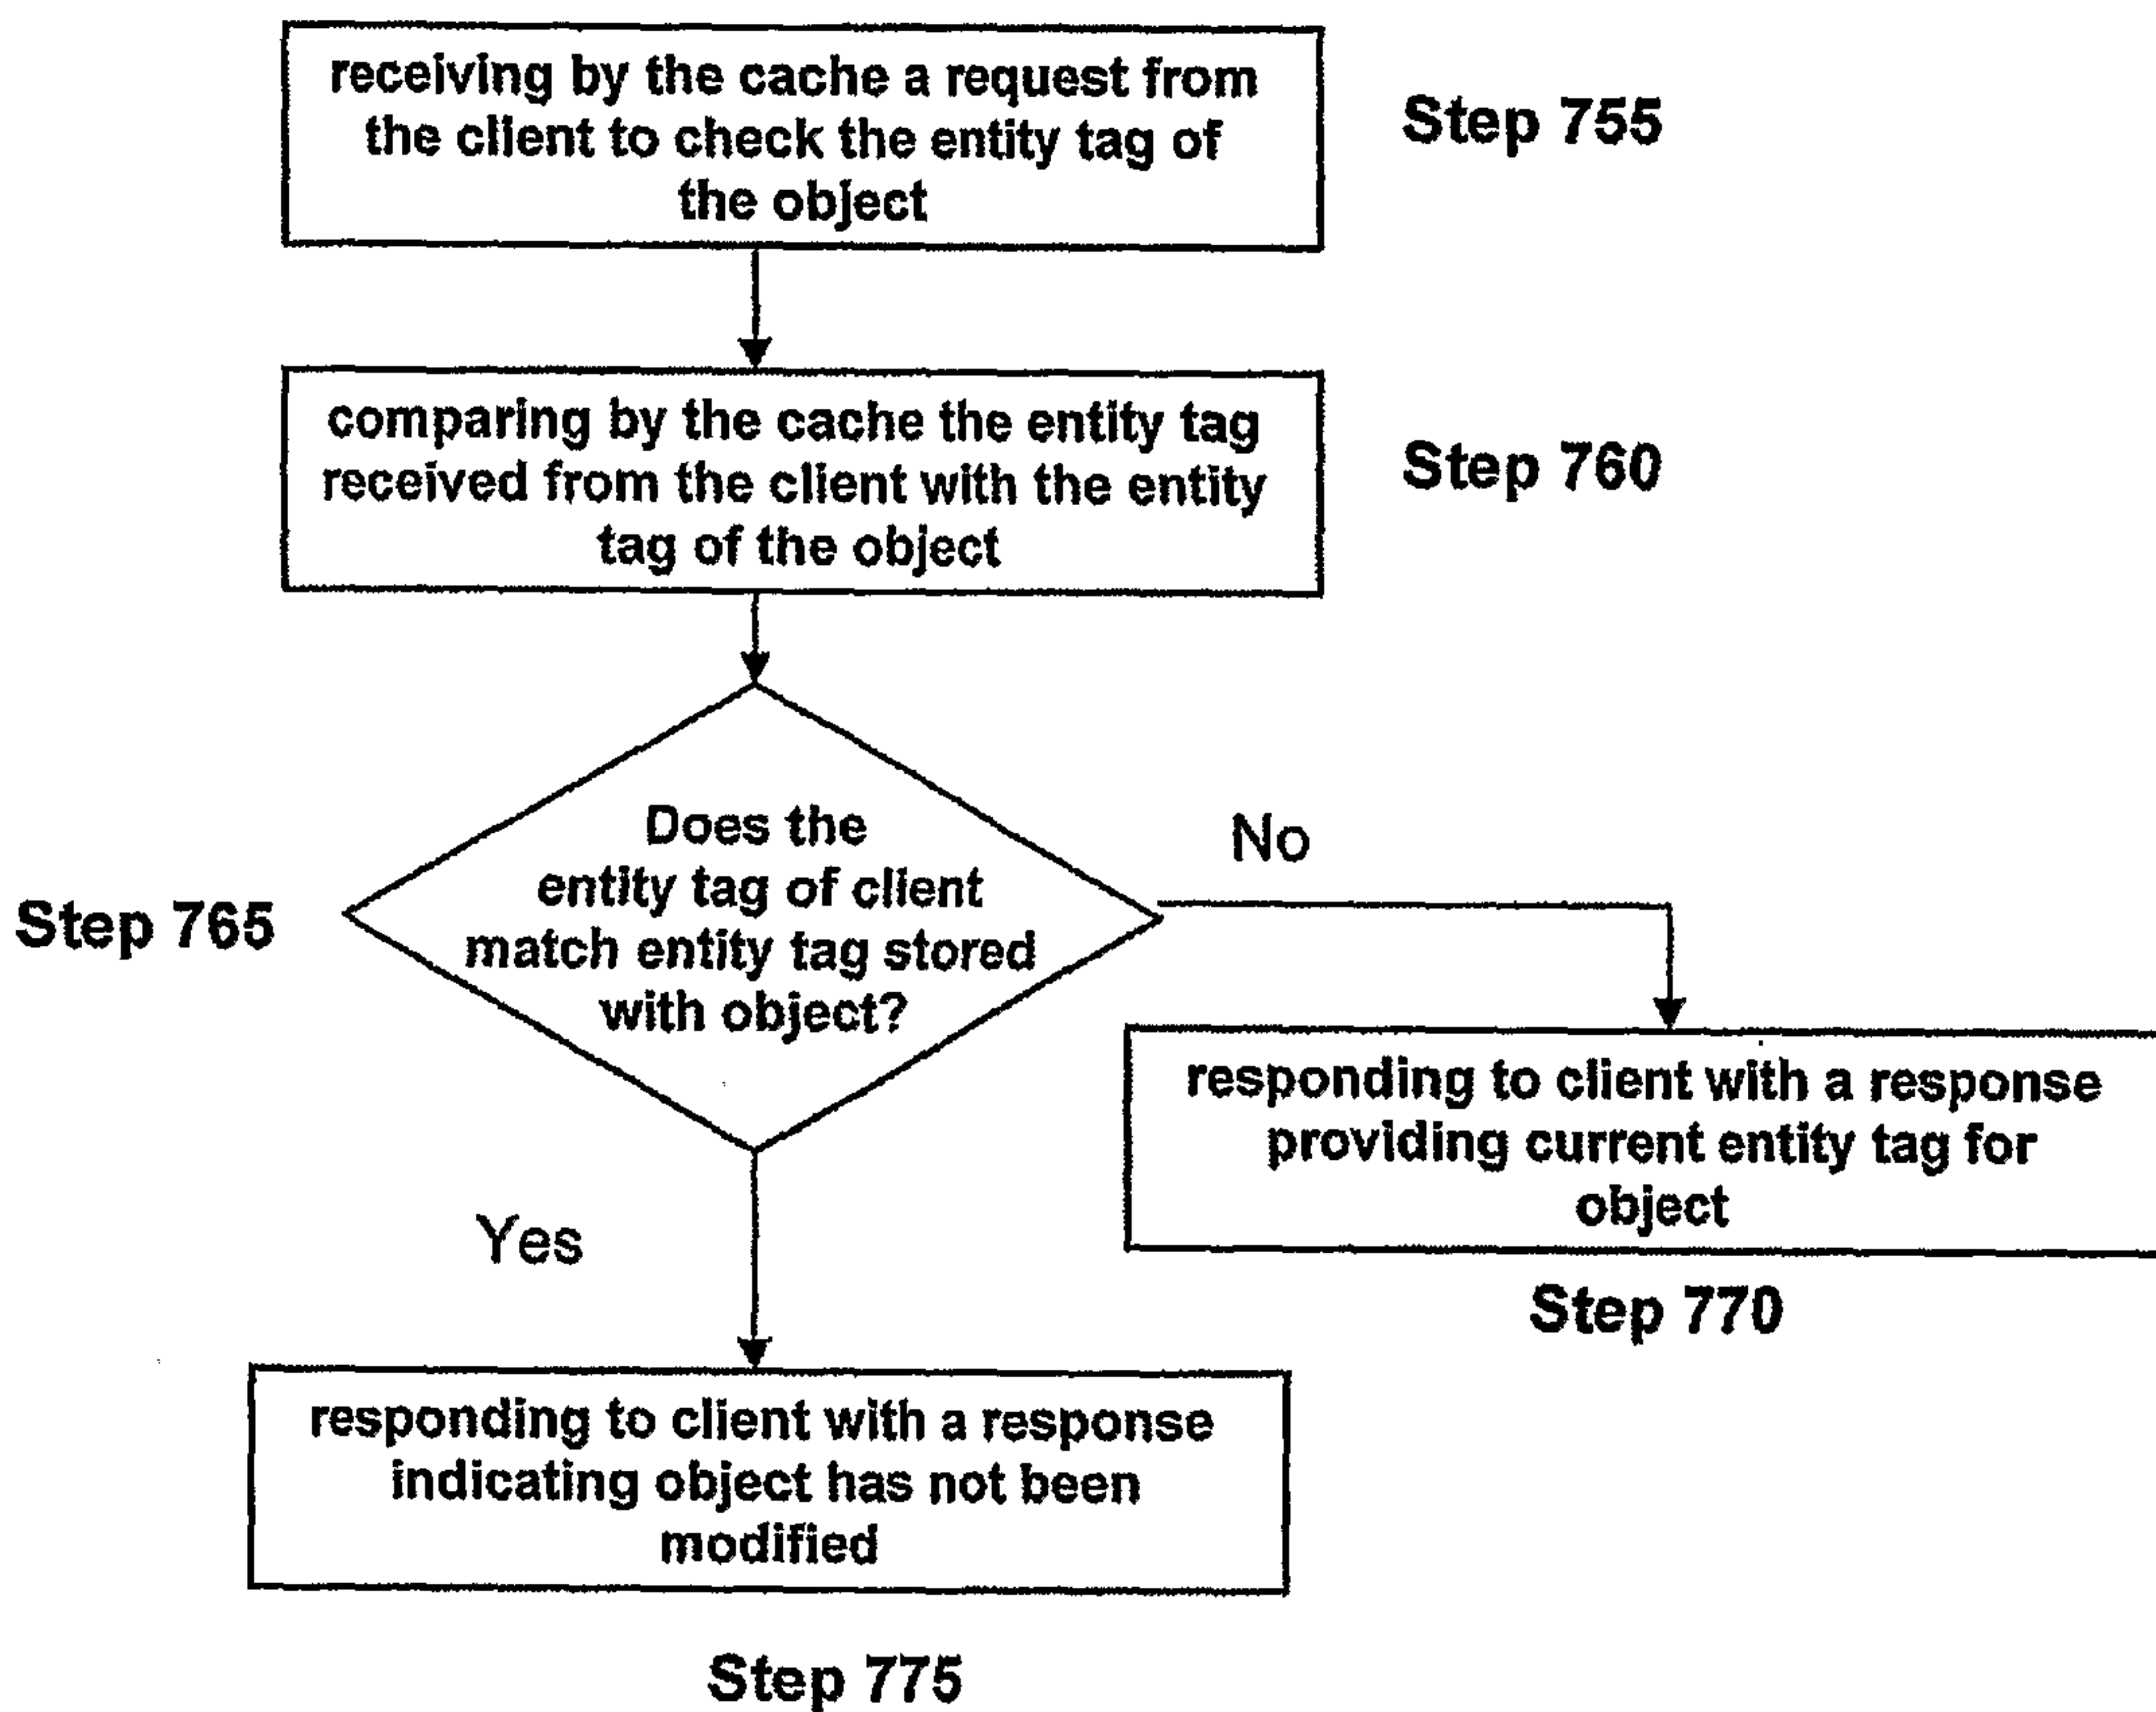
600

FIG. 6



700 ~

FIG. 7A



750 ~

FIG. 7B

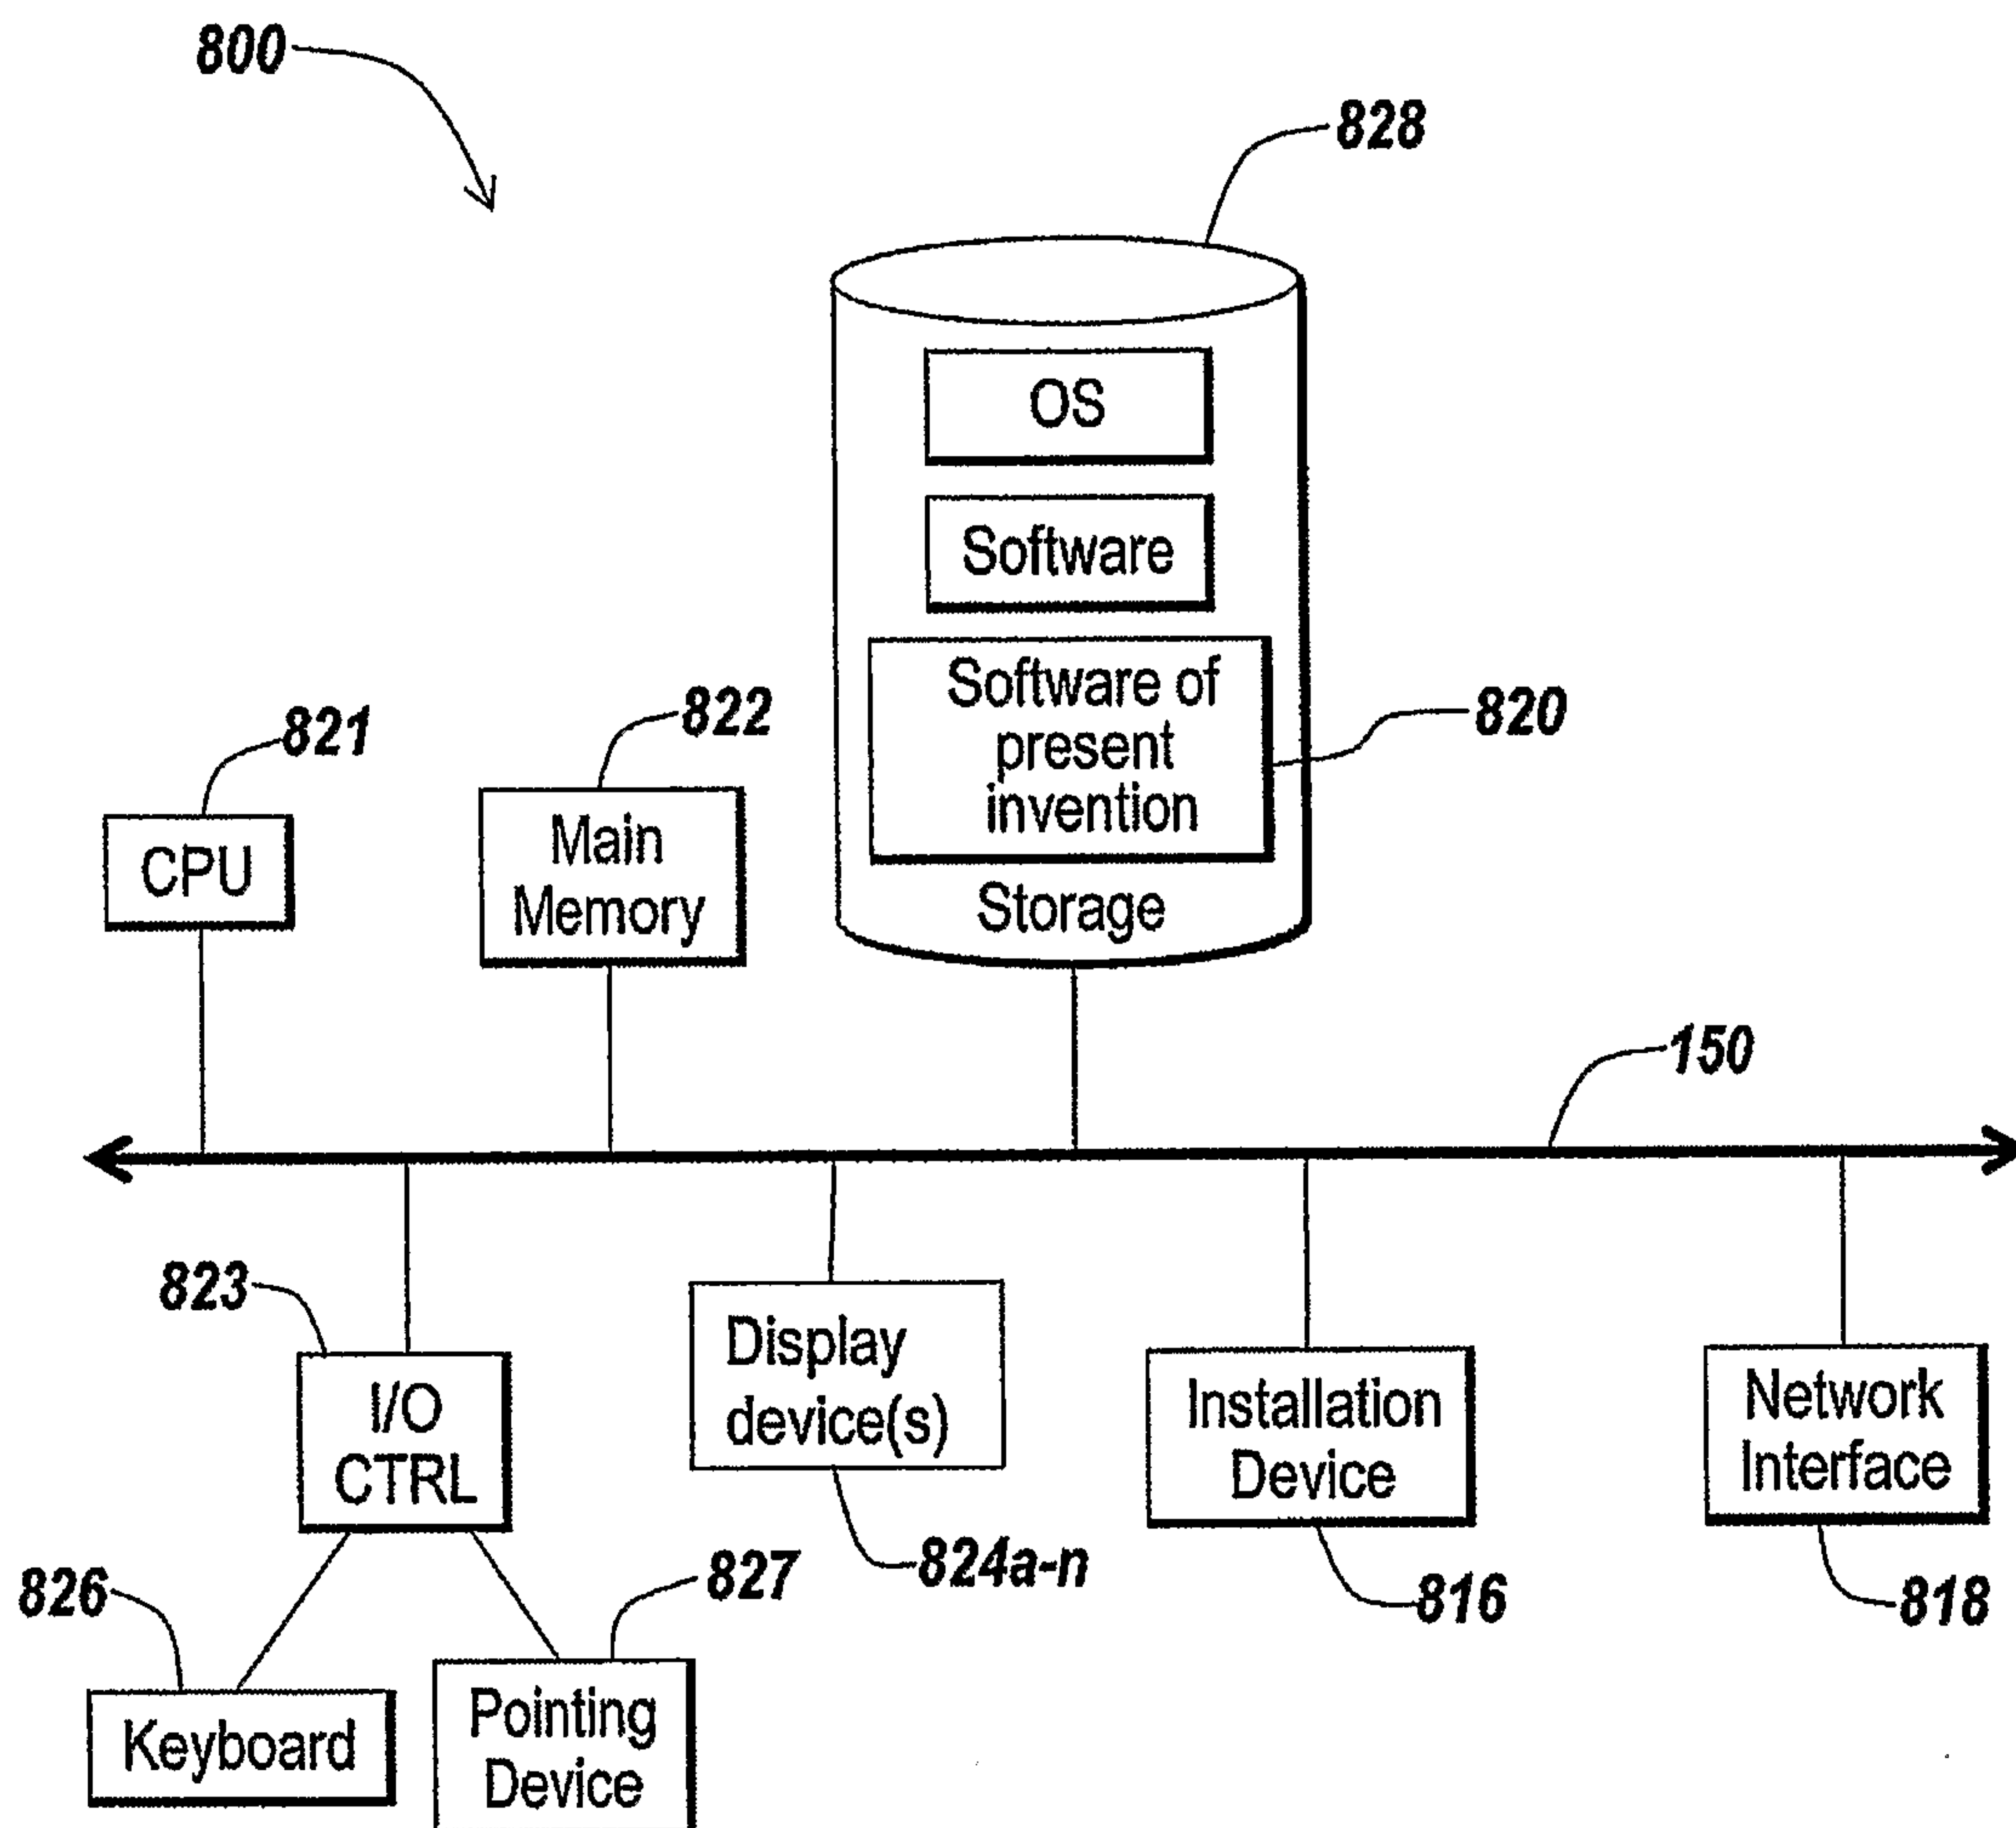


Fig. 8A


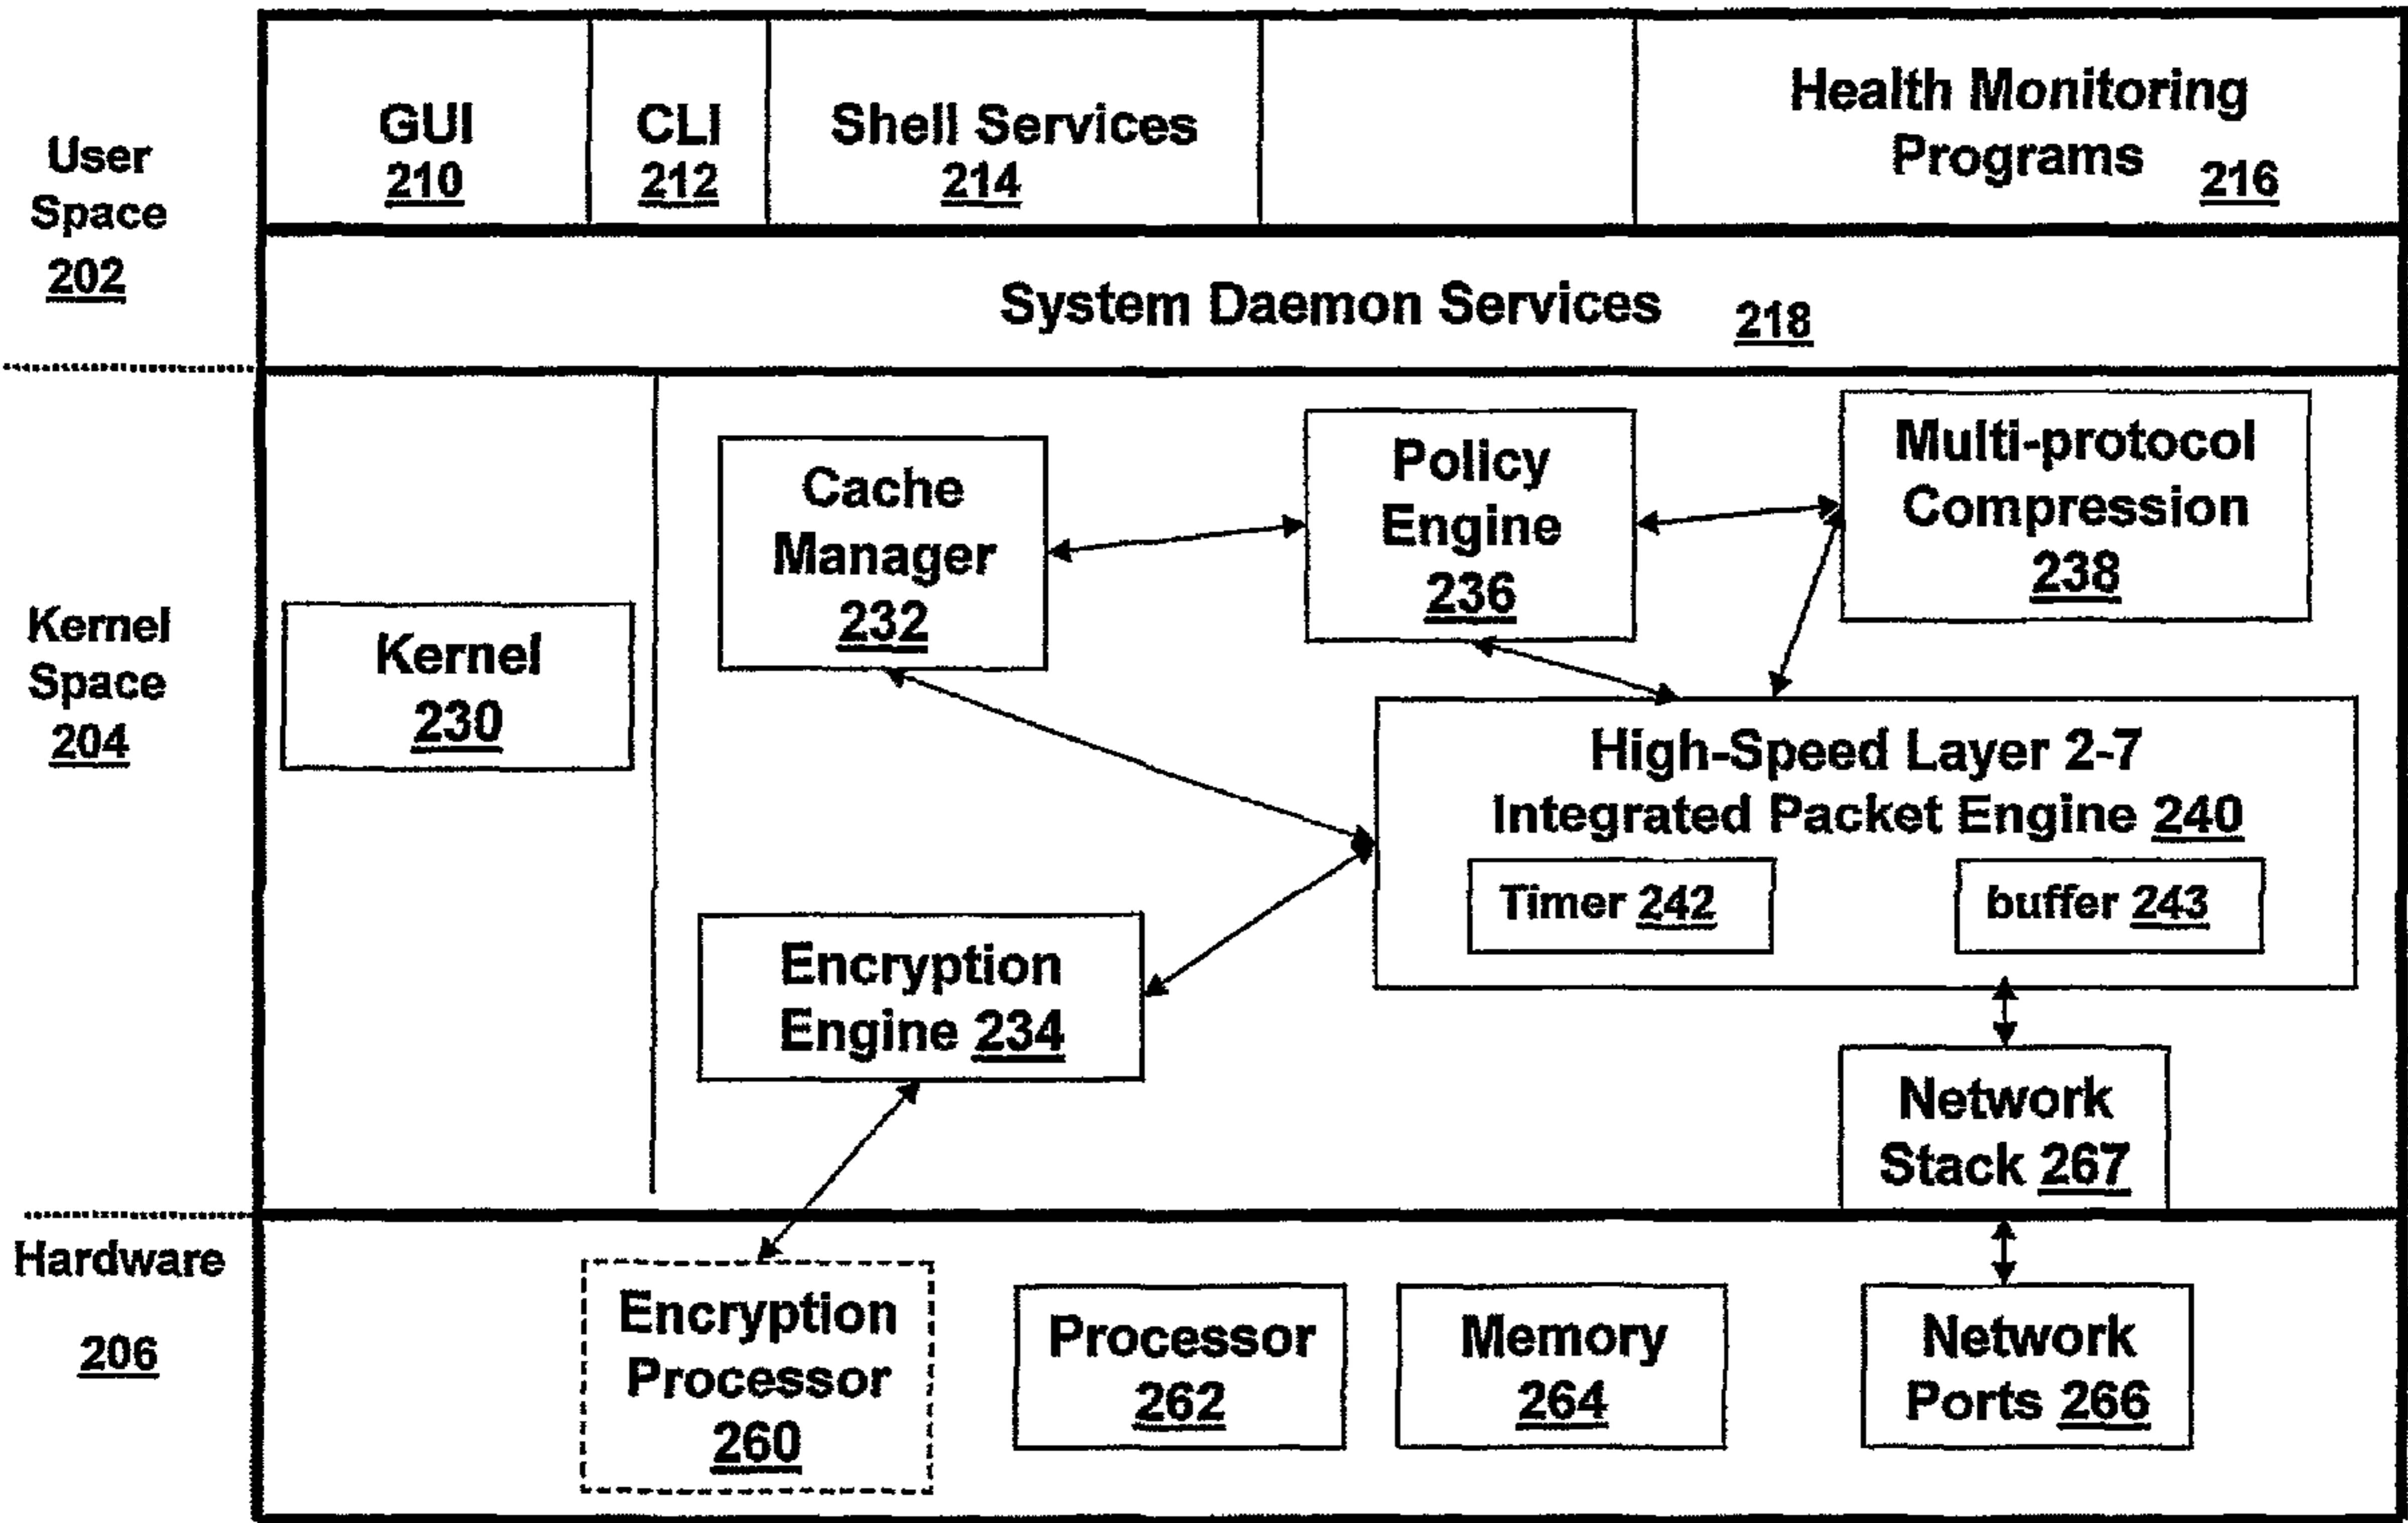


Fig. 8B



200