



(19) **United States**

(12) **Patent Application Publication**
Arolovitch et al.

(10) **Pub. No.: US 2012/0102148 A1**

(43) **Pub. Date: Apr. 26, 2012**

(54) **METHODS AND SYSTEMS FOR TRANSMISSION OF DATA OVER COMPUTER NETWORKS**

Publication Classification

(51) **Int. Cl.**
G06F 15/16 (2006.01)
(52) **U.S. Cl.** 709/217
(57) **ABSTRACT**

(75) Inventors: **Alan Arolovitch**, Brookline, MA (US); **Shmuel Bachar**, Herzliyya (IL); **Dror M. Gavish**, Shoham (IL); **Shahar G. Grin**, Ramat HaSheron (IL); **Shay Shemer**, Hod Hasharon (IL)

(73) Assignee: **PeerApp Ltd.**, Newton Upper Falls, MA (US)

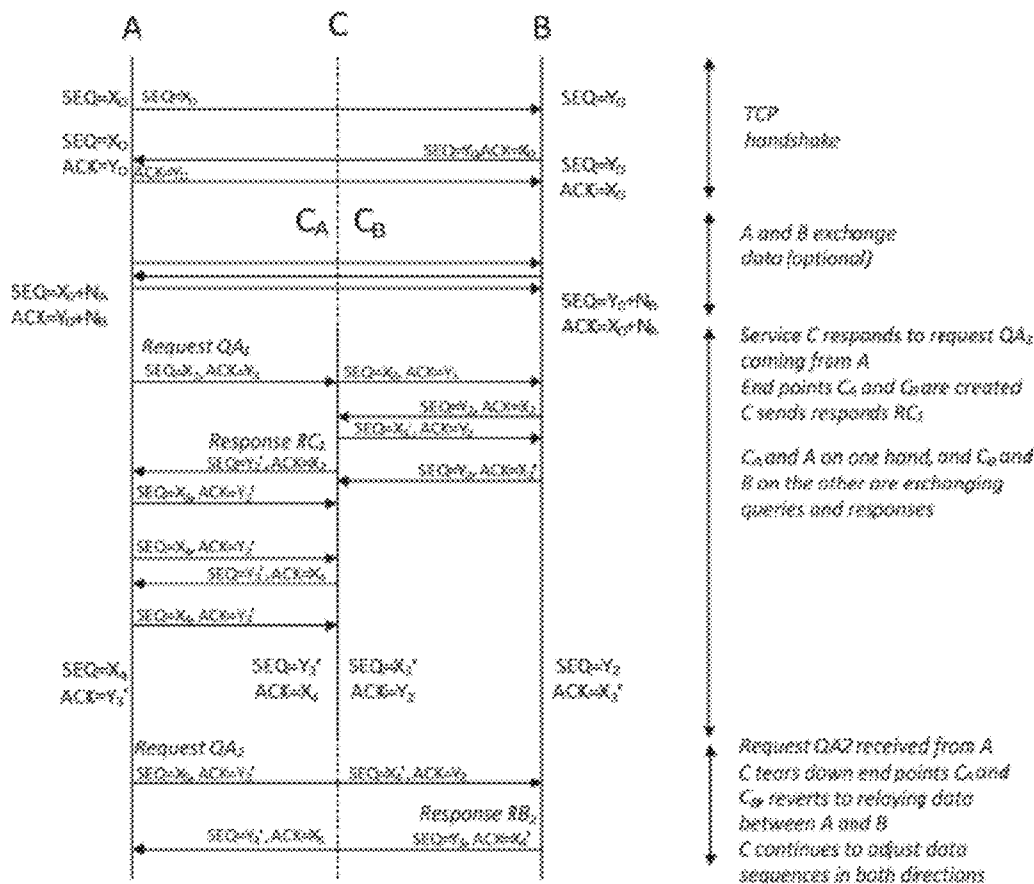
(21) Appl. No.: **13/341,409**


(22) Filed: **Dec. 30, 2011**

Related U.S. Application Data

(60) Provisional application No. 61/428,527, filed on Dec. 30, 2010.

A computer-implemented method is provided for transparently optimizing data transmission between a first endpoint and a second endpoint in a computer network. The endpoints have a directly established data session therebetween. The data session is identified by each endpoint at least to itself in the same way throughout the session. The method includes the steps of: relaying data between the endpoints transparently in the session using a network optimization service; and transparently modifying or storing at least some of the data transmitted from the second endpoint to the first endpoint using the network optimization service in order to optimize data communications between the endpoints, wherein transparently modifying at least some of the data comprises changing the data, replacing the data, or inserting additional data such that the first endpoint receives different data than was sent by the second endpoint.





A and B establish session between themselves

A and B establish session between themselves

FIG. 1

FIG. 2A

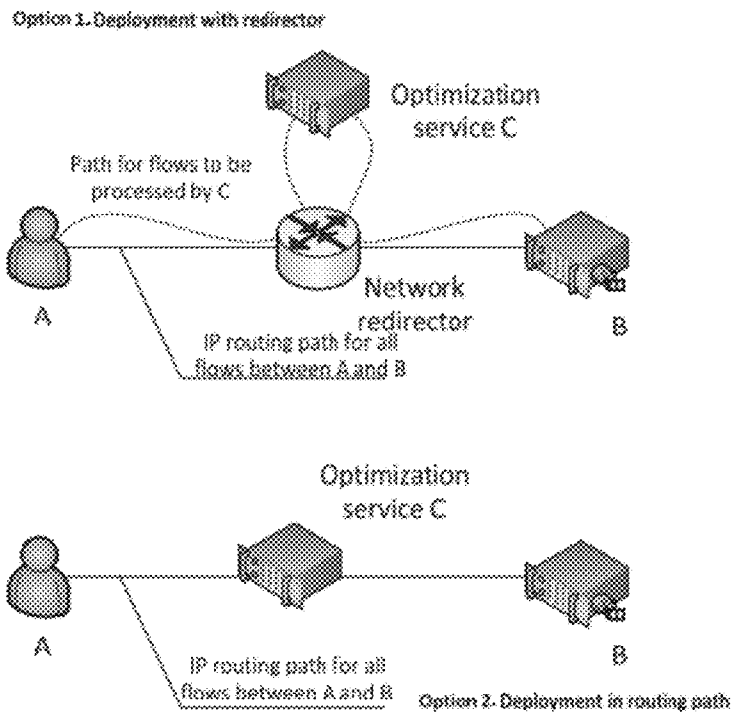


FIG. 2B




FIG. 3

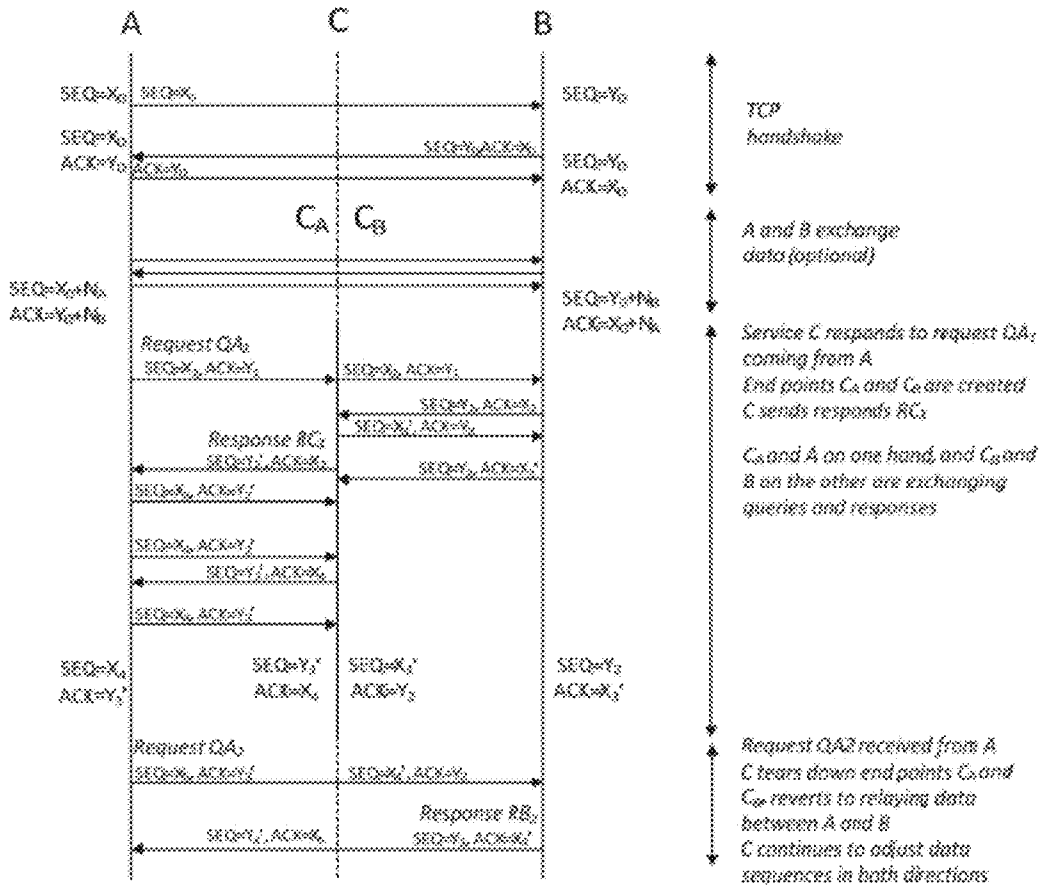


FIG. 4

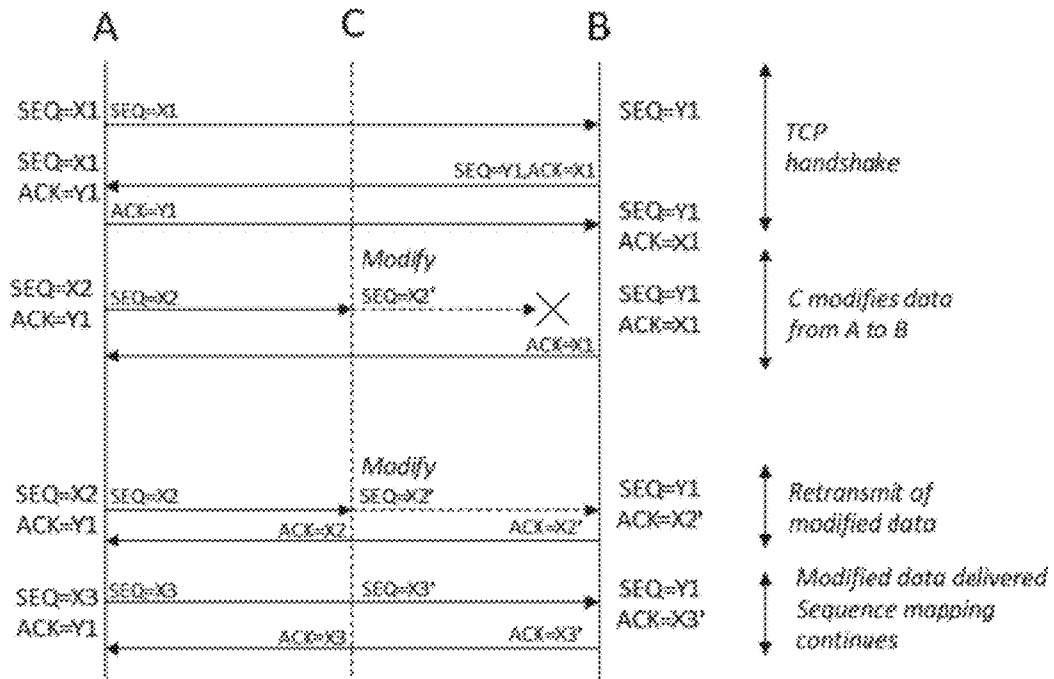


FIG. 5

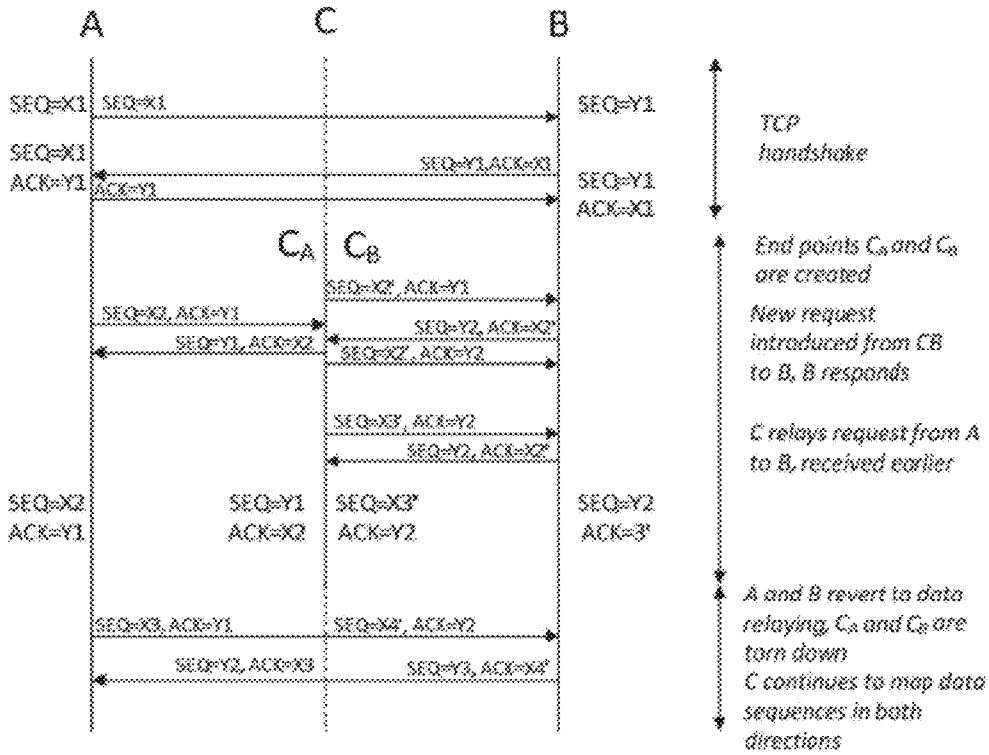


FIG. 6

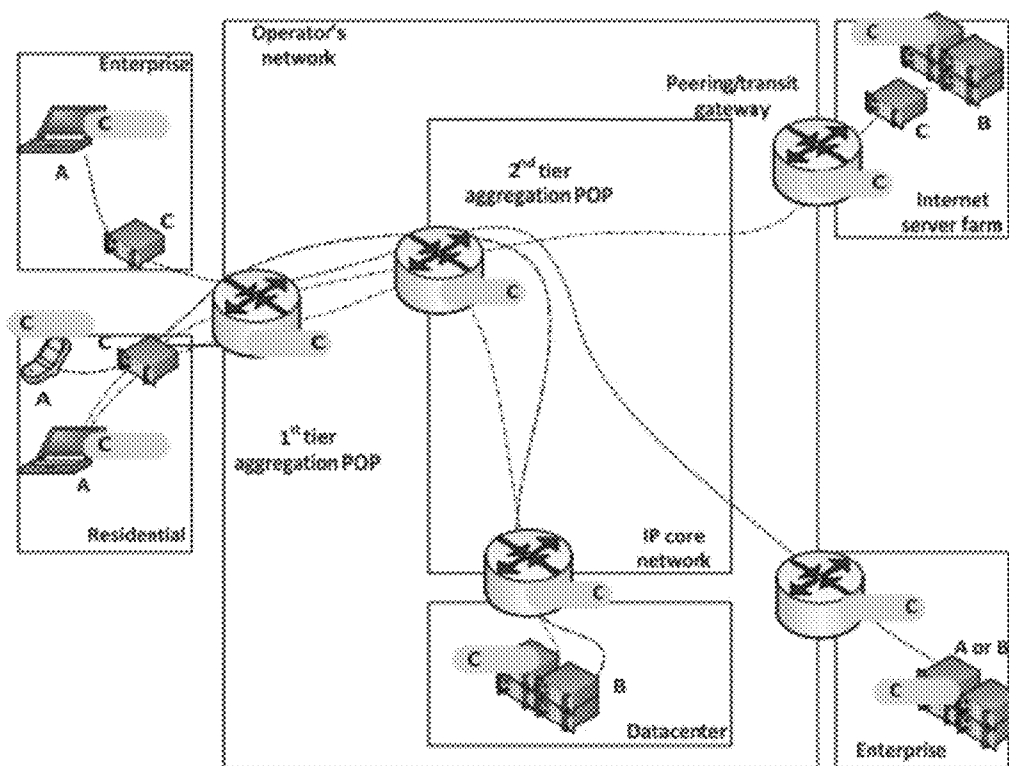


FIG. 7

METHODS AND SYSTEMS FOR TRANSMISSION OF DATA OVER COMPUTER NETWORKS

CROSS REFERENCE TO RELATED APPLICATION

[0001] This application claims priority from U.S. Provisional Patent Application No. 61/428,527, filed on Dec. 30, 2010, entitled METHODS AND SYSTEMS FOR TRANSMISSION OF DATA OVER COMPUTER NETWORKS, which is hereby incorporated by reference.

BACKGROUND

[0002] The present application relates generally to transmissions of data over a computer network such as, e.g., the Internet, a local area network, a wide area network, a wireless network, and others.

[0003] Both enterprise and consumer broadband networks have undergone significant and continuous growth of traffic volumes for the last 3-5 years. The traffic growth is driven by introduction of faster end user connectivity options, adoption of various bandwidth-intensive applications, and introduction of various Internet-connected consumer electronics products.

[0004] To respond to network congestion, degradation of application performance, and, the need to continuously upgrade their networks, caused by the broadband growth, broadband network operators have introduced various network optimization solutions and services aimed at controlling their network costs, containing growth of network scale, improving performance and security of Internet applications, and creating new revenue sources for the operators.

[0005] Such solutions include content caching, video transcoding and transrating, content adaptation, content filtering, intrusion detection and prevention, among others.

[0006] All these solution classes share several common deployment requirements. They should be deployed in a transparent way, so that Internet applications may operate without change.

[0007] It is also common for some network optimization solutions to modify the Internet content flow and/or content payload itself.

[0008] Furthermore, the network optimization solutions should address the scale requirements of modern broadband networks that frequently operate on 10 Gbps, 40 Gbps and 100 Gbps scale.

[0009] Common solution architecture for network-based optimization involves a network optimization platform deployed in conjunction with a network element (e.g., routing, switching or dedicated DPI equipment) that sits in data path and redirects traffic to the network optimization platform.

[0010] Network elements typically employ selective redirection of network traffic, matching types of traffic flows to the network optimization service used.

[0011] Network optimization services commonly use application proxy architecture. A connection that otherwise would be established between two endpoints 'A' and 'B' (e.g., an Internet browser and a Web server), is terminated by proxy 'P' and two distinct transport sessions (TCP or UDP) are created between A and P on one hand, and P and B on the other. Following the connection setup, the proxy P relays data between the two sessions at application level.

[0012] The proxy architecture carries significant performance penalties due to the need to maintain transport (TCP or UDP) stack for all sessions flowing across the network, to copy data to relay all data at application level, and perform conversion from data frames to application buffers and back,

[0013] As a result of these limitations, the proxy architecture limits throughput of network optimization applications to 1-2 Gbps per standard Intel-based server, and number of concurrently supported flows to tens of thousands. The performance limitation effectively blocks the network optimization solutions from scaling to 10/40/100 Gbps network scale in an economical fashion.

[0014] Thus, there exists a need in for an alternative architecture for network optimization platforms that would eliminate the above bottlenecks of the application proxy architecture.

BRIEF SUMMARY OF THE DISCLOSURE

[0015] In accordance with one or more embodiments, a computer-implemented method is provided for transparently optimizing data transmission between a first endpoint and a second endpoint in a computer network. The endpoints have a directly established data session therebetween. The data session is identified by each endpoint at least to itself in the same way throughout the session. The method includes the steps of: relaying data between the endpoints transparently in the session using a network optimization service; and transparently modifying or storing at least some of the data transmitted from the second endpoint to the first endpoint using the network optimization service in order to optimize data communications between the endpoints, wherein transparently modifying at least some of the data comprises changing the data, replacing the data, or inserting additional data such that the first endpoint receives different data than was sent by the second endpoint.

[0016] In accordance with one or more further embodiments, an optimization service is provided for transparently optimizing data transmission between a first endpoint and a second endpoint in a computer network. The endpoints have a directly established data session therebetween. The data session is identified by each endpoint at least to itself in the same way throughout the session. The optimization service is configured to: relay data between the endpoints transparently in the session using a network optimization service; and transparently modify or store at least some of the data transmitted from the second endpoint to the first endpoint using the network optimization service in order to optimize data communications between the endpoints, wherein modification of data comprises changing the data, replacing the data, or inserting additional data such that the first endpoint receives different data than was sent by the second endpoint.

[0017] Various embodiments of the invention are provided in the following detailed description. As will be realized, the invention is capable of other and different embodiments, and its several details may be capable of modifications in various respects, all without departing from the invention. Accordingly, the drawings and description are to be regarded as illustrative in nature and not in a restrictive or limiting sense, with the scope of the application being indicated in the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] FIG. 1 is a flow diagram illustrating creation of a session between endpoints in accordance with one or more embodiments.

[0019] FIGS. 2A and 2B (collectively FIG. 2) are simplified diagrams illustrating deployment of an optimization service in accordance with one or more embodiments.

[0020] FIG. 3 is a simplified diagram illustrating deployment of an optimization service operating in a tunnel in accordance with one or more embodiments.

[0021] FIG. 4 is a flow diagram illustrating response caching in accordance with one or more embodiments.

[0022] FIG. 5 is a flow diagram illustrating data modification in accordance with one or more embodiments.

[0023] FIG. 6 is a flow diagram illustrating new request introduction in accordance with one or more embodiments.

[0024] FIG. 7 is a simplified diagram illustrating an exemplary network architecture in accordance with one or more embodiments.

[0025] Like or identical reference characters are used to identify common or similar elements.

DETAILED DESCRIPTION

[0026] Various embodiments disclosed herein are directed to a service for optimizing data transmission in a computer network between endpoints having a directly established session therebetween. The optimization service transparently modifies or stores at least some of the data transmitted between the endpoints or introduces a new request to an endpoint in order to optimize data communications between the endpoints. Each endpoint identifies the session to itself in the same way throughout the session.

[0027] As used herein, the term “network node” refers to any device, connected to an IP-based network, including, without limitation, computer servers, personal computers (including desktop, notebook, and tablet computers), smart phones, and other network connected devices.

[0028] As used herein, the term “endpoint” refers to an endpoint of a bi-directional inter-process communication flow across an IP-based network, residing on a network node connected to such network. Examples of endpoints include, without limitation, TCP sockets, SCTP sockets, UDP sockets, and raw IP sockets.

[0029] The optimization service operates as part of device involved in relaying of data between network nodes on an IP-based network. Examples of such devices include, without limitation, residential home gateways, WiFi hotspots, firewalls, routers, Metro Ethernet switches, optical switches, DPI devices, computer servers, application gateways, cable modem termination systems (CMTS), optical line terminals (OLT), broadband network gateways (BNG), broadband access servers (BRAS), DSL access multiplexers (DRAM), gateway GPRS support nodes (GGSN), and PDN gateways (PGW).

[0030] As shown in FIGS. 1 and 2, two endpoints ‘A’ and ‘B’ on an IP-based computer network, e.g., ISP subscriber and Internet-based web server, establish data session ‘S’ between each other. In case of TCP protocol, the session setup phase involves a TCP session handshake, including negotiation of network and transport parameters.

[0031] The session S between endpoints A and B involves data queries sent by A to B and, in some cases, data responses sent to B by A. The session may optionally include queries and responses sent by both endpoints.

[0032] Each respective endpoint typically identifies the data session S with at least a 5-tuple: IP address and port of the local endpoint, IP address and port of the remote endpoint and protocol used (e.g., TCP, UDP, or other). The definition of the

session S by the endpoints A and B may not be identical in case of network address translation (NAT) taking place in the network between A and B.

[0033] The endpoints A and B optionally keep track of the data sent and received, by counting bytes and/or frames sent and received. The endpoints A and B may further keep track of the data sent and received by the remote endpoint, for purposes of packet loss detection and retransmission, congestion avoidance, congestion control, among others.

[0034] The identification of the session S by each respective endpoint does not change throughout the session lifetime.

[0035] In accordance with one or more embodiments, following an establishment of the session S between endpoints A and B, an optimization service ‘C’ creates a transparent endpoint C_A facing the endpoint A, that would be appear to endpoint A as endpoint B, at both network and transport levels, as defined by TCP/IP model per RFC1122.

[0036] As illustrated in FIG. 1, the service C may optionally create two transparent endpoints C_A and C_B , with the endpoint C_A appearing to endpoint A as an endpoint B, and the endpoint C_B appearing to endpoint B as an endpoint A,

[0037] In accordance with one or more embodiments, the service C creates transparent endpoints per [0035-0036] in only some sessions it processes, with the decision being taken by C based on at least one variable including, e.g., temporal information, ordinal information, frequency information, endpoint identification information, session identification information, network state information, and external policy information.

[0038] In accordance with one or more embodiments, the service C may relay all data frames in the session S between A and B, either by being in data path between the session endpoints, or through use of one or more dedicated redirection devices (e.g., a load balancer, router, DPI device, etc.) that sit in data path and redirect specific data sessions to the service C, as depicted in FIGS. 2A and 2B.

[0039] In accordance with other embodiments, service C may relay only portion of data frames in the session between A and B. For example, the redirection device may redirect the session S to the service C starting from certain frame within the session, using Layer 7 analysis of the session to determine whether the session should be redirected to service C.

[0040] When the service C relays data frames between endpoint A and endpoint B without creating transparent endpoints, it may do so at a physical level (e.g., by switching data frames from port to port), or at link level (e.g., by changing MAC addresses and/or VLAN tags), or by combination of the above.

[0041] In accordance with one or more embodiments, the service C may optionally perform a network address translation (NAT) of the session it processes.

[0042] As part of the relaying, C is continuously tracking and storing the state of the connection, including all or some variables from the following group:

[0043] static session identifiers (endpoint addresses and port numbers, transport protocol used) dynamic transport state of each endpoint, including but not limited to sequence identifiers of data sent and acknowledged by each endpoint

[0044] negotiated transport attributes of the session and individual endpoints, including but not limited to TCP options, such as selective ACK, timestamp, scaled window and others

[0045] dynamic network-level attributes of data frames sent in each direction, including but not limited to IP DSCP, IP TOS, IPv6 flow label

[0046] dynamic link-level attributes of data frames sent in each direction, including but not limited to source and destination MAC addresses, 802.1Q VLAN tags, 802.1P priority bits, QinQ stacked VLAN tags

[0047] dynamic circuit-level attributes of data frames sent in each direction, including but not limited to identity of ingress and egress ports, physical port properties

[0048] According to one or more embodiments, the service C provides data modification and caching services for one or more data sessions S' between A and B, that traverse the service C in a tunnel established between two endpoints T_1 and T_2 , as illustrated in FIG. 3.

[0049] Tunnel protocols supported by the service C can include, but are not limited to, L2TP, PPPoE, PPPoA, L2TP, GRE, IP in IP, MPLS, Teredo, 6RD, 6to4, and PMIP.

[0050] According to one or more embodiments, as shown in FIG. 3, the service C tracks the state of the tunneled session between endpoints T_1 and T_2 , across multiple connections between endpoints A and B that traverse the tunnel.

[0051] As discussed in further detail below, the service C in accordance with various embodiments provides a number of session modification and other capabilities, including (a) data response caching, (b) modification of data queries and data responses (c) introduction of new requests.

[0052] (a) Response Caching

[0053] Following an establishment of data session S between endpoints A and B as described above in [0030-0034], the data query from endpoint A to endpoint B reaches the service C.

[0054] According to one or more embodiments, the service C analyzes the data query to match it with previously stored data responses. To do so, C analyzes the query received from endpoint A based on at least one variable, selected from the group consisting of temporal information, ordinal information, frequency information, client information, and identification information.

[0055] If a matching response is found in storage, C delivers the stored response to the endpoint A by itself.

[0056] According to one or more embodiments, should a matching response be previously stored by the service C, the service C does not relay the query to the endpoint B, but rather responds to the query by itself.

[0057] In accordance with one or more alternate embodiments, the service C relays the query received from endpoint A to endpoint B, receives a response or portion of it from endpoint B, and matches the data query received from endpoint A and the data response, or portion of it, received from endpoint B, against data responses previously stored by C.

[0058] In this case, should a matching stored response be identified by the service C, it delivers the stored response, or portion of it, to the endpoint A. In this case, service C blocks relaying of the response received from endpoint B.

[0059] In case of the data response being delivered by C, it may cause endpoint B to terminate the data session S on its end or stall delivery of the response.

[0060] When sending new data frames to endpoints A and B within session S, that were not received from the opposite endpoint, the service C utilizes the IP and port address of the opposite endpoint as well as the session state that is continuously stored by it, as described above in [0042].

[0061] Assume that endpoints A and B started sequencing their data streams in session S starting with X_0 and Y_0 , respectively, as depicted in FIG. 4. By the time service C receives the data query from endpoint A that C responds to, service C may have relayed N_A bytes of data from A to B and N_B bytes of data from B to A, where N_A and N_B can be larger than or equal to zero.

[0062] Service C keeps track of the sequences of both endpoints A and B and data acknowledged by each endpoint. When service C starts delivering its response to endpoint A, it starts sequencing its data with Y_0+N_B , in continuation of data sequences used by endpoint B earlier, while expecting new data from endpoint A starting from X_0+N_A , in continuation to sequences sent by endpoint A earlier. It can be said that C initializes an endpoint C_A with TCP sequence number Y_0+N_B and acknowledgement number X_0+N_A .

[0063] As a result, the data delivered by C appears to endpoint A as a seamless continuation of the session S between A and B.

[0064] In case of packets sent by endpoint B to endpoint A sequenced between Y_0 and Y_0+N_B are lost in the network segment between C and A, the endpoint A responds by sending back frames with acknowledgment sequence smaller than Y_0+N_B after service C started sending its own data sequenced Y_0+N_B and higher. According to one or more embodiments, C relays such packets to endpoint B, causing endpoint B to re-transmit the lost packets. In this case, C shall only relay back to endpoint A the data in the range between Y_0 and Y_0+N_B .

[0065] Similarly, when sending its own data (i.e., not data received from the other endpoint) to endpoint B, the service C utilizes the current state of endpoint A within session S, as seen by service C.

[0066] The description of sequencing of sent and received data done by service C in [0054-0059] applies equally to TCP-like semantics based on individual bytes of data as well as other semantics, including but not limited to sequencing of individual frames exchanged between the two endpoints.

[0067] Service C can apply same method of sequencing data as described in [[0054-0059], to multiple protocol layers within same session, including but not limited to TCP/IP session over PPP and PPP-like protocols, TCP/IP session over UDP/IP tunnel etc., session created in IPv6 over IPv4 tunnel, utilizing the data stored using multi-level session tracking as described above in [0043-0044].

[0068] According to one or more embodiments, to deliver the previously stored data response or other data, the service C transparently creates a transport endpoint C_A (e.g., TCP/IP UDP/IP socket), allowing it to deal with packet loss and retransmission, congestion detection and avoidance, and other aspects of transport data transmission, as done by the endpoints A and B.

[0069] Service C may create a single endpoint C_A facing endpoint A, or a pair of endpoints C_A and C_B , facing A and B respectively. The endpoint C_A facing endpoint A shall have an address of an opposite endpoint B (IP address IP_B and port P_B per [0054]) and the transport state of endpoint B as stored by service C as a result of session tracking prior to creation of endpoint C_A . In the same way, the endpoint C_B shall have the attributes of endpoint A (IP address IP_A and port P_A).

[0070] According to one or more embodiments, the service C stores data queries and data responses as they are relayed between endpoints A and B, without becoming a transport-level endpoint.

[0071] According to other embodiments, the service C may retrieve the data responses from one of the endpoints, or receive it from another data source.

[0072] According to one or more embodiments, the service C may respond to data queries from both endpoints A and B.

[0073] According to one or more embodiments, the service C responds to data queries from endpoints A and/or B, based on at least one variable from the following group: configuration information, temporal information, frequency information, ordinal information, system load information, network state information, client information and identification information.

[0074] According to one or more embodiments, endpoint A sends query QA₁ to endpoint B to which service C responds by sending previously stored response RC₁. Upon receiving response RC₁, endpoint A sends another data query QA₂. If service C does not have a matching response to the query QA₂ stored, it relays the query to endpoint B, receives response RB₂ and relays it to endpoint A.

[0075] As a result of response RC₁ delivered to endpoint A by service C, the counters of sent and received data of endpoint A and B may be in disagreement. To allow switch back to relay mode, where queries and responses are again relayed between A and B, service C performs an ongoing modification of sequences for data frames it relays between A and B as shown, e.g., in FIG. 4. For example, when request QA₂ is received from endpoint A, endpoint A reports receiving data up to Y₃' that reflects the data received from endpoint C_A, as part of RC₁ response. At the same time, endpoint B has sent data up to sequence Y₂, as part of its communication with endpoint C_B.

[0076] Similarly the counters of data that endpoint B received and endpoint A sent (X₃' and X₄ respectively), do not match either.

[0077] To eliminate the delta between Y₃' and Y₂ as well, as between X₃' and X₄, service C modifies sequences of sent and received data when relaying data between A and B, in both directions.

[0078] Furthermore, endpoint A may initiate another query QA₃, which can be replied to by service C, using previously stored response RC₃.

[0079] In other words, service C may alternate between responding to endpoint queries from one or both endpoints, and relaying queries and responses between two endpoints.

[0080] (b) Queries and Responses Modification

[0081] According to one or more embodiments, following establishment of session S, as described above in [0028-0031], the service C modifies data queries and/or data responses as relayed between two endpoints A and B, as illustrated by way of example in FIG. 5.

[0082] In accordance with one or more embodiments, service C does not utilize a transport endpoint for purposes of sending the modified data, but rather continues to track the transport state of endpoint A and B and relies on the sending endpoint to re-send the data in case of packet loss.

[0083] As part of modification of the relayed data, service C may need to change protocol checksums of the frames to reflect the new payload.

[0084] In case of modification by C of data sent from A to B, the packet loss of the modified data may occur between C and B. In this case, service C relays the data frames reflecting such loss from B to A, causing endpoint A to retransmit the lost frames and service C to re-apply modification again. C

tracks the re-transmitted frames using the stored session status information and re-applies the modification again.

[0085] According to other embodiments, to deliver modified data to endpoint B, service C creates a new transport endpoint C_B facing endpoint B. Such endpoint C_B utilizes IP address IP_A and port P_A of endpoint A, and relays the modified data in continuation of the frames previously relayed from endpoint A to endpoint B.

[0086] When service C creates an endpoint C_B to deliver modified data to endpoint B, service C may optionally create an endpoint C_A to facilitate communication with endpoint A, for example, for purposes of receiving data responses from it. Similarly to the endpoint C_B, the endpoint C_A utilizes IP address IP_B and port P_B of opposite endpoint B, and communicates with endpoint B in continuation of the frames previously relayed from endpoint B to endpoint A.

[0087] Upon completion of delivery of the modified data, service C may fall back to relaying frames between A and B, while making necessary adjustments for sequences of sent and received data, as described in [0069-0071].

[0088] According to one or more embodiments, service C modifies data queries and/or responses relayed from endpoint A to endpoint B, as part of negotiation of endpoint capabilities, in order to affect format, protocol, or other attribute of session S.

[0089] According to one or more embodiments, service C modifies parameters of a data query sent by endpoint A to negate a capability reported by endpoint A. For example, service C may modify the capability to receive a response in a compressed format reported by A, causing the opposite endpoint B to transmit its response in a compressed format.

[0090] Service C subsequently receives the compressed response RB₁₀, modifies it by decompressing the payload and delivers to endpoint A in a modified form, resulting in optimization of network between B and C and improved performance.

[0091] According to one or more embodiments, service C modifies response RB₁₁ received from endpoint B, including but not limited to rendering of the textual data in different format, image adaptation to endpoint device capabilities, change in video quality, and transcoding of audio and/or video data into different format, among others.

[0092] The modification of responses as described in [0084] can be done for a number of purposes, including improving utilization of network resources between the service C and the endpoint receiving the modified data, adapting the data responses to the endpoint application capabilities, improving application performance, among others.

[0093] According to one or more embodiments, the service C may modify data responses relayed between endpoints A and B, that pertain to data items or portions of data items available at one of or both endpoints, e.g., as utilized in peer-to-peer protocols like Bittorrent, eDonkey, and others.

[0094] (c) Introduction of New Requests

[0095] According to one or more embodiments, the service C may introduce new requests to endpoint A and/or endpoint B within session S, in addition and/or instead of queries sent by respective endpoints, as depicted in FIG. 6.

[0096] According to one or more embodiments, the service C may utilize an endpoint approach to transmission of new queries and reception of responses from endpoints A and B, as described in [0035-0037] and [0054-0059].

[0097] According to one or more embodiments, the service C combines caching of responses and response modification,

introduction of new requests with response caching, and relaying of data between endpoints in the same session S.

[0098] According to one or more embodiments, the service C modifies data availability responses in combination with as reported by one or both endpoints to improve the cache hit ratio of service C by including in it such data items (or portions of items) that are stored by the service C, and/or excluding such data items (or portions of item) that are not.

[0099] According to additional embodiments, the service C modifies the data availability information as reported by one or both endpoints to force the endpoints to transfer such data items (or portions of items) that are currently not stored by the service C, as a way to populate the cache managed by the service C.

[0100] According to one or more embodiments, the service C modifies data queries between endpoint A and endpoint B to disable use of end-to-end encryption, to allow subsequent caching of data responses.

[0101] According to one or more embodiments, the service C stores modified data responses as delivered by it to the endpoints, and may retrieve a stored copy of modified data response, rather than perform the modification on the fly.

[0102] According to one or more embodiments, the service C may deliver a data response stored through response caching mechanism as described above, if the stored copy of data response matches the needs for modification,

[0103] According to one or more embodiments, the service C may utilize a stored copy of a data response, stored through a response storing mechanism as described above, as an input for data modification, rather than allowing the full data response to be delivered from the endpoint B.

[0104] According to one or more embodiments, the service C may introduce new requests into session S, in order to trigger endpoints responses needed for optimal response caching.

[0105] Such data responses may include, but are not limited to, missing portions of content objects already stored in service C, content objects that have been identified as popular, however have not been stored by service C yet, content objects associated with other objects known to the service C (e.g., objects referenced by HTML page or additional playback levels for adaptive bitrate video etc.).

[0106] According to one or more embodiments, a system is provided for transparent modification of at least one data communications session between two endpoints A and B, in a way that requires endpoints A and B establish a data session between each other first, which includes at least one node of an IP network, designed and configured to provide at least one of services (a) to (c), as described above.

[0107] According to one or more embodiments, the optimization system can reside in single or multiple service provider networks, dedicated hosting location, datacenters, and enterprise or at residential premises as described in FIG. 7 below.

[0108] According to one or more embodiments, the system comprises multiple components in different physical locations.

[0109] According to one or more embodiments, multiple systems can reside in the data path of same connection S between two endpoints A and B in series.

[0110] According to one or more embodiments, the optimization service can operate at the same network node, on which one of the endpoints resides.

[0111] According to one or more embodiments, multiple optimization services can operate in series, as illustrated in FIG. 7.

[0112] According to one or more embodiments, multiple optimization services can operate in parallel, e.g., as part of load balancing of sessions done by redirecting device.

[0113] According to one or more embodiments, multiple instances of optimization services can operate in series and/or in parallel, wherein each instance of optimization carries out different and/or same data modification and storing operations.

[0114] According to one or more embodiments, multiple instances of optimization services can operate in series and/or in parallel, wherein each instance of optimization carries out different data modification and storing operations, in coordination with one another.

[0115] The processes of the optimization service described above may be implemented in software, hardware, firmware, or any combination thereof. The processes are preferably implemented in one or more computer programs executing on a programmable device including a processor, a storage medium readable by the processor (including, e.g., volatile and non-volatile memory and/or storage elements), and input and output devices. Each computer program can be a set of instructions (program code) in a code module resident in the random access memory of the device. Until required by the device, the set of instructions may be stored in another computer memory (e.g., in a hard disk drive, or in a removable memory such as an optical disk, external hard drive, memory card, or flash drive) or stored on another computer system and downloaded via the Internet or other network.

[0116] Having thus described several illustrative embodiments, it is to be appreciated that various alterations, modifications, and improvements will readily occur to those skilled in the art. Such alterations, modifications, and improvements are intended to form a part of this disclosure, and are intended to be within the spirit and scope of this disclosure. While some examples presented herein involve specific combinations of functions or structural elements, it should be understood that those functions and elements may be combined in other ways according to the present disclosure to accomplish the same or different objectives. In particular, acts, elements, and features discussed in connection with one embodiment are not intended to be excluded from similar or other roles in other embodiments.

[0117] Additionally, elements and components described herein may be further divided into additional components or joined together to form fewer components for performing the same functions. For example, the optimization service may comprise one or more physical machines, or virtual machines running on one or more physical machines. In addition, the optimization service may comprise a duster of computers or numerous distributed computers that are connected by the Internet or another network.

[0118] Accordingly, the foregoing description and attached drawings are by way of example only, and are not intended to be limiting.

What is claimed is:

1. A computer-implemented method for transparently optimizing data transmission between a first endpoint and a second endpoint in a computer network, said endpoints having a directly established data session therebetween, the data session being identified by each endpoint at least to itself in the same way throughout the session, the method comprising:

- relaying data between the endpoints transparently in the session using a network optimization service; and transparently modifying or storing at least some of the data transmitted from the second endpoint to the first endpoint using the network optimization service in order to optimize data communications between the endpoints, wherein transparently modifying at least some of the data comprises changing the data, replacing the data, or inserting additional data such that the first endpoint receives different data than was sent by the second endpoint.
2. The method of claim 1, wherein the session comprises a TCP or UDP session.
3. The method of claim 1, wherein the session is identified by identifiers including an IP address and transport port of the first endpoint, an IP address and transport port of the second endpoint, and a transport protocol used, and wherein the identifiers do not change throughout a lifetime of the session.
4. The method of claim 1, wherein the optimization service or a redirection service communicating with the optimization service is in a data path between the first and second endpoints.
5. The method of claim 1, wherein the step of relaying data includes creating a first transparent endpoint facing the first endpoint, wherein first transparent endpoint appears to the first endpoint as the second endpoint, such that the first endpoint communicates with the first transparent endpoint.
6. The method of claim 5, wherein the step of relaying data further includes creating a second transparent endpoint facing the second endpoint, wherein the second transparent endpoint appears to the second endpoint as the first endpoint, such that the second endpoint communicates with the second transparent endpoint.
7. The method of claim 1, wherein the step of relaying data is performed without creating a separate transparent endpoint, and instead by tracking the transport state of the endpoints, and relying on the endpoint sending data to resend data in case of packet loss, and modifying the resent data to be sent to the second endpoint.
8. The method of claim 1, wherein the step of relaying data includes receiving data communicated between the endpoints through a redirection device, and wherein the step of relaying data comprises relaying only a portion of data frames in the session between the first and second endpoints.
9. The method of claim 1, wherein, said step of relaying data includes generally continuously tracking and storing a state of the session, including tracking what data has been sent and received by each endpoint and transport-level attributes, link level attributes, or network-level attributes of the session.
10. The method of claim 9, wherein tracking involves monitoring multiple tiers of tunneling.
11. The method of claim 1, wherein the step of transparently modifying comprises analyzing a data query from the first endpoint directed to the second endpoint and/or a data response to the data query received from the second endpoint, matching the data query and the data response with a previously stored data response, and delivering the previously stored data response to the first endpoint, while blocking relaying of the data response from the second endpoint.
12. The method of claim 11, further comprising causing the second endpoint to terminate the session on its end or to stall delivery of the data response.
13. The method of claim 11, wherein delivering the previously stored data response comprises utilizing the IP and port address of the second endpoint and the session state of the session as a continuation of the session between the first and second endpoints.
14. The method of claim 11, wherein the step of matching is based on at least one variable selected from the group consisting of temporal information, ordinal information, frequency information, client information, and identification information.
15. The method of claim 1, wherein the step of transparently modifying comprises;
- analyzing a first data query from the first endpoint directed to the second endpoint, matching the first data query with a previously stored data response, and delivering the previously stored data response to the first endpoint through a separate transparent endpoint; and
- receiving a second data query from the first endpoint directed to the second endpoint, determining that a data response responding to the second data query has not been stored at the optimization service, relaying the second data query to the second endpoint, receiving the data response corresponding to the second query from the second endpoint, and relaying the data response from the second endpoint to the first endpoint without use of a separate transparent endpoint, while modifying the transport parameters of such response so that it appears as a continuation of the session between the first and second endpoints.
16. The method of claim 1, wherein the step of transparently modifying data comprises rendering textual data in a different format, adapting data to endpoint device capabilities, changing video quality, or transcoding audio and/or video data into a different format.
17. The method of claim 1, wherein the step of transparently modifying comprises modifying parameters of a data query or response from one endpoint to another to affect an attribute of the session.
18. The method of claim 17, wherein the step of relaying data includes creating a first transparent endpoint facing the first endpoint, wherein first transparent endpoint appears to the first endpoint as the second endpoint, such that the first endpoint communicates with the first transparent endpoint.
19. The method of claim 17, wherein the step of relaying data is performed without creating a separate transparent endpoint, and instead by tracking the transport state of the endpoints, and relying on the endpoint sending data to resend data in case of packet loss, and modifying the resent data to be sent to the second endpoint.
20. The method of claim 1, where step of transparently modifying data comprises modifying parameters of a data query or response from one endpoint to another to affect an attribute of subsequent data queries or responses so that the cacheability of data queries in the session is improved.
21. The method of claim 1, where step of transparently modifying data comprises utilizing previously stored data responses for subsequent modification of a data response.
22. The method of claim 1, wherein the step of transparently modifying data comprises modifying data availability responses reported by an endpoint to improve a cache hit ratio by the optimization service or to cause an endpoint to transfer such data items that are currently not stored at the optimization service, if the endpoints exchange information on data availability.

23. The method of claim **1**, wherein the step of introducing a new request is performed to cause an endpoint to deliver a data response needed for data caching.

24. An optimization service for transparently optimizing data transmission between a first endpoint and a second endpoint in a computer network, said endpoints having a directly established data session therebetween, the data session being identified by each endpoint at least to itself in the same way throughout the session, the optimization service being configured to:

relay data between the endpoints transparently in the session using a network optimization service; and

transparently modify or store at least some of the data transmitted from the second endpoint to the first endpoint using the network optimization service in order to optimize data communications between the endpoints, wherein modification of data comprises changing the data, replacing the data, or inserting additional data such that the first endpoint receives different data than was sent by the second endpoint.

25. The optimization service of claim **24**, wherein the optimization service is implemented in a device involved in relaying of data between network nodes on an IP-based network.

26. The optimization service of claim **24**, wherein the optimization service operates on a network node, on which one of the endpoints resides.

27. The optimization service of claim **24**, wherein the optimization service comprises one of a plurality of optimization services operating in series.

28. The optimization service of claim **24**, wherein the optimization service comprises one of a plurality of optimization services operating in parallel.

29. The optimization service of claim **24**, wherein the optimization service comprises multiple instances operating in series and/or in parallel, wherein each instance of the optimization service carries out different and/or same data modification and storing operations.

30. The optimization service of claim **24**, wherein the optimization service comprises multiple instances operating in series and/or in parallel, wherein each instance of the optimization service carries out different data modification and storing operations, in coordination with one another.

* * * * *