



- (51) **International Patent Classification:**  
*H04L 12/56* (2006.01)
- (21) **International Application Number:**  
PCT/US2012/027906
- (22) **International Filing Date:**  
6 March 2012 (06.03.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
61/451,874 11 March 2011 (11.03.2011) US
- (71) **Applicant (for all designated States except US):** CITRIX SYSTEMS, INC. [US/US]; 851 West Cypress Creek Road, Fort Lauderdale, FL 33309 (US).
- (72) **Inventor; and**
- (75) **Inventor/Applicant (for US only):** OVSIIANNIKOV, Michael [US/US]; c/o Citrix Systems, Inc., 4988 Great America Parkway, Santa Clara, CA 95054 (US).
- (74) **Agent:** MCKENNA, Christopher, J.; Foley & Lardner, LLP, 111 Huntington Avenue, Boston, MA 02199 (US).


(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

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

(54) **Title:** SYSTEMS AND METHODS OF QoS FOR SINGLE STREAM ICA

**FIG. 4A**

(57) **Abstract:** The present solution is directed to a system for providing quality of service (QoS) for a stream of protocol data units via a single transport layer connection. The system includes a device intermediary to a sender and receiver. The device receives via a single transport layer connection a plurality of packets carrying a plurality of protocol data units. Each protocol data unit of the plurality of protocol data units identifying a priority. One or more of the plurality of protocol data units having a priority different than the priority of other protocol data units of the plurality of protocol data units. The device may include a filter for determining an average priority for a predetermined window of protocol data units of the plurality of protocol data units and an engine for assigning the average priority as a connection priority of the single transport layer connection. The device transmits via the single transport layer connection the plurality of packets carrying those protocol data units within the predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the average priority for the predetermined window of protocol data units.

## **Systems and Methods of QoS for Single Stream ICA**

### Related Application

This application claims the benefit of and priority to U.S. Provisional Application No. 61/451874, entitled “Systems and Methods of QoS for Single Stream ICA”, and filed on  
5 March 11, 2011, which is incorporated herein by reference in its entirety for all purposes.

### Field of the Invention

The present application generally relates to data communication networks. In particular, the present application relates to systems and methods for applying quality of service to multi-priority stream of data carried via a single transport layer connection.

### Background of the Invention

A single transport layer connection between a remote display client and a server may be used for all communications. One remote display client is an ICA client. In ICA protocol the communication latency requirements may formulated in terms of “virtual streams or channels” which are assigned priorities. A high priority virtual channels may expect to have  
15 low communication latency. For example it may be desirable to maintain low latency for GUI interactions to improve human perceived responsiveness. To achieve this, GUI interaction events may be sent in the high priority virtual channel, while bulk data transfers (file copy for example) may be sent in the low priority channel.

### Summary of the Invention

20 The present solution addresses providing quality of service (QoS) to a transport layer connection given the different priorities of individual in flight packets within the connection. The ICA protocol enables QoS devices to extract priority information or tags from the network packets, in order to send packets according to the stream’s priority. Typically, a single network packet contains payload only of one virtual channel for ICA. For a single  
25 connection carrying ICA, out of order packet delivery based on virtual channel priority is desirable. As multiple virtual streams are time and space multiplexed into a single transport layer connection. Furthermore, the transport layer connection of Transport Control Protocol

(TCP) does not support out of order delivery. To make performing quality of service more challenging, an attempt to reorder packets within single TCP stream may increase latency and wastes bandwidth due the retransmissions algorithm of TCP implementations. Typical TCP retransmission algorithms assume small, normally constant, reordering window, and invokes “fast” retransmissions in case of packet reordering. With this ICA QoS may be performed on per connection basis, such as by giving bandwidth preference to certain connections at the expense of the other connections.

One approach to per connection QoS for these type of protocols is to assign as the connection priority the maximum priority of all the buffered packets or protocol data units (PDUs). For example, in the case of a typical file copy to local drive or printing, the corresponding connection will have a mix of bulk lower priority traffic and higher priority traffic due to user interface updates caused by user interface feedback, such as progress reports and animations. Due to multiple channels multiplexing over the same connection, assigning the maximum priority might not be optimal. The frequent user interface updates would effectively cause “flushing” of all buffered packets. This flushing may not produce any perceived improvement for the bulk transfer connection: the progress bar and animation updates would still be queued behind the bulk traffic. On the other hand, flushing may have less desirable and noticeable effect on the competing connections with interactive traffic. With this approach, the performance may be reciprocal of the number of PDUs buffered. The more PDUs that are buffered (e.g., waiting to be sent) the less effective the approach may be. The buffering consideration can be more important with split connections maintained by intermediary devices, such as WAN optimization applications. The amount of buffering in such applications is expected to be higher than in strictly QoS applications.

Another approach is to use an average of priorities of packets within the single connection, such as the average of buffered packets waiting to be sent or in flight packets. In other words, the solution may give more bandwidth to the connections that are on average “mostly” interactive, at the expense of the connections that are on average are “mostly” bulk. In some aspects, the average or “mostly” approach is useful. The solution may not try to improve bulk printing or file copy connection latency, as it is not attainable with TCP in order delivery, but instead to maintain acceptable latency for interactive connections. The averaging window can also be the same, large or smaller than the number of buffered PDUs. Furthermore, link utilization may be a factor to make “averaging” non linear or adaptive in view of link utilization of the connections.

In some aspects, embodiments of the present solution may implement an averaging approach or policy using a number of various techniques. In some embodiments, a straight average may be implemented such as via a “boxcar” FIR (Finite Impulse Response) type filter. A IIR (Infinite Impulse Response) filter may be employed to save space with a large window. In some embodiments, the system does not store PDUs but just their priorities. For example, with an embodiment of 8 different priorities, a 3 bit state per PDU may be used. Furthermore, any type of FIR and/or IIR filters may be employed. For example, in some embodiments, hysteresis, symmetric or asymmetric filtering may be used. In some embodiments, a “quick raise” effect may be used while in some embodiments, a “long fall” effect may be used. For example, the system may quickly lower priority of the connection when encountering low priority PDU(s), and then slowly raise the priority of the connection based on time and number of PDUs with specific priorities.

In some embodiments, PDUs may be replaced with bytes. Instead of sampling / counting PDUs, the system may sample and/or count bytes in each virtual channel can be counted. Counting bytes instead of PDUs will make the algorithm agnostic to the differences in PDU sizes. In further embodiments, the connection priorities can be translated into the standard TOC/diff serve QoS bits / packet marking to employ and work with other external QoS devices.

When employing standard QoS devices, the reordering effects within the single connection may be considered when assigning different priority tags within the same TCP session. Typically, standard QoS devices do not maintain per TCP connection state. Instead these standard QoS devices may assume that all packets within a single connection will have the same priority. To address this assumption, one approach is to explicitly or implicitly “rate limit” connection priority changes. For example choosing large filter window would implicitly limit the rate. Another approach is to change the priority from low to high when there are no packets in flight.

In some aspects, the present solution is directed to a method for providing quality of service (QoS) for a stream of protocol data units via a single transport layer connection. The method includes receiving, by a device via a single transport layer connection, a plurality of packets carrying a plurality of protocol data units. Each protocol data unit of the plurality of protocol data units identifying a priority. At least one or more of the plurality of protocol data units comprising a priority different than the priority of other protocol data units of the plurality of protocol data units. The method may include determining, by the device, an average priority for a predetermined window of protocol data units of the plurality of protocol

data units and assigning the average priority as a connection priority of the single transport layer connection. The method also includes transmitting, by the device via the single transport layer connection, the plurality of packets carrying those protocol data units within the predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the average priority for the predetermined window of protocol data units.

In some embodiments, the method includes receiving, by the device via the single transport layer connection the plurality of packets carrying the plurality of protocol data units of a remote display protocol. Each protocol data unit may correspond to a virtual channel of the remote display protocol and the virtual channel may have or be assigned the priority via the remote display protocol. In some embodiments, the method includes buffering, by the device, the plurality of packets until the connection priority for the plurality of packets is assigned to the transport layer connection.

In some embodiments, the method includes determining, by a filter of the device, the average priority based on finite impulse response (FIR) filtering. In some embodiments, the method includes determining, by a filter of the device, the average priority based on infinite impulse response (IIR) filtering. In some embodiments, the method includes identifying, by the device, the predetermined window as comprising a predetermined number of protocol data units. In some embodiments, the method includes setting, by the device, a Differentiated Services Code Point (DSCP) field in the plurality of packets to correspond to the connection priority assigned to the single transport layer connection.

In some embodiments, the method includes transmitting, by the device, a first packet of the plurality of data packets with the connection priority identified via one of an internet protocol (IP layer) or transport layer and the first packet carrying a first protocol data unit having the priority identified via a protocol above the transport layer. In some embodiments, the method includes determining, by the device, a second average priority for a second predetermined window of protocol data units of a second plurality of protocol data units and assigning, by the device, the second average priority as the connection priority of the single transport layer connection. The method may further include transmitting, by the device via the single transport layer connection, a second plurality of packets carrying those protocol data units within the second predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the second average priority for the second predetermined window of protocol data units.

In some aspects, the present solution is directed to a system for providing quality of service (QoS) for a stream of protocol data units via a single transport layer connection. The system includes a device intermediary to a sender and receiver. The device receives via a single transport layer connection a plurality of packets carrying a plurality of protocol data units. Each protocol data unit of the plurality of protocol data units identifying a priority. One or more of the plurality of protocol data units having a priority different than the priority of other protocol data units of the plurality of protocol data units. The device may include a filter for determining an average priority for a predetermined window of protocol data units of the plurality of protocol data units and an engine for assigning the average priority as a connection priority of the single transport layer connection. The device transmits via the single transport layer connection the plurality of packets carrying those protocol data units within the predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the average priority for the predetermined window of protocol data units.

In some embodiments, the device receives via the single transport layer connection the plurality of packets carrying the plurality of protocol data units of a remote display protocol, each protocol data unit corresponding to a virtual channel, the virtual channel assigned the priority. In some embodiments, the engine buffers the plurality of packets until the connection priority for the plurality of packets is assigned to the transport layer connection.

In some embodiments, the filter determines the average priority based on finite impulse response (FIR) filtering. In some embodiments, the filter determines the average priority based on infinite impulse response (IIR) filtering. In some embodiments, the engine identifies the predetermined window as comprising a predetermined number of protocol data units. In some embodiments, the engine sets the Differentiated Services Code Point (DSCP) field in the plurality of packets to correspond to the connection priority assigned to the single transport layer connection.

In some embodiments, the device transmits a first packet of the plurality of data packets with the connection priority identified via one of an internet protocol (IP layer) or transport layer and the first packet carrying a first protocol data unit having the priority identified via a protocol above the transport layer. In some embodiments, the filter determines a second average priority for a second predetermined window of protocol data units of a second plurality of protocol data units and the engine assigns the second average priority as the connection priority of the single transport layer connection. The device

transmits via the single transport layer connection, a second plurality of packets carrying those protocol data units within the second predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the second average priority for the second predetermined window of protocol data units.

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

### Brief Description of the Figures

The foregoing and other objects, aspects, features, and advantages of the invention will become more apparent and better understood by referring to the following description  
10 taken in conjunction with the accompanying drawings, in which:

FIG. 1A is a block diagram of an embodiment of a network environment for a client to access a server via one or more network optimization appliances;

FIG. 1B is a block diagram of another embodiment of a network environment for a client to access a server via one or more network optimization appliances in conjunction with  
15 other network appliances;

FIG. 1C is a block diagram of another embodiment of a network environment for a client to access a server via a single network optimization appliance deployed stand-alone or in conjunction with other network appliances;

FIGs. 1D and 1E are block diagrams of embodiments of a computing device;

20           FIG. 2A is a block diagram of an embodiment of an appliance for processing communications between a client and a server;

FIG. 2B is a block diagram of another embodiment of a client and/or server deploying the network optimization features of the appliance;

25           FIG. 3 is a block diagram of an embodiment of a client for communicating with a server using the network optimization feature;

FIG. 4A is a block diagram of an embodiment of a system for providing Quality of Service (QoS) to data streams of different priority packets within a single connection; and

FIG. 4B is a flow diagram of an embodiment of a method for providing Quality of Service (QoS) to data streams of different priority packets within a single connection.

30           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 of the Invention

For purposes of reading the description of the various embodiments of the present invention below, the following descriptions of the sections of the specification and their respective contents may be helpful:

- Section A describes a network environment and computing environment useful for practicing an embodiment of the present invention;
- Section B describes embodiments of a system and appliance architecture for accelerating delivery of a computing environment to a remote user;
- Section C describes embodiments of a client agent for accelerating communications between a client and a server; and
- Section D describes embodiments of systems and methods for Quality of Service of single connection have multiple priority data stream.

#### A. Network and Computing Environment

Prior to discussing the specifics of embodiments of the systems and methods of an appliance and/or client, it may be helpful to discuss the network and computing environments in which such embodiments may be deployed. Referring now to Figure 1A, an embodiment of a network environment is depicted. In brief overview, the network environment has one or more clients 102a-102n (also generally referred to as local machine(s) 102, or client(s) 102) in communication with one or more servers 106a-106n (also generally referred to as server(s) 106, or remote machine(s) 106) via one or more networks 104, 104', 104''. In some embodiments, a client 102 communicates with a server 106 via one or more network optimization appliances 200, 200' (generally referred to as appliance 200). In one embodiment, the network optimization appliance 200 is designed, configured or adapted to optimize Wide Area Network (WAN) network traffic. In some embodiments, a first appliance 200 works in conjunction or cooperation with a second appliance 200' to optimize network traffic. For example, a first appliance 200 may be located between a branch office and a WAN connection while the second appliance 200' is located between the WAN and a corporate Local Area Network (LAN). The appliances 200 and 200' may work together to

optimize the WAN related network traffic between a client in the branch office and a server on the corporate LAN.

Although FIG. 1A shows a network 104, network 104' and network 104'' (generally referred to as network(s) 104) between the clients 102 and the servers 106, the clients 102 and the servers 106 may be on the same network 104. The networks 104, 104', 104'' can be the same type of network or different types of networks. The network 104 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 networks 104, 104', 104'' can be a private or public network. In one embodiment, network 104' or network 104'' may be a private network and network 104 may be a public network. In some embodiments, network 104 may be a private network and network 104' and/or network 104'' a public network. In another embodiment, networks 104, 104', 104'' may be private networks. In some embodiments, clients 102 may be located at a branch office of a corporate enterprise communicating via a WAN connection over the network 104 to the servers 106 located on a corporate LAN in a corporate data center.

The network 104 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. In some embodiments, the network 104 may comprise a wireless link, such as an infrared channel or satellite band. The topology of the network 104 may be a bus, star, or ring network topology. The network 104 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 described herein.

As depicted in FIG. 1A, a first network optimization appliance 200 is shown between networks 104 and 104' and a second network optimization appliance 200' is also between networks 104' and 104''. In some embodiments, the appliance 200 may be located on network 104. For example, a corporate enterprise may deploy an appliance 200 at the branch office. In other embodiments, the appliance 200 may be located on network 104'. In some embodiments, the appliance 200' may be located on network 104' or network 104''. For example, an appliance 200 may be located at a corporate data center. In one embodiment, the appliance 200 and 200' are on the same network. In another embodiment, the appliance 200 and 200' are on different networks.

In one embodiment, the appliance 200 is a device for accelerating, optimizing or otherwise improving the performance, operation, or quality of service of any type and form of network traffic. In some embodiments, the appliance 200 is a performance enhancing proxy. In other embodiments, the appliance 200 is any type and form of WAN optimization or acceleration device, sometimes also referred to as a WAN optimization controller. In one embodiment, the appliance 200 is any of the product embodiments referred to as WANScaler manufactured by Citrix Systems, Inc. of Ft. Lauderdale, Florida. In other embodiments, the appliance 200 includes any of the product embodiments referred to as BIG-IP link controller and WANjet manufactured by F5 Networks, Inc. of Seattle, Washington. In another embodiment, the appliance 200 includes any of the WX and WXC WAN acceleration device platforms manufactured by Juniper Networks, Inc. of Sunnyvale, California. In some embodiments, the appliance 200 includes any of the steelhead line of WAN optimization appliances manufactured by Riverbed Technology of San Francisco, California. In other embodiments, the appliance 200 includes any of the WAN related devices manufactured by Expand Networks Inc. of Roseland, New Jersey. In one embodiment, the appliance 200 includes any of the WAN related appliances manufactured by Packeteer Inc. of Cupertino, California, such as the PacketShaper, iShared, and SkyX product embodiments provided by Packeteer. In yet another embodiment, the appliance 200 includes any WAN related appliances and/or software manufactured by Cisco Systems, Inc. of San Jose, California, such as the Cisco Wide Area Network Application Services software and network modules, and Wide Area Network engine appliances.

In some embodiments, the appliance 200 provides application and data acceleration services for branch-office or remote offices. In one embodiment, the appliance 200 includes optimization of Wide Area File Services (WAFS). In another embodiment, the appliance 200 accelerates the delivery of files, such as via the Common Internet File System (CIFS) protocol. In other embodiments, the appliance 200 provides caching in memory and/or storage to accelerate delivery of applications and data. In one embodiment, the appliance 200 provides compression of network traffic at any level of the network stack or at any protocol or network layer. In another embodiment, the appliance 200 provides transport layer protocol optimizations, flow control, performance enhancements or modifications and/or management to accelerate delivery of applications and data over a WAN connection. For example, in one embodiment, the appliance 200 provides Transport Control Protocol (TCP) optimizations. In other embodiments, the appliance 200 provides optimizations, flow control, performance enhancements or modifications and/or management for any session or application layer

protocol. Further details of the optimization techniques, operations and architecture of the appliance 200 are discussed below in Section B.

Still referring to FIG. 1A, the network environment may include multiple, logically-grouped servers 106. In these embodiments, the logical group of servers may be referred to as a server farm 38. In some of these embodiments, the servers 106 may be geographically dispersed. In some cases, a farm 38 may be administered as a single entity. In other embodiments, the server farm 38 comprises a plurality of server farms 38. In one embodiment, the server farm executes one or more applications on behalf of one or more clients 102.

The servers 106 within each farm 38 can be heterogeneous. One or more of the servers 106 can operate according to one type of operating system platform (e.g., WINDOWS NT, manufactured by Microsoft Corp. of Redmond, Washington), while one or more of the other servers 106 can operate according to another type of operating system platform (e.g., Unix or Linux). The servers 106 of each farm 38 do not need to be physically proximate to another server 106 in the same farm 38. Thus, the group of servers 106 logically grouped as a farm 38 may be interconnected using a wide-area network (WAN) connection or metropolitan-area network (MAN) connection. For example, a farm 38 may include servers 106 physically located in different continents or different regions of a continent, country, state, city, campus, or room. Data transmission speeds between servers 106 in the farm 38 can be increased if the servers 106 are connected using a local-area network (LAN) connection or some form of direct connection.

Servers 106 may be referred to as a file server, application server, web server, proxy server, or gateway server. In some embodiments, a server 106 may have the capacity to function as either an application server or as a master application server. In one embodiment, a server 106 may include an Active Directory. The clients 102 may also be referred to as client nodes or endpoints. In some embodiments, a client 102 has the capacity to function as both a client node seeking access to applications on a server and as an application server providing access to hosted applications for other clients 102a-102n.

In some embodiments, a client 102 communicates with a server 106. In one embodiment, the client 102 communicates directly with one of the servers 106 in a farm 38. In another embodiment, the client 102 executes a program neighborhood application to communicate with a server 106 in a farm 38. In still another embodiment, the server 106 provides the functionality of a master node. In some embodiments, the client 102 communicates with the server 106 in the farm 38 through a network 104. Over the network

104, the client 102 can, for example, request execution of various applications hosted by the servers 106a-106n in the farm 38 and receive output of the results of the application execution for display. In some embodiments, only the master node provides the functionality required to identify and provide address information associated with a server 106' hosting a requested application.

In one embodiment, the server 106 provides functionality of a web server. In another embodiment, the server 106a receives requests from the client 102, forwards the requests to a second server 106b and responds to the request by the client 102 with a response to the request from the server 106b. In still another embodiment, the server 106 acquires an enumeration of applications available to the client 102 and address information associated with a server 106 hosting an application identified by the enumeration of applications. In yet another embodiment, the server 106 presents the response to the request to the client 102 using a web interface. In one embodiment, the client 102 communicates directly with the server 106 to access the identified application. In another embodiment, the client 102 receives application output data, such as display data, generated by an execution of the identified application on the server 106.

#### Deployed With Other Appliances.

Referring now to FIG. 1B, another embodiment of a network environment is depicted in which the network optimization appliance 200 is deployed with one or more other appliances 205, 205' (generally referred to as appliance 205 or second appliance 205) such as a gateway, firewall or acceleration appliance. For example, in one embodiment, the appliance 205 is a firewall or security appliance while appliance 205' is a LAN acceleration device. In some embodiments, a client 102 may communicate to a server 106 via one or more of the first appliances 200 and one or more second appliances 205.

One or more appliances 200 and 205 may be located at any point in the network or network communications path between a client 102 and a server 106. In some embodiments, a second appliance 205 may be located on the same network 104 as the first appliance 200. In other embodiments, the second appliance 205 may be located on a different network 104 as the first appliance 200. In yet another embodiment, a first appliance 200 and second appliance 205 is on the same network, for example network 104, while the first appliance 200' and second appliance 205' is on the same network, such as network 104'.

In one embodiment, the second appliance 205 includes any type and form of transport control protocol or transport later terminating device, such as a gateway or firewall device.

In one embodiment, the appliance 205 terminates the transport control protocol by establishing a first transport control protocol connection with the client and a second

5 transport control connection with the second appliance or server. In another embodiment, the appliance 205 terminates the transport control protocol by changing, managing or controlling the behavior of the transport control protocol connection between the client and the server or second appliance. For example, the appliance 205 may change, queue, forward or transmit network packets in manner to effectively terminate the transport control protocol connection  
10 or to act or simulate as terminating the connection.

In some embodiments, the second appliance 205 is a performance enhancing proxy.

In one embodiment, the appliance 205 provides a virtual private network (VPN) connection.

In some embodiments, the appliance 205 provides a Secure Socket Layer VPN (SSL VPN) connection. In other embodiments, the appliance 205 provides an IPsec (Internet Protocol

15 Security) based VPN connection. In some embodiments, the appliance 205 provides any one or more of the following functionality: compression, acceleration, load-balancing, switching/routing, caching, and Transport Control Protocol (TCP) acceleration.

In one embodiment, the appliance 205 is any of the product embodiments referred to as Access Gateway, Application Firewall, Application Gateway, or NetScaler manufactured

20 by Citrix Systems, Inc. of Ft. Lauderdale, Florida. As such, in some embodiments, the appliance 205 includes any logic, functions, rules, or operations to perform services or functionality such as SSL VPN connectivity, SSL offloading, switching/load balancing, Domain Name Service resolution, LAN acceleration and an application firewall.

In some embodiments, the appliance 205 provides a SSL VPN connection between a

25 client 102 and a server 106. For example, a client 102 on a first network 104 requests to establish a connection to a server 106 on a second network 104'. In some embodiments, the second network 104'' is not routable from the first network 104. In other embodiments, the client 102 is on a public network 104 and the server 106 is on a private network 104', such as a corporate network. In one embodiment, a client agent intercepts communications of the  
30 client 102 on the first network 104, encrypts the communications, and transmits the communications via a first transport layer connection to the appliance 205. The appliance 205 associates the first transport layer connection on the first network 104 to a second transport layer connection to the server 106 on the second network 104. The appliance 205 receives the intercepted communication from the client agent, decrypts the communications,

and transmits the communication to the server 106 on the second network 104 via the second transport layer connection. The second transport layer connection may be a pooled transport layer connection. In one embodiment,, the appliance 205 provides an end-to-end secure transport layer connection for the client 102 between the two networks 104, 104'

5 In one embodiment, the appliance 205 hosts an intranet internet protocol or intranetIP address of the client 102 on the virtual private network 104. The client 102 has a local network identifier, such as an internet protocol (IP) address and/or host name on the first network 104. When connected to the second network 104' via the appliance 205, the appliance 205 establishes, assigns or otherwise provides an IntranetIP, which is network  
10 identifier, such as IP address and/or host name, for the client 102 on the second network 104'. The appliance 205 listens for and receives on the second or private network 104' for any communications directed towards the client 102 using the client's established IntranetIP. In one embodiment, the appliance 205 acts as or on behalf of the client 102 on the second private network 104.

15 In some embodiment, the appliance 205 has an encryption engine providing 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 encrypts and decrypts network packets, or any portion thereof, communicated via the appliance 205. The encryption engine may also setup or establish SSL or TLS connections  
20 on behalf of the client 102a-102n, server 106a-106n, or appliance 200, 205. As such, the encryption engine provides offloading and acceleration of SSL processing. In one embodiment, the encryption engine 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 uses an encryption processor. In other embodiments, the encryption  
25 engine includes executable instructions running on an encryption processor.

In some embodiments, the appliance 205 provides one or more of the following acceleration techniques to communications between the client 102 and server 106: 1) compression, 2) decompression, 3) Transmission Control Protocol pooling, 4) Transmission Control Protocol multiplexing, 5) Transmission Control Protocol buffering, and 6) caching.

30 In one embodiment, the appliance 200 relieves servers 106 of much of the processing load caused by repeatedly opening and closing transport layers connections to clients 102 by opening one or more transport layer connections with each server 106 and maintaining these connections to allow repeated data accesses by clients via the Internet. This technique is referred to herein as "connection pooling".

In some embodiments, in order to seamlessly splice communications from a client 102 to a server 106 via a pooled transport layer connection, the appliance 205 translates or multiplexes communications by modifying sequence number and acknowledgment numbers at the transport layer protocol level. This is referred to as “connection multiplexing”. In some  
5 embodiments, no application layer protocol interaction is required. For example, in the case of an in-bound packet (that is, a packet received from a client 102), the source network address of the packet is changed to that of an output port of appliance 205, and the destination network address is changed to that of the intended server. In the case of an outbound packet (that is, one received from a server 106), the source network address is changed from that of  
10 the server 106 to that of an output port of appliance 205 and the destination address is changed from that of appliance 205 to that of the requesting client 102. The sequence numbers and acknowledgment numbers of the packet are also translated to sequence numbers and acknowledgement expected by the client 102 on the appliance’s 205 transport layer connection to the client 102. In some embodiments, the packet checksum of the transport  
15 layer protocol is recalculated to account for these translations.

In another embodiment, the appliance 205 provides switching or load-balancing functionality for communications between the client 102 and server 106. In some embodiments, the appliance 205 distributes traffic and directs client requests to a server 106 based on layer 4 payload or application-layer request data. In one embodiment, although the  
20 network layer or layer 2 of the network packet identifies a destination server 106, the appliance 205 determines the server 106 to distribute the network packet by application information and data carried as payload of the transport layer packet. In one embodiment, a health monitoring program of the appliance 205 monitors the health of servers to determine the server 106 for which to distribute a client’s request. In some embodiments, if the  
25 appliance 205 detects a server 106 is not available or has a load over a predetermined threshold, the appliance 205 can direct or distribute client requests to another server 106.

In some embodiments, the appliance 205 acts as a Domain Name Service (DNS) resolver or otherwise provides resolution of a DNS request from clients 102. In some embodiments, the appliance intercepts’ a DNS request transmitted by the client 102. In one  
30 embodiment, the appliance 205 responds to a client’s DNS request with an IP address of or hosted by the appliance 205. In this embodiment, the client 102 transmits network communication for the domain name to the appliance 200. In another embodiment, the appliance 200 responds to a client’s DNS request with an IP address of or hosted by a second

appliance 200'. In some embodiments, the appliance 205 responds to a client's DNS request with an IP address of a server 106 determined by the appliance 200.

In yet another embodiment, the appliance 205 provides application firewall functionality for communications between the client 102 and server 106. In one embodiment, a policy engine 295' provides rules for detecting and blocking illegitimate requests. In some embodiments, the application firewall protects against denial of service (DoS) attacks. In other embodiments, the appliance inspects the content of intercepted requests to identify and block application-based attacks. In some embodiments, the rules/policy engine includes one or more application firewall or security control policies for providing protections against various classes and types of web or Internet based vulnerabilities, such as one or more of the following: 1) buffer overflow, 2) CGI-BIN parameter manipulation, 3) form/hidden field manipulation, 4) forceful browsing, 5) cookie or session poisoning, 6) broken access control list (ACLs) or weak passwords, 7) cross-site scripting (XSS), 8) command injection, 9) SQL injection, 10) error triggering sensitive information leak, 11) insecure use of cryptography, 12) server misconfiguration, 13) back doors and debug options, 14) website defacement, 15) platform or operating systems vulnerabilities, and 16) zero-day exploits. In an embodiment, the application firewall of the appliance provides HTML form field protection in the form of inspecting or analyzing the network communication for one or more of the following: 1) required fields are returned, 2) no added field allowed, 3) read-only and hidden field enforcement, 4) drop-down list and radio button field conformance, and 5) form-field max-length enforcement. In some embodiments, the application firewall of the appliance 205 ensures cookies are not modified. In other embodiments, the appliance 205 protects against forceful browsing by enforcing legal URLs.

In still yet other embodiments, the application firewall appliance 205 protects any confidential information contained in the network communication. The appliance 205 may inspect or analyze any network communication in accordance with the rules or policies of the policy engine to identify any confidential information in any field of the network packet. In some embodiments, the application firewall identifies in the network communication one or more occurrences of a credit card number, password, social security number, name, patient code, contact information, and age. The encoded portion of the network communication may include these occurrences or the confidential information. Based on these occurrences, in one embodiment, the application firewall may take a policy action on the network communication, such as prevent transmission of the network communication. In another

embodiment, the application firewall may rewrite, remove or otherwise mask such identified occurrence or confidential information.

Although generally referred to as a network optimization or first appliance 200 and a second appliance 205, the first appliance 200 and second appliance 205 may be the same type and form of appliance. In one embodiment, the second appliance 205 may perform the same functionality, or portion thereof, as the first appliance 200, and vice-versa. For example, the first appliance 200 and second appliance 205 may both provide acceleration techniques. In one embodiment, the first appliance may perform LAN acceleration while the second appliance performs WAN acceleration, or vice-versa. In another example, the first appliance 200 may also be a transport control protocol terminating device as with the second appliance 205. Furthermore, although appliances 200 and 205 are shown as separate devices on the network, the appliance 200 and/or 205 could be a part of any client 102 or server 106.

Referring now to FIG. 1C, other embodiments of a network environment for deploying the appliance 200 are depicted. In another embodiment as depicted on the top of FIG. 1C, the appliance 200 may be deployed as a single appliance or single proxy on the network 104. For example, the appliance 200 may be designed, constructed or adapted to perform WAN optimization techniques discussed herein without a second cooperating appliance 200'. In other embodiments as depicted on the bottom of FIG. 1C, a single appliance 200 may be deployed with one or more second appliances 205. For example, a WAN acceleration first appliance 200, such as a Citrix WANScaler appliance, may be deployed with a LAN accelerating or Application Firewall second appliance 205, such as a Citrix NetScaler appliance.

### Computing Device

The client 102, server 106, and appliance 200 and 205 may be deployed as and/or executed on any type and form of computing device, such as a computer, network device or appliance capable of communicating on any type and form of network and performing the operations described herein. FIGs. 1C and 1D depict block diagrams of a computing device 100 useful for practicing an embodiment of the client 102, server 106 or appliance 200. As shown in FIGs. 1C and 1D, each computing device 100 includes a central processing unit 101, and a main memory unit 122. As shown in FIG. 1C, a computing device 100 may include a visual display device 124, a keyboard 126 and/or a pointing device 127, such as a mouse. Each computing device 100 may also include additional optional elements, such as

one or more input/output devices 130a-130b (generally referred to using reference numeral 130), and a cache memory 140 in communication with the central processing unit 101.

The central processing unit 101 is any logic circuitry that responds to and processes instructions fetched from the main memory unit 122. 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; the RS/6000 processor, those manufactured by International Business Machines of White Plains, New York; or those manufactured by Advanced Micro Devices of Sunnyvale, California. The computing device 100 may be based on any of these processors, or any other processor capable of operating as described herein.

Main memory unit 122 may be one or more memory chips capable of storing data and allowing any storage location to be directly accessed by the microprocessor 101, 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 122 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. 1C, the processor 101 communicates with main memory 122 via a system bus 150 (described in more detail below). FIG. 1C depicts an embodiment of a computing device 100 in which the processor communicates directly with main memory 122 via a memory port 103. For example, in FIG. 1D the main memory 122 may be DRDRAM.

FIG. 1D depicts an embodiment in which the main processor 101 communicates directly with cache memory 140 via a secondary bus, sometimes referred to as a backside bus. In other embodiments, the main processor 101 communicates with cache memory 140 using the system bus 150. Cache memory 140 typically has a faster response time than main memory 122 and is typically provided by SRAM, BSRAM, or EDRAM. In the embodiment shown in FIG. 1C, the processor 101 communicates with various I/O devices 130 via a local system bus 150. Various busses may be used to connect the central processing unit 101 to any of the I/O devices 130, 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 124, the processor 101 may use an Advanced Graphics Port (AGP) to communicate with the display 124. FIG. 1D depicts an embodiment of a computer 100 in which the main processor 101 communicates directly with I/O device 130 via HyperTransport, Rapid I/O, or InfiniBand. FIG. 1D also depicts an embodiment in which local busses and direct communication are mixed: the processor 101 communicates with I/O device 130 using a local interconnect bus while communicating with I/O device 130 directly.

The computing device 100 may support any suitable installation device 116, 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 client agent 120, or portion thereof. The computing device 100 may further comprise a storage device 128, 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 client agent 120. Optionally, any of the installation devices 116 could also be used as the storage device 128. Additionally, the operating system and the software can be run from a bootable medium, for example, a bootable CD, such as KNOPPIX®, a bootable CD for GNU/Linux that is available as a GNU/Linux distribution from knoppix.net.

Furthermore, the computing device 100 may include a network interface 118 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 118 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 100 to any type of network capable of communication and performing the operations described herein.

A wide variety of I/O devices 130a-130n may be present in the computing device 100. 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 130 may be controlled by an I/O controller 123 as shown in FIG. 1C. The I/O controller may control one or more I/O devices such as a

keyboard 126 and a pointing device 127, e.g., a mouse or optical pen. Furthermore, an I/O device may also provide storage 128 and/or an installation medium 116 for the computing device 100. In still other embodiments, the computing device 100 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.

In some embodiments, the computing device 100 may comprise or be connected to multiple display devices 124a-124n, which each may be of the same or different type and/or form. As such, any of the I/O devices 130a-130n and/or the I/O controller 123 may comprise any type and/or form of suitable hardware, software, or combination of hardware and software to support, enable or provide for the connection and use of multiple display devices 124a-124n by the computing device 100. For example, the computing device 100 may include any type and/or form of video adapter, video card, driver, and/or library to interface, communicate, connect or otherwise use the display devices 124a-124n. In one embodiment, a video adapter may comprise multiple connectors to interface to multiple display devices 124a-124n. In other embodiments, the computing device 100 may include multiple video adapters, with each video adapter connected to one or more of the display devices 124a-124n. In some embodiments, any portion of the operating system of the computing device 100 may be configured for using multiple displays 124a-124n. In other embodiments, one or more of the display devices 124a-124n may be provided by one or more other computing devices, such as computing devices 100a and 100b connected to the computing device 100, for example, via a network. These embodiments may include any type of software designed and constructed to use another computer's display device as a second display device 124a for the computing device 100. One ordinarily skilled in the art will recognize and appreciate the various ways and embodiments that a computing device 100 may be configured to have multiple display devices 124a-124n.

In further embodiments, an I/O device 130 may be a bridge 170 between the system bus 150 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.

A computing device 100 of the sort depicted in FIGs. 1C and 1D typically operate under the control of operating systems, which control scheduling of tasks and access to system resources. The computing device 100 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 real-time operating system, any open source operating system, any proprietary operating system, any operating systems for mobile  
5 computing 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  
10 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.

In other embodiments, the computing device 100 may have different processors,  
15 operating systems, and input devices consistent with the device. For example, in one embodiment the computer 100 is a Treo 180, 270, 1060, 600 or 650 smart phone manufactured by Palm, Inc. In this embodiment, the Treo smart phone is operated under the control of the PalmOS operating system and includes a stylus input device as well as a five-way navigator device. Moreover, the computing device 100 can be any workstation, desktop  
20 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 described herein.

## 25 B. System and Appliance Architecture

Referring now to FIG. 2A, an embodiment of a system environment and architecture of an appliance 200 for delivering and/or operating a computing environment on a client is depicted. In some embodiments, a server 106 includes an application delivery system 290 for  
30 delivering a computing environment or an application and/or data file to one or more clients 102. In brief overview, a client 102 is in communication with a server 106 via network 104 and appliance 200. For example, the client 102 may reside in a remote office of a company, e.g., a branch office, and the server 106 may reside at a corporate data center. The client 102 has a client agent 120, and a computing environment 215. The computing environment 215

may execute or operate an application that accesses, processes or uses a data file. The computing environment 215, application and/or data file may be delivered via the appliance 200 and/or the server 106.

In some embodiments, the appliance 200 accelerates delivery of a computing environment 215, or any portion thereof, to a client 102. In one embodiment, the appliance 200 accelerates the delivery of the computing environment 215 by the application delivery system 290. For example, the embodiments described herein may be used to accelerate delivery of a streaming application and data file processable by the application from a central corporate data center to a remote user location, such as a branch office of the company. In another embodiment, the appliance 200 accelerates transport layer traffic between a client 102 and a server 106. In another embodiment, the appliance 200 controls, manages, or adjusts the transport layer protocol to accelerate delivery of the computing environment. In some embodiments, the appliance 200 uses caching and/or compression techniques to accelerate delivery of a computing environment.

In some embodiments, the application delivery management system 290 provides application delivery techniques to deliver a computing environment to a desktop of a user, remote or otherwise, based on a plurality of execution methods and based on any authentication and authorization policies applied via a policy engine 295. With these techniques, a remote user may obtain a computing environment and access to server stored applications and data files from any network connected device 100. In one embodiment, the application delivery system 290 may reside or execute on a server 106. In another embodiment, the application delivery system 290 may reside or execute on a plurality of servers 106a-106n. In some embodiments, the application delivery system 290 may execute in a server farm 38. In one embodiment, the server 106 executing the application delivery system 290 may also store or provide the application and data file. In another embodiment, a first set of one or more servers 106 may execute the application delivery system 290, and a different server 106n may store or provide the application and data file. In some embodiments, each of the application delivery system 290, the application, and data file may reside or be located on different servers. In yet another embodiment, any portion of the application delivery system 290 may reside, execute or be stored on or distributed to the appliance 200, or a plurality of appliances.

The client 102 may include a computing environment 215 for executing an application that uses or processes a data file. The client 102 via networks 104, 104' and appliance 200 may request an application and data file from the server 106. In one

embodiment, the appliance 200 may forward a request from the client 102 to the server 106. For example, the client 102 may not have the application and data file stored or accessible locally. In response to the request, the application delivery system 290 and/or server 106 may deliver the application and data file to the client 102. For example, in one embodiment, the server 106 may transmit the application as an application stream to operate in computing environment 215 on client 102.

In some embodiments, the application delivery system 290 comprises any portion of the Citrix Access Suite™ by Citrix Systems, Inc., such as the MetaFrame or Citrix Presentation Server™ and/or any of the Microsoft® Windows Terminal Services manufactured by the Microsoft Corporation. In one embodiment, the application delivery system 290 may deliver one or more applications to clients 102 or users via a remote-display protocol or otherwise via remote-based or server-based computing. In another embodiment, the application delivery system 290 may deliver one or more applications to clients or users via streaming of the application.

In one embodiment, the application delivery system 290 includes a policy engine 295 for controlling and managing the access to, selection of application execution methods and the delivery of applications. In some embodiments, the policy engine 295 determines the one or more applications a user or client 102 may access. In another embodiment, the policy engine 295 determines how the application should be delivered to the user or client 102, e.g., the method of execution. In some embodiments, the application delivery system 290 provides a plurality of delivery techniques from which to select a method of application execution, such as a server-based computing, streaming or delivering the application locally to the client 120 for local execution.

In one embodiment, a client 102 requests execution of an application program and the application delivery system 290 comprising a server 106 selects a method of executing the application program. In some embodiments, the server 106 receives credentials from the client 102. In another embodiment, the server 106 receives a request for an enumeration of available applications from the client 102. In one embodiment, in response to the request or receipt of credentials, the application delivery system 290 enumerates a plurality of application programs available to the client 102. The application delivery system 290 receives a request to execute an enumerated application. The application delivery system 290 selects one of a predetermined number of methods for executing the enumerated application, for example, responsive to a policy of a policy engine. The application delivery system 290 may select a method of execution of the application enabling the client 102 to receive

application-output data generated by execution of the application program on a server 106. The application delivery system 290 may select a method of execution of the application enabling the client or local machine 102 to execute the application program locally after retrieving a plurality of application files comprising the application. In yet another  
5 embodiment, the application delivery system 290 may select a method of execution of the application to stream the application via the network 104 to the client 102.

A client 102 may execute, operate or otherwise provide an application, which can be any type and/or form of software, program, or executable instructions such as any type and/or form of web browser, web-based client, client-server application, a thin-client computing  
10 client, an ActiveX control, or a Java applet, or any other type and/or form of executable instructions capable of executing on client 102. In some embodiments, the application may be a server-based or a remote-based application executed on behalf of the client 102 on a server 106. In one embodiment the server 106 may display output to the client 102 using any thin-client or remote-display protocol, such as the Independent Computing Architecture  
15 (ICA) protocol manufactured by Citrix Systems, Inc. of Ft. Lauderdale, Florida or the Remote Desktop Protocol (RDP) manufactured by the Microsoft Corporation of Redmond, Washington. The application can use any type of protocol and it can be, for example, an HTTP client, an FTP client, an Oscar client, or a Telnet client. In other embodiments, the application comprises any type of software related to VoIP communications, such as a soft IP  
20 telephone. In further embodiments, the application comprises any application related to real-time data communications, such as applications for streaming video and/or audio.

In some embodiments, the server 106 or a server farm 38 may be running one or more applications, such as an application providing a thin-client computing or remote display presentation application. In one embodiment, the server 106 or server farm 38 executes as an  
25 application, any portion of the Citrix Access Suite™ by Citrix Systems, Inc., such as the MetaFrame or Citrix Presentation Server™, and/or any of the Microsoft® Windows Terminal Services manufactured by the Microsoft Corporation. In one embodiment, the application is an ICA client, developed by Citrix Systems, Inc. of Fort Lauderdale, Florida. In other embodiments, the application includes a Remote Desktop (RDP) client, developed by  
30 Microsoft Corporation of Redmond, Washington. Also, the server 106 may run an application, which for example, may be an application server providing email services such as Microsoft Exchange manufactured by the Microsoft Corporation of Redmond, Washington, a web or Internet server, or a desktop sharing server, or a collaboration server. In some embodiments, any of the applications may comprise any type of hosted service or

products, such as GoToMeeting™ provided by Citrix Online Division, Inc. of Santa Barbara, California, WebEx™ provided by WebEx, Inc. of Santa Clara, California, or Microsoft Office Live Meeting provided by Microsoft Corporation of Redmond, Washington.

5 Example Appliance Architecture

FIG. 2A also illustrates an example embodiment of the appliance 200. The architecture of the appliance 200 in FIG. 2A is provided by way of illustration only and is not intended to be limiting in any manner. The appliance 200 may include any type and form of computing device 100, such as any element or portion described in conjunction with FIGs.

10 1D and 1E above. In brief overview, the appliance 200 has one or more network ports 266A-226N and one or more networks stacks 267A-267N for receiving and/or transmitting communications via networks 104. The appliance 200 also has a network optimization engine 250 for optimizing, accelerating or otherwise improving the performance, operation, or quality of any network traffic or communications traversing the appliance 200.

15 The appliance 200 includes or is under the control of an operating system. The operating system of the appliance 200 may be any type and/or form of Unix operating system although the invention is not so limited. As such, the appliance 200 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®  
20 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 200 and performing the operations described herein.

25 The operating system of appliance 200 allocates, manages, or otherwise segregates the available system memory into what is referred to as kernel or system space, and user or application space. The kernel space is typically reserved for running the kernel, including any device drivers, kernel extensions or other kernel related software. As known to those skilled in the art, the kernel is the core of the operating system, and provides access, control,  
30 and management of resources and hardware-related elements of the appliance 200. In accordance with an embodiment of the appliance 200, the kernel space also includes a number of network services or processes working in conjunction with the network optimization engine 250, or any portion thereof. Additionally, the embodiment of the kernel will depend on the embodiment of the operating system installed, configured, or otherwise

used by the device 200. In contrast to kernel space, user space 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 directly and uses service calls in order to access kernel services. The operating system uses the user or application space for executing or running applications and provisioning of user level programs, services, processes and/or tasks.

The appliance 200 has one or more network ports 266 for transmitting and receiving data over a network 104. The network port 266 provides a physical and/or logical interface between the computing device and a network 104 or another device 100 for transmitting and receiving network communications. The type and form of network port 266 depends on the type and form of network and type of medium for connecting to the network. Furthermore, any software of, provisioned for or used by the network port 266 and network stack 267 may run in either kernel space or user space.

In one embodiment, the appliance 200 has one network stack 267, such as a TCP/IP based stack, for communicating on a network 105, such with the client 102 and/or the server 106. In one embodiment, the network stack 267 is used to communicate with a first network, such as network 104, and also with a second network 104'. In another embodiment, the appliance 200 has two or more network stacks, such as first network stack 267A and a second network stack 267N. The first network stack 267A may be used in conjunction with a first port 266A to communicate on a first network 104. The second network stack 267N may be used in conjunction with a second port 266N to communicate on a second network 104'. In one embodiment, the network stack(s) 267 has one or more buffers for queuing one or more network packets for transmission by the appliance 200.

The network stack 267 includes any type and form of software, or hardware, or any combinations thereof, for providing connectivity to and communications with a network. In one embodiment, the network stack 267 includes a software implementation for a network protocol suite. The network stack 267 may have one or more network layers, such as any networks layers of the Open Systems Interconnection (OSI) communications model as those skilled in the art recognize and appreciate. As such, the network stack 267 may have any type and form of protocols for any of the following layers of the OSI model: 1) physical link layer, 2) data link layer, 3) network layer, 4) transport layer, 5) session layer, 6) presentation layer, and 7) application layer. In one embodiment, the network stack 267 includes a transport control protocol (TCP) over the network layer protocol of the internet protocol (IP), generally referred to as TCP/IP. In some embodiments, the TCP/IP protocol may be carried

over the Ethernet protocol, which may comprise any of the family of IEEE wide-area-network (WAN) or local-area-network (LAN) protocols, such as those protocols covered by the IEEE 802.3. In some embodiments, the network stack 267 has any type and form of a wireless protocol, such as IEEE 802.11 and/or mobile internet protocol.

5 In view of a TCP/IP based network, any TCP/IP based protocol may be used, 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  
10 Voice Over IP (VoIP) protocol. In another embodiment, the network stack 267 comprises any type and form of transport control protocol, such as a modified transport control protocol, for example a 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. In other embodiments, any type and  
15 form of user datagram protocol (UDP), such as UDP over IP, may be used by the network stack 267, such as for voice communications or real-time data communications.

Furthermore, the network stack 267 may include one or more network drivers supporting the one or more layers, such as a TCP driver or a network layer driver. The network drivers may be included as part of the operating system of the computing device 100  
20 or as part of any network interface cards or other network access components of the computing device 100. In some embodiments, any of the network drivers of the network stack 267 may be customized, modified or adapted to provide a custom or modified portion of the network stack 267 in support of any of the techniques described herein.

In one embodiment, the appliance 200 provides for or maintains a transport layer  
25 connection between a client 102 and server 106 using a single network stack 267. In some embodiments, the appliance 200 effectively terminates the transport layer connection by changing, managing or controlling the behavior of the transport control protocol connection between the client and the server. In these embodiments, the appliance 200 may use a single network stack 267. In other embodiments, the appliance 200 terminates a first transport layer  
30 connection, such as a TCP connection of a client 102, and establishes a second transport layer connection to a server 106 for use by or on behalf of the client 102, e.g., the second transport layer connection is terminated at the appliance 200 and the server 106. The first and second transport layer connections may be established via a single network stack 267. In other embodiments, the appliance 200 may use multiple network stacks, for example 267A and

267N. In these embodiments, the first transport layer connection may be established or terminated at one network stack 267A, and the second transport layer connection may be established or terminated on the second network stack 267N. For example, one network stack may be for receiving and transmitting network packets on a first network, and another network stack for receiving and transmitting network packets on a second network.

As shown in FIG. 2A, the network optimization engine 250 includes one or more of the following elements, components or modules: network packet processing engine 240, LAN/WAN detector 210, flow controller 220, QoS engine 236, protocol accelerator 234, compression engine 238, cache manager 232 and policy engine 295'. The network optimization engine 250, or any portion thereof, may include software, hardware or any combination of software and hardware. Furthermore, any software of, provisioned for or used by the network optimization engine 250 may run in either kernel space or user space. For example, in one embodiment, the network optimization engine 250 may run in kernel space. In another embodiment, the network optimization engine 250 may run in user space. In yet another embodiment, a first portion of the network optimization engine 250 runs in kernel space while a second portion of the network optimization engine 250 runs in user space.

#### Network Packet Processing Engine

The network packet engine 240, also generally referred to as a packet processing engine or packet engine, is responsible for controlling and managing the processing of packets received and transmitted by appliance 200 via network ports 266 and network stack(s) 267. The network packet engine 240 may operate at any layer of the network stack 267. In one embodiment, the network packet engine 240 operates at layer 2 or layer 3 of the network stack 267. In some embodiments, the packet engine 240 intercepts or otherwise receives packets at the network layer, such as the IP layer in a TCP/IP embodiment. In another embodiment, the packet engine 240 operates at layer 4 of the network stack 267. For example, in some embodiments, the packet engine 240 intercepts or otherwise receives packets at the transport layer, such as intercepting packets as the TCP layer in a TCP/IP embodiment. In other embodiments, the packet engine 240 operates at any session or application layer above layer 4. For example, in one embodiment, the packet engine 240 intercepts or otherwise receives network packets above the transport layer protocol layer, such as the payload of a TCP packet in a TCP embodiment.

The packet engine 240 may include a buffer for queuing one or more network packets during processing, such as for receipt of a network packet or transmission of a network packet. Additionally, the packet engine 240 is in communication with one or more network stacks 267 to send and receive network packets via network ports 266. The packet engine 240 may include a packet processing timer. In one embodiment, the packet processing timer 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 packet engine 240 processes network packets responsive to the timer. The packet processing timer 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 operates in the order of milliseconds, such as for example 100ms, 50ms, 25ms, 10ms, 5ms or 1ms.

During operations, the packet engine 240 may be interfaced, integrated or be in communication with any portion of the network optimization engine 250, such as the LAN/WAN detector 210, flow controller 220, QoS engine 236, protocol accelerator 234, compression engine 238, cache manager 232 and/or policy engine 295'. As such, any of the logic, functions, or operations of the LAN/WAN detector 210, flow controller 220, QoS engine 236, protocol accelerator 234, compression engine 238, cache manager 232 and policy engine 295' may be performed responsive to the packet processing timer and/or the packet engine 240. In some embodiments, 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, for example, at a time interval of less than or equal to 10ms. For example, in one embodiment, the cache manager 232 may perform expiration of any cached objects responsive to the 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, such as at every 10 ms.

### Cache Manager

The cache manager 232 may include software, hardware or any combination of software and hardware to store data, information and objects to a cache in memory or storage, provide cache access, and control and manage the cache. The data, objects or content processed and stored by the cache manager 232 may include data in any format, such as a

markup language, or any type of data 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 or storage element. Once the data is stored in the cache, 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 may comprise a data object in memory of the appliance 200. In another embodiment, the cache may comprise any type and form of storage element of the appliance 200, such as a portion of a hard disk. In some embodiments, the processing unit of the device may provide cache memory for use by the cache manager 232. 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.

Furthermore, the cache manager 232 includes any logic, functions, rules, or operations to perform any caching techniques of the appliance 200. In some embodiments, the cache manager 232 may operate as an application, library, program, service, process, thread or task. 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).

### Policy Engine

The policy engine 295' includes any logic, function or operations for providing and applying one or more policies or rules to the function, operation or configuration of any portion of the appliance 200. The policy engine 295' may include, for example, an intelligent statistical engine or other programmable application(s). In one embodiment, the policy engine 295 provides a configuration mechanism to allow a user to identify, specify, define or configure a policy for the network optimization engine 250, or any portion thereof. For example, the policy engine 295 may provide policies for what data to cache, when to cache the data, for whom to cache the data, when to expire an object in cache or refresh the cache. In other embodiments, the policy engine 236 may include any logic, rules, functions or operations to determine and provide access, control and management of objects, data or content being cached by the appliance 200 in addition to access, control and management of security, network traffic, network access, compression or any other function or operation performed by the appliance 200.

In some embodiments, the policy engine 295' provides and applies one or more policies based on any one or more of the following: a user, identification of the client, identification of the server, the type of connection, the time of the connection, the type of network, or the contents of the network traffic. In one embodiment, the policy engine 295' provides and applies a policy based on any field or header at any protocol layer of a network packet. In another embodiment, the policy engine 295' provides and applies a policy based on any payload of a network packet. For example, in one embodiment, the policy engine 295' applies a policy based on identifying a certain portion of content of an application layer protocol carried as a payload of a transport layer packet. In another example, the policy engine 295' applies a policy based on any information identified by a client, server or user certificate. In yet another embodiment, the policy engine 295' applies a policy based on any attributes or characteristics obtained about a client 102, such as via any type and form of endpoint detection (see for example the collection agent of the client agent discussed below).

In one embodiment, the policy engine 295' works in conjunction or cooperation with the policy engine 295 of the application delivery system 290. In some embodiments, the policy engine 295' is a distributed portion of the policy engine 295 of the application delivery system 290. In another embodiment, the policy engine 295 of the application delivery system 290 is deployed on or executed on the appliance 200. In some embodiments, the policy engines 295, 295' both operate on the appliance 200. In yet another embodiment, the policy engine 295', or a portion thereof, of the appliance 200 operates on a server 106.

#### Multi-Protocol and Multi-Layer Compression Engine

The compression engine 238 includes 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 appliance 200. The compression engine 238 may also be referred to as a multi-protocol compression engine 238 in that it may be designed, constructed or capable of compressing a plurality of protocols. In one embodiment, the compression engine 238 applies context insensitive compression, which is compression applied to data without knowledge of the type of data. In another embodiment, the compression engine 238 applies context-sensitive compression. In this embodiment, the compression engine 238 utilizes knowledge of the data type to select a specific compression algorithm from a suite of suitable algorithms. In some embodiments, knowledge of the

specific protocol is used to perform context-sensitive compression. In one embodiment, the appliance 200 or compression engine 238 can use port numbers (e.g., well-known ports), as well as data from the connection itself to determine the appropriate compression algorithm to use. Some protocols use only a single type of data, requiring only a single compression  
5 algorithm that can be selected when the connection is established. Other protocols contain different types of data at different times. For example, POP, IMAP, SMTP, and HTTP all move files of arbitrary types interspersed with other protocol data.

In one embodiment, the compression engine 238 uses a delta-type compression algorithm. In another embodiment, the compression engine 238 uses first site compression as  
10 well as searching for repeated patterns among data stored in cache, memory or disk. In some embodiments, the compression engine 238 uses a lossless compression algorithm. In other embodiments, the compression engine uses a lossy compression algorithm. In some cases, knowledge of the data type and, sometimes, permission from the user are required to use a lossy compression algorithm. Compression is not limited to the protocol payload. The  
15 control fields of the protocol itself may be compressed. In some embodiments, the compression engine 238 uses a different algorithm than that used for the payload.

In some embodiments, the compression engine 238 compresses at one or more layers of the network stack 267. In one embodiment, the compression engine 238 compresses at a transport layer protocol. In another embodiment, the compression engine 238 compresses at  
20 an application layer protocol. In some embodiments, the compression engine 238 compresses at a layer 2-4 protocol. In other embodiments, the compression engine 238 compresses at a layer 5-7 protocol. In yet another embodiment, the compression engine compresses a transport layer protocol and an application layer protocol. In some embodiments, the compression engine 238 compresses a layer 2-4 protocol and a layer 5-7 protocol.

In some embodiments, the compression engine 238 uses memory-based compression, cache-based compression or disk-based compression or any combination thereof. As such,  
25 the compression engine 238 may be referred to as a multi-layer compression engine. In one embodiment, the compression engine 238 uses a history of data stored in memory, such as RAM. In another embodiment, the compression engine 238 uses a history of data stored in a  
30 cache, such as L2 cache of the processor. In other embodiments, the compression engine 238 uses a history of data stored to a disk or storage location. In some embodiments, the compression engine 238 uses a hierarchy of cache-based, memory-based and disk-based data history. The compression engine 238 may first use the cache-based data to determine one or more data matches for compression, and then may check the memory-based data to determine

one or more data matches for compression. In another case, the compression engine 238 may check disk storage for data matches for compression after checking either the cache-based and/or memory-based data history.

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 200 to appliance 200 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.

As such, the multi-protocol compression engine 238 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 by integrating with packet processing engine 240 accessing the network stack 267 is able to compress any of the protocols carried by a transport layer protocol, such as any application layer protocol.

#### LAN/WAN Detector

The LAN/WAN detector 238 includes any logic, business rules, function or operations for automatically detecting a slow side connection (e.g., a wide area network (WAN) connection such as an Intranet) and associated port 267, and a fast side connection (e.g., a local area network (LAN) connection) and an associated port 267. In some embodiments, the LAN/WAN detector 238 monitors network traffic on the network ports 267 of the appliance 200 to detect a synchronization packet, sometimes referred to as a “tagged”

network packet. The synchronization packet identifies a type or speed of the network traffic. In one embodiment, the synchronization packet identifies a WAN speed or WAN type connection. The LAN/WAN detector 238 also identifies receipt of an acknowledgement packet to a tagged synchronization packet and on which port it is received. The appliance  
5 200 then configures itself to operate the identified port on which the tagged synchronization packet arrived so that the speed on that port is set to be the speed associated with the network connected to that port. The other port is then set to the speed associated with the network connected to that port.

For ease of discussion herein, reference to "fast" side will be made with respect to  
10 connection with a wide area network (WAN), e.g., the Internet, and operating at a network speed of the WAN. Likewise, reference to "slow" side will be made with respect to connection with a local area network (LAN) and operating at a network speed the LAN. However, it is noted that "fast" and "slow" sides in a network can change on a per-connection basis and are relative terms to the speed of the network connections or to the type of network  
15 topology. Such configurations are useful in complex network topologies, where a network is "fast" or "slow" only when compared to adjacent networks and not in any absolute sense.

In one embodiment, the LAN/WAN detector 238 may be used to allow for auto-discovery by an appliance 200 of a network to which it connects. In another embodiment, the LAN/WAN detector 238 may be used to detect the existence or presence of a second  
20 appliance 200' deployed in the network 104. For example, an auto-discovery mechanism in operation in accordance with FIG. 1A functions as follows: appliance 200 and 200' are placed in line with the connection linking client 102 and server 106. The appliances 200 and 200' are at the ends of a low-speed link, e.g., Internet, connecting two LANs. In one example embodiment, appliances 200 and 200' each include two ports--one to connect with the  
25 "lower" speed link and the other to connect with a "higher" speed link, e.g., a LAN. Any packet arriving at one port is copied to the other port. Thus, appliance 200 and 200' are each configured to function as a bridge between the two networks 104.

When an end node, such as the client 102, opens a new TCP connection with another end node, such as the server 106, the client 102 sends a TCP packet with a synchronization  
30 (SYN) header bit set, or a SYN packet, to the server 106. In the present example, client 102 opens a transport layer connection to server 106. When the SYN packet passes through appliance 200, the appliance 200 inserts, attaches or otherwise provides a characteristic TCP header option to the packet, which announces its presence. If the packet passes through a second appliance, in this example appliance 200' the second appliance notes the header

option on the SYN packet. The server 106 responds to the SYN packet with a synchronization acknowledgment (SYN-ACK) packet. When the SYN-ACK packet passes through appliance 200', a TCP header option is tagged (e.g., attached, inserted or added) to the SYN-ACK packet to announce appliance 200' presence to appliance 200. When  
5 appliance 200 receives this packet, both appliances 200, 200' are now aware of each other and the connection can be appropriately accelerated.

Further to the operations of the LAN/WAN detector 238, a method or process for detecting "fast" and "slow" sides of a network using a SYN packet is described. During a transport layer connection establishment between a client 102 and a server 106, the appliance  
10 200 via the LAN/WAN detector 238 determines whether the SYN packet is tagged with an acknowledgement (ACK). If it is tagged, the appliance 200 identifies or configures the port receiving the tagged SYN packet (SYN-ACK) as the "slow" side. In one embodiment, the appliance 200 optionally removes the ACK tag from the packet before copying the packet to the other port. If the LAN/WAN detector 238 determines that the packet is not tagged, the  
15 appliance 200 identifies or configures the port receiving the untagged packet as the "fast" side. The appliance 200 then tags the SYN packet with an ACK and copies the packet to the other port.

In another embodiment, the LAN/WAN detector 238 detects fast and slow sides of a network using a SYN-ACK packet. The appliance 200 via the LAN/WAN detector 238  
20 determines whether the SYN-ACK packet is tagged with an acknowledgement (ACK). If it is tagged, the appliance 200 identifies or configures the port receiving the tagged SYN packet (SYN-ACK) as the "slow" side. In one embodiment, the appliance 200 optionally removes the ACK tag from the packet before copying the packet to the other port. If the LAN/WAN detector 238 determines that the packet is not tagged, the appliance 200 identifies or  
25 configures the port receiving the untagged packet as the "fast" side. The LAN/WAN detector 238 determines whether the SYN packet was tagged. If the SYN packet was not tagged, the appliance 200 copied the packet to the other port. If the SYN packet was tagged, the appliance tags the SYN-ACK packet before copying it to the other port.

The appliance 200, 200' may add, insert, modify, attach or otherwise provide any  
30 information or data in the TCP option header to provide any information, data or characteristics about the network connection, network traffic flow, or the configuration or operation of the appliance 200. In this manner, not only does an appliance 200 announce its presence to another appliance 200' or tag a higher or lower speed connection, the appliance 200 provides additional information and data via the TCP option headers about the appliance

or the connection. The TCP option header information may be useful to or used by an appliance in controlling, managing, optimizing, acceleration or improving the network traffic flow traversing the appliance 200, or to otherwise configure itself or operation of a network port.

5           Although generally described in conjunction with detecting speeds of network connections or the presence of appliances, the LAN/WAN detector 238 can be used for applying any type of function, logic or operation of the appliance 200 to a port, connection or flow of network traffic. In particular, automated assignment of ports can occur whenever a device performs different functions on different ports, where the assignment of a port to a  
10 task can be made during the unit's operation, and/or the nature of the network segment on each port is discoverable by the appliance 200.

#### Flow Control

15           The flow controller 220 includes any logic, business rules, function or operations for optimizing, accelerating or otherwise improving the performance, operation or quality of service of transport layer communications of network packets or the delivery of packets at the transport layer. A flow controller, also sometimes referred to as a flow control module, regulates, manages and controls data transfer rates. In some embodiments, the flow  
20 controller 220 is deployed at or connected at a bandwidth bottleneck in the network 104. In one embodiment, the flow controller 220 effectively regulates, manages and controls bandwidth usage or utilization. In other embodiments, the flow control modules may also be deployed at points on the network of latency transitions (low latency to high latency) and on links with media losses (such as wireless or satellite links).

25           In some embodiments, a flow controller 220 may include a receiver-side flow control module for controlling the rate of receipt of network transmissions and a sender-side flow control module for the controlling the rate of transmissions of network packets. In other embodiments, a first flow controller 220 includes a receiver-side flow control module and a second flow controller 220' includes a sender-side flow control module. In some  
30 embodiments, a first flow controller 220 is deployed on a first appliance 200 and a second flow controller 220' is deployed on a second appliance 200'. As such, in some embodiments, a first appliance 200 controls the flow of data on the receiver side and a second appliance 200' controls the data flow from the sender side. In yet another embodiment, a single appliance 200 includes flow control for both the receiver-side and sender-side of network communications traversing the appliance 200.

In one embodiment, a flow control module 220 is configured to allow bandwidth at the bottleneck to be more fully utilized, and in some embodiments, not over utilized. In some embodiments, the flow control module 220 transparently buffers (or rebuffers data already buffered by, for example, the sender) network sessions that pass between nodes having associated flow control modules 220. When a session passes through two or more flow control modules 220, one or more of the flow control modules controls a rate of the session(s).

In one embodiment, the flow control module 200 is configured with predetermined data relating to bottleneck bandwidth. In another embodiment, the flow control module 220 may be configured to detect the bottleneck bandwidth or data associated therewith. Unlike conventional network protocols such as TCP, a receiver-side flow control module 220 controls the data transmission rate. The receiver-side flow control module controls 220 the sender-side flow control module, e.g., 220, data transmission rate by forwarding transmission rate limits to the sender-side flow control module 220. In one embodiment, the receiver-side flow control module 220 piggybacks these transmission rate limits on acknowledgement (ACK) packets (or signals) sent to the sender, e.g., client 102, by the receiver, e.g., server 106. The receiver-side flow control module 220 does this in response to rate control requests that are sent by the sender side flow control module 220'. The requests from the sender-side flow control module 220' may be "piggybacked" on data packets sent by the sender 106.

In some embodiments, the flow controller 220 manipulates, adjusts, simulates, changes, improves or otherwise adapts the behavior of the transport layer protocol to provide improved performance or operations of delivery, data rates and/or bandwidth utilization of the transport layer. The flow controller 220 may implement a plurality of data flow control techniques at the transport layer, including but not limited to 1) pre-acknowledgements, 2) window virtualization, 3) reconnection techniques, 3) local retransmission techniques, 4) wavefront detection and disambiguation, 5) transport control protocol selective acknowledgements, 6) transaction boundary detection techniques and 7) repacketization.

Although a sender may be generally described herein as a client 102 and a receiver as a server 106, a sender may be any end point such as a server 106 or any computing device 100 on the network 104. Likewise, a receiver may be a client 102 or any other computing device on the network 104.

#### Pre-Acknowledgements

In brief overview of a pre-acknowledgement flow control technique, the flow controller 220, in some embodiments, handles the acknowledgements and retransmits for a sender, effectively terminating the sender's connection with the downstream portion of a network connection. In reference to FIG. 1B, one possible deployment of an appliance 200 into a network architecture to implement this feature is depicted. In this example environment, a sending computer or client 102 transmits data on network 104, for example, via a switch, which determines that the data is destined for VPN appliance 205. Because of the chosen network topology, all data destined for VPN appliance 205 traverses appliance 200, so the appliance 200 can apply any necessary algorithms to this data.

Continuing further with the example, the client 102 transmits a packet, which is received by the appliance 200. When the appliance 200 receives the packet, which is transmitted from the client 102 to a recipient via the VPN appliance 205 the appliance 200 retains a copy of the packet and forwards the packet downstream to the VPN appliance 205. The appliance 200 then generates an acknowledgement packet (ACK) and sends the ACK packet back to the client 102 or sending endpoint. This ACK, a pre-acknowledgment, causes the sender 102 to believe that the packet has been delivered successfully, freeing the sender's resources for subsequent processing. The appliance 200 retains the copy of the packet data in the event that a retransmission of the packet is required, so that the sender 102 does not have to handle retransmissions of the data. This early generation of acknowledgements may be called "preacking."

If a retransmission of the packet is required, the appliance 200 retransmits the packet to the sender. The appliance 200 may determine whether retransmission is required as a sender would in a traditional system, for example, determining that a packet is lost if an acknowledgement has not been received for the packet after a predetermined amount of time. To this end, the appliance 200 monitors acknowledgements generated by the receiving endpoint, e.g., server 106 (or any other downstream network entity) so that it can determine whether the packet has been successfully delivered or needs to be retransmitted. If the appliance 200 determines that the packet has been successfully delivered, the appliance 200 is free to discard the saved packet data. The appliance 200 may also inhibit forwarding acknowledgements for packets that have already been received by the sending endpoint.

In the embodiment described above, the appliance 200 via the flow controller 220 controls the sender 102 through the delivery of pre-acknowledgements, also referred to as "preacks", as though the appliance 200 was a receiving endpoint itself. Since the appliance 200 is not an endpoint and does not actually consume the data, the appliance 200 includes a

mechanism for providing overflow control to the sending endpoint. Without overflow control, the appliance 200 could run out of memory because the appliance 200 stores packets that have been preacked to the sending endpoint but not yet acknowledged as received by the receiving endpoint. Therefore, in a situation in which the sender 102 transmits packets to the appliance 200 faster than the appliance 200 can forward the packets downstream, the memory available in the appliance 200 to store unacknowledged packet data can quickly fill. A mechanism for overflow control allows the appliance 200 to control transmission of the packets from the sender 102 to avoid this problem.

In one embodiment, the appliance 200 or flow controller 220 includes an inherent "self-clocking" overflow control mechanism. This self-clocking is due to the order in which the appliance 200 may be designed to transmit packets downstream and send ACKs to the sender 102 or 106. In some embodiments, the appliance 200 does not preack the packet until after it transmits the packet downstream. In this way, the sender 102 will receive the ACKs at the rate at which the appliance 200 is able to transmit packets rather than the rate at which the appliance 200 receives packets from the sender 100. This helps to regulate the transmission of packets from a sender 102.

#### Window Virtualization

Another overflow control mechanism that the appliance 200 may implement is to use the TCP window size parameter, which tells a sender how much buffer the receiver is permitting the sender to fill up. A nonzero window size (e.g., a size of at least one Maximum Segment Size (MSS)) in a preack permits the sending endpoint to continue to deliver data to the appliance, whereas a zero window size inhibits further data transmission. Accordingly, the appliance 200 may regulate the flow of packets from the sender, for example when the appliance's 200 buffer is becoming full, by appropriately setting the TCP window size in each preack.

Another technique to reduce this additional overhead is to apply hysteresis. When the appliance 200 delivers data to the slower side, the overflow control mechanism in the appliance 200 can require that a minimum amount of space be available before sending a nonzero window advertisement to the sender. In one embodiment, the appliance 200 waits until there is a minimum of a predetermined number of packets, such as four packets, of space available before sending a nonzero window packet, such as a window size of four packet). This reduces the overhead by approximately a factor four, since only two ACK

packets are sent for each group of four data packets, instead of eight ACK packets for four data packets.

Another technique the appliance 200 or flow controller 220 may use for overflow control is the TCP delayed ACK mechanism, which skips ACKs to reduce network traffic.

5 The TCP delayed ACKs automatically delay the sending of an ACK, either until two packets are received or until a fixed timeout has occurred. This mechanism alone can result in cutting the overhead in half; moreover, by increasing the numbers of packets above two, additional overhead reduction is realized. But merely delaying the ACK itself may be insufficient to control overflow, and the appliance 200 may also use the advertised window mechanism on  
10 the ACKs to control the sender. When doing this, the appliance 200 in one embodiment avoids triggering the timeout mechanism of the sender by delaying the ACK too long.

In one embodiment, the flow controller 220 does not preack the last packet of a group of packets. By not preacking the last packet, or at least one of the packets in the group, the appliance avoids a false acknowledgement for a group of packets. For example, if the  
15 appliance were to send a preack for a last packet and the packet were subsequently lost, the sender would have been tricked into thinking that the packet is delivered when it was not. Thinking that the packet had been delivered, the sender could discard that data. If the appliance also lost the packet, there would be no way to retransmit the packet to the recipient. By not preacking the last packet of a group of packets, the sender will not discard the packet  
20 until it has been delivered.

In another embodiment, the flow controller 220 may use a window virtualization technique to control the rate of flow or bandwidth utilization of a network connection. Though it may not immediately be apparent from examining conventional literature such as RFC 1323, there is effectively a send window for transport layer protocols such as TCP. The  
25 send window is similar to the receive window, in that it consumes buffer space (though on the sender). The sender's send window consists of all data sent by the application that has not been acknowledged by the receiver. This data must be retained in memory in case retransmission is required. Since memory is a shared resource, some TCP stack implementations limit the size of this data. When the send window is full, an attempt by an  
30 application program to send more data results in blocking the application program until space is available. Subsequent reception of acknowledgements will free send-window memory and unblock the application program. In some embodiments, this window size is known as the socket buffer size in some TCP implementations.

In one embodiment, the flow control module 220 is configured to provide access to increased window (or buffer) sizes. This configuration may also be referenced to as window virtualization. In the embodiment of TCP as the transport layer protocol, the TCP header includes a bit string corresponding to a window scale. In one embodiment, "window" may  
5 be referenced in a context of send, receive, or both.

One embodiment of window virtualization is to insert a preacking appliance 200 into a TCP session. In reference to any of the environments of FIG. 1A or 1B, initiation of a data communication session between a source node, e.g., client 102 (for ease of discussion, now referenced as source node 102), and a destination node, e.g., server 106 (for ease of  
10 discussion, now referenced as destination node 106) is established. For TCP communications, the source node 102 initially transmits a synchronization signal ("SYN") through its local area network 104 to first flow control module 220. The first flow control module 220 inserts a configuration identifier into the TCP header options area. The configuration identifier identifies this point in the data path as a flow control module.

The appliances 200 via a flow control module 220 provide window (or buffer) to  
15 allow increasing data buffering capabilities within a session despite having end nodes with small buffer sizes, e.g., typically 16 k bytes. However, RFC 1323 requires window scaling for any buffer sizes greater than 64 k bytes, which must be set at the time of session initialization (SYN, SYN-ACK signals). Moreover, the window scaling corresponds to the lowest common  
20 denominator in the data path, often an end node with small buffer size. This window scale often is a scale of 0 or 1, which corresponds to a buffer size of up to 64 k or 128 k bytes. Note that because the window size is defined as the window field in each packet shifted over by the window scale, the window scale establishes an upper limit for the buffer, but does not guarantee the buffer is actually that large. Each packet indicates the current available buffer  
25 space at the receiver in the window field.

In one embodiment of scaling using the window virtualization technique, during connection establishment (i.e., initialization of a session) when the first flow control module 220 receives from the source node 102 the SYN signal (or packet), the flow control module 220 stores the windows scale of the source node 102 (which is the previous node) or stores a  
30 0 for window scale if the scale of the previous node is missing. The first flow control module 220 also modifies the scale, e.g., increases the scale to 4 from 0 or 1, in the SYN-FCM signal. When the second flow control module 220 receives the SYN signal, it stores the increased scale from the first flow control signal and resets the scale in the SYN signal back to the source node 103 scale value for transmission to the destination node 106. When the second

flow controller 220 receives the SYN-ACK signal from the destination node 106, it stores the scale from the destination node 106 scale, e.g., 0 or 1, and modifies it to an increased scale that is sent with the SYN-ACK-FCM signal. The first flow control node 220 receives and notes the received window scale and revises the windows scale sent back to the source node 102 back down to the original scale, e.g., 0 or 1. Based on the above window shift conversation during connection establishment, the window field in every subsequent packet, e.g., TCP packet, of the session must be shifted according to the window shift conversion.

The window scale, as described above, expresses buffer sizes of over 64 k and may not be required for window virtualization. Thus, shifts for window scale may be used to express increased buffer capacity in each flow control module 220. This increase in buffer capacity in may be referenced as window (or buffer) virtualization. The increase in buffer size allows greater packet through put from and to the respective end nodes 102 and 106. Note that buffer sizes in TCP are typically expressed in terms of bytes, but for ease of discussion "packets" may be used in the description herein as it relates to virtualization.

By way of example, a window (or buffer) virtualization performed by the flow controller 220 is described. In this example, the source node 102 and the destination node 106 are configured similar to conventional end nodes having a limited buffer capacity of 16 k bytes, which equals approximately 10 packets of data. Typically, an end node 102, 106 must wait until the packet is transmitted and confirmation is received before a next group of packets can be transmitted. In one embodiment, using increased buffer capacity in the flow control modules 220, when the source node 103 transmits its data packets, the first flow control module 220 receives the packets, stores it in its larger capacity buffer, e.g., 512 packet capacity, and immediately sends back an acknowledgement signal indicating receipt of the packets ("REC-ACK") back to the source node 102. The source node 102 can then "flush" its current buffer, load it with 10 new data packets, and transmit those onto the first flow control module 220. Again, the first flow control module 220 transmits a REC-ACK signal back to the source node 102 and the source node 102 flushes its buffer and loads it with 10 more new packets for transmission.

As the first flow control module 220 receives the data packets from the source nodes, it loads up its buffer accordingly. When it is ready the first flow control module 220 can begin transmitting the data packets to the second flow control module 230, which also has an increased buffer size, for example, to receive 512 packets. The second flow control module 220' receives the data packets and begins to transmit 10 packets at a time to the destination node 106. Each REC-ACK received at the second flow control node 220 from the destination

node 106 results in 10 more packets being transmitted to the destination node 106 until all the data packets are transferred. Hence, the present invention is able to increase data transmission throughput between the source node (sender) 102 and the destination node (receiver) 106 by taking advantage of the larger buffer in the flow control modules 220, 220' between the

5 devices.

It is noted that by "preacking" the transmission of data as described previously, a sender (or source node 102) is allowed to transmit more data than is possible without the preacks, thus affecting a larger window size. For example, in one embodiment this technique is effective when the flow control module 220, 220' is located "near" a node (e.g., source

10 node 102 or destination node 106) that lacks large windows.

### Recongestion

Another technique or algorithm of the flow controller 220 is referred to as recongestion. The standard TCP congestion avoidance algorithms are known to perform

15 poorly in the face of certain network conditions, including: large RTTs (round trip times), high packet loss rates, and others. When the appliance 200 detects a congestion condition such as long round trip times or high packet loss, the appliance 200 intervenes, substituting an alternate congestion avoidance algorithm that better suits the particular network condition. In one embodiment, the recongestion algorithm uses preacks to effectively terminate the

20 connection between the sender and the receiver. The appliance 200 then resends the packets from itself to the receiver, using a different congestion avoidance algorithm. Recongestion algorithms may be dependent on the characteristics of the TCP connection. The appliance 200 monitors each TCP connection, characterizing it with respect to the different dimensions, selecting a recongestion algorithm that is appropriate for the current characterization.

25 In one embodiment, upon detecting a TCP connection that is limited by round trip times (RTT), a recongestion algorithm is applied which behaves as multiple TCP connections. Each TCP connection operates within its own performance limit but the aggregate bandwidth achieves a higher performance level. One parameter in this mechanism is the number of parallel connections that are applied (N). Too large a value of N and the

30 connection bundle achieves more than its fair share of bandwidth. Too small a value of N and the connection bundle achieves less than its fair share of bandwidth. One method of establishing "N" relies on the appliance 200 monitoring the packet loss rate, RTT, and packet size of the actual connection. These numbers are plugged into a TCP response curve formula to provide an upper limit on the performance of a single TCP connection in the present

configuration. If each connection within the connection bundle is achieving substantially the same performance as that computed to be the upper limit, then additional parallel connections are applied. If the current bundle is achieving less performance than the upper limit, the number of parallel connections is reduced. In this manner, the overall fairness of the system is maintained since individual connection bundles contain no more parallelism than is required to eliminate the restrictions imposed by the protocol itself. Furthermore, each individual connection retains TCP compliance.

Another method of establishing "N" is to utilize a parallel flow control algorithm such as the TCP "Vegas" algorithm or its improved version "Stabilized Vegas." In this method, the network information associated with the connections in the connection bundle (e.g., RTT, loss rate, average packet size, etc.) is aggregated and applied to the alternate flow control algorithm. The results of this algorithm are in turn distributed among the connections of the bundle controlling their number (i.e., N). Optionally, each connection within the bundle continues using the standard TCP congestion avoidance algorithm.

In another embodiment, the individual connections within a parallel bundle are virtualized, i.e., actual individual TCP connections are not established. Instead the congestion avoidance algorithm is modified to behave as though there were N parallel connections. This method has the advantage of appearing to transiting network nodes as a single connection. Thus the QOS, security and other monitoring methods of these nodes are unaffected by the recongestion algorithm. In yet another embodiment, the individual connections within a parallel bundle are real, i.e., a separate TCP connection is established for each of the parallel connections within a bundle. The congestion avoidance algorithm for each TCP connection need not be modified.

## Retransmission

In some embodiments, the flow controller 220 may apply a local retransmission technique. One reason for implementing preacks is to prepare to transit a high-loss link (e.g., wireless). In these embodiments, the preacking appliance 200 or flow control module 220 is located most beneficially "before" the wireless link. This allows retransmissions to be performed closer to the high loss link, removing the retransmission burden from the remainder of the network. The appliance 200 may provide local retransmission, in which case, packets dropped due to failures of the link are retransmitted directly by the appliance 200. This is advantageous because it eliminates the retransmission burden upon an end node, such as server 106, and infrastructure of any of the networks 104. With appliance 200

providing local retransmissions, the dropped packet can be retransmitted across the high loss link without necessitating a retransmit by an end node and a corresponding decrease in the rate of data transmission from the end node.

Another reason for implementing preacks is to avoid a receive time out (RTO)

5 penalty. In standard TCP there are many situations that result in an RTO, even though a large percentage of the packets in flight were successfully received. With standard TCP algorithms, dropping more than one packet within an RTT window would likely result in a timeout. Additionally, most TCP connections experience a timeout if a retransmitted packet is dropped. In a network with a high bandwidth delay product, even a relatively small packet  
10 loss rate will cause frequent Retransmission timeouts (RTOs). In one embodiment, the appliance 200 uses a retransmit and timeout algorithm to avoid premature RTOs. The appliance 200 or flow controller 220 maintains a count of retransmissions is maintained on a per-packet basis. Each time that a packet is retransmitted, the count is incremented by one and the appliance 200 continues to transmit packets. In some embodiments, only if a packet  
15 has been retransmitted a predetermined number of times is an RTO declared.

#### Wavefront Detection and Disambiguation

In some embodiments, the appliance 200 or flow controller 220 uses wavefront detection and disambiguation techniques in managing and controlling flow of network traffic.

20 In this technique, the flow controller 220 uses transmit identifiers or numbers to determine whether particular data packets need to be retransmitted. By way of example, a sender transmits data packets over a network, where each instance of a transmitted data packet is associated with a transmit number. It can be appreciated that the transmit number for a packet is not the same as the packet's sequence number, since a sequence number references the data  
25 in the packet while the transmit number references an instance of a transmission of that data. The transmit number can be any information usable for this purpose, including a timestamp associated with a packet or simply an increasing number (similar to a sequence number or a packet number). Because a data segment may be retransmitted, different transmit numbers may be associated with a particular sequence number.

30 As the sender transmits data packets, the sender maintains a data structure of acknowledged instances of data packet transmissions. Each instance of a data packet transmission is referenced by its sequence number and transmit number. By maintaining a transmit number for each packet, the sender retains the ordering of the transmission of data packets. When the sender receives an ACK or a SACK, the sender determines the highest

transmit number associated with packets that the receiver indicated has arrived (in the received acknowledgement). Any outstanding unacknowledged packets with lower transmit numbers are presumed lost.

5 In some embodiments, the sender is presented with an ambiguous situation when the arriving packet has been retransmitted: a standard ACK/SACK does not contain enough information to allow the sender to determine which transmission of the arriving packet has triggered the acknowledgement. After receiving an ambiguous acknowledgement, therefore, the sender disambiguates the acknowledgement to associate it with a transmit number. In various embodiments, one or a combination of several techniques may be used to resolve this  
10 ambiguity.

In one embodiment, the sender includes an identifier with a transmitted data packet, and the receiver returns that identifier or a function thereof with the acknowledgement. The identifier may be a timestamp (e.g., a TCP timestamp as described in RFC 1323), a sequential number, or any other information that can be used to resolve between two or more instances  
15 of a packet's transmission. In an embodiment in which the TCP timestamp option is used to disambiguate the acknowledgement, each packet is tagged with up to 32-bits of unique information. Upon receipt of the data packet, the receiver echoes this unique information back to the sender with the acknowledgement. The sender ensures that the originally sent packet and its retransmitted version or versions contain different values for the timestamp  
20 option, allowing it to unambiguously eliminate the ACK ambiguity. The sender may maintain this unique information, for example, in the data structure in which it stores the status of sent data packets. This technique is advantageous because it complies with industry standards and is thus likely to encounter little or no interoperability issues. However, this technique may require ten bytes of TCP header space in some implementations, reducing the effective  
25 throughput rate on the network and reducing space available for other TCP options.

In another embodiment, another field in the packet, such as the IP ID field, is used to disambiguate in a way similar to the TCP timestamp option described above. The sender arranges for the ID field values of the original and the retransmitted version or versions of the packet to have different ID fields in the IP header. Upon reception of the data packet at the  
30 receiver, or a proxy device thereof, the receiver sets the ID field of the ACK packet to a function of the ID field of the packet that triggers the ACK. This method is advantageous, as it requires no additional data to be sent, preserving the efficiency of the network and TCP header space. The function chosen should provide a high degree of likelihood of providing disambiguation. In a preferred embodiment, the sender selects IP ID values with the most

significant bit set to 0. When the receiver responds, the IP ID value is set to the same IP ID value with the most significant bit set to a one.

In another embodiment, the transmit numbers associated with non-ambiguous acknowledgements are used to disambiguate an ambiguous acknowledgement. This technique is based on the principle that acknowledgements for two packets will tend to be received closer in time as the packets are transmitted closer in time. Packets that are not retransmitted will not result in ambiguity, as the acknowledgements received for such packets can be readily associated with a transmit number. Therefore, these known transmit numbers are compared to the possible transmit numbers for an ambiguous acknowledgement received near in time to the known acknowledgement. The sender compares the transmit numbers of the ambiguous acknowledgement against the last known received transmit number, selecting the one closest to the known received transmit number. For example, if an acknowledgement for data packet 1 is received and the last received acknowledgement was for data packet 5, the sender resolves the ambiguity by assuming that the third instance of data packet 1 caused the acknowledgement.

#### Selective Acknowledgements

Another technique of the appliance 200 or flow controller 220 is to implement an embodiment of transport control protocol selective acknowledgements, or TCP SACK, to determine what packets have or have not been received. This technique allows the sender to determine unambiguously a list of packets that have been received by the receiver as well as an accurate list of packets not received. This functionality may be implemented by modifying the sender and/or receiver, or by inserting sender- and receiver-side flow control modules 220 in the network path between the sender and receiver. In reference to FIG. 1A or FIG. 1B, a sender, e.g., client 102, is configured to transmit data packets to the receiver, e.g., server 106, over the network 104. In response, the receiver returns a TCP Selective Acknowledgment option, referred to as SACK packet to the sender. In one embodiment, the communication is bi-directional, although only one direction of communication is discussed here for simplicity. The receiver maintains a list, or other suitable data structure, that contains a group of ranges of sequence numbers for data packets that the receiver has actually received. In some embodiments,, the list is sorted by sequence number in an ascending or descending order. The receiver also maintains a left-off pointer, which comprises a reference into the list and indicates the left-off point from the previously generated SACK packet.

Upon reception of a data packet, the receiver generates and transmits a SACK packet back to the sender. In some embodiments, the SACK packet includes a number of fields, each of which can hold a range of sequence numbers to indicate a set of received data packets. The receiver fills this first field of the SACK packet with a range of sequence numbers that  
5 includes the landing packet that triggered the SACK packet. The remaining available SACK fields are filled with ranges of sequence numbers from the list of received packets. As there are more ranges in the list than can be loaded into the SACK packet, the receiver uses the left-off pointer to determine which ranges are loaded into the SACK packet. The receiver inserts the SACK ranges consecutively from the sorted list, starting from the range referenced  
10 by the pointer and continuing down the list until the available SACK range space in the TCP header of the SACK packet is consumed. The receiver wraps around to the start of the list if it reaches the end. In some embodiments, two or three additional SACK ranges can be added to the SACK range information.

Once the receiver generates the SACK packet, the receiver sends the  
15 acknowledgement back to the sender. The receiver then advances the left-off pointer by one or more SACK range entries in the list. If the receiver inserts four SACK ranges, for example, the left-off pointer may be advanced two SACK ranges in the list. When the advanced left-off pointer reaches at the end of the list, the pointer is reset to the start of the list, effectively wrapping around the list of known received ranges. Wrapping around the list  
20 enables the system to perform well, even in the presence of large losses of SACK packets, since the SACK information that is not communicated due to a lost SACK packet will eventually be communicated once the list is wrapped around.

It can be appreciated, therefore, that a SACK packet may communicate several details about the condition of the receiver. First, the SACK packet indicates that, upon generation of  
25 the SACK packet, the receiver had just received a data packet that is within the first field of the SACK information. Secondly, the second and subsequent fields of the SACK information indicate that the receiver has received the data packets within those ranges. The SACK information also implies that the receiver had not, at the time of the SACK packet's generation, received any of the data packets that fall between the second and subsequent  
30 fields of the SACK information. In essence, the ranges between the second and subsequent ranges in the SACK information are "holes" in the received data, the data therein known not to have been delivered. Using this method, therefore, when a SACK packet has sufficient space to include more than two SACK ranges, the receiver may indicate to the sender a range of data packets that have not yet been received by the receiver.

In another embodiment, the sender uses the SACK packet described above in combination with the retransmit technique described above to make assumptions about which data packets have been delivered to the receiver. For example, when the retransmit algorithm (using the transmit numbers) declares a packet lost, the sender considers the packet to be only  
5 conditionally lost, as it is possible that the SACK packet identifying the reception of this packet was lost rather than the data packet itself. The sender thus adds this packet to a list of potentially lost packets, called the presumed lost list. Each time a SACK packet arrives, the known missing ranges of data from the SACK packet are compared to the packets in the presumed lost list. Packets that contain data known to be missing are declared actually lost  
10 and are subsequently retransmitted. In this way, the two schemes are combined to give the sender better information about which packets have been lost and need to be retransmitted.

#### Transaction Boundary Detection

In some embodiments, the appliance 200 or flow controller 220 applies a technique  
15 referred to as transaction boundary detection. In one embodiment, the technique pertains to ping-pong behaved connections. At the TCP layer, ping-pong behavior is when one communicant – a sender- sends data and then waits for a response from the other communicant – the receiver. Examples of ping-pong behavior include remote procedure call, HTTP and others. The algorithms described above use retransmission timeout (RTO) to  
20 recover from the dropping of the last packet or packets associated with the transaction. Since the TCP RTO mechanism is extremely coarse in some embodiments, for example requiring a minimum one second value in all cases), poor application behavior may be seen in these situations.

In one embodiment, the sender of data or a flow control module 220 coupled to the  
25 sender detects a transaction boundary in the data being sent. Upon detecting a transaction boundary, the sender or a flow control module 220 sends additional packets, whose reception generates additional ACK or SACK responses from the receiver. Insertion of the additional packets is preferably limited to balance between improved application response time and network capacity utilization. The number of additional packets that is inserted may be  
30 selected according to the current loss rate associated with that connection, with more packets selected for connections having a higher loss rate.

One method of detecting a transaction boundary is time based. If the sender has been sending data and ceases, then after a period of time the sender or flow control module 200 declares a transaction boundary. This may be combined with other techniques. For example,

the setting of the PSH (TCP Push) bit by the sender in the TCP header may indicate a transaction boundary. Accordingly, combining the time-based approach with these additional heuristics can provide for more accurate detection of a transaction boundary. In another technique, if the sender or flow control module 220 understands the application protocol, it can parse the protocol data stream and directly determine transaction boundaries. In some embodiment, this last behavior can be used independent of any time-based mechanism.

Responsive to detecting a transaction boundary, the sender or flow control module 220 transmits additional data packets to the receiver to cause acknowledgements therefrom. The additional data packets should therefore be such that the receiver will at least generate an ACK or SACK in response to receiving the data packet. In one embodiment, the last packet or packets of the transaction are simply retransmitted. This has the added benefit of retransmitting needed data if the last packet or packets had been dropped, as compared to merely sending dummy data packets. In another embodiment, fractions of the last packet or packets are sent, allowing the sender to disambiguate the arrival of these packets from their original packets. This allows the receiver to avoid falsely confusing any reordering adaptation algorithms. In another embodiment, any of a number of well-known forward error correction techniques can be used to generate additional data for the inserted packets, allowing for the reconstruction of dropped or otherwise missing data at the receiver.

In some embodiments, the boundary detection technique described herein helps to avoid a timeout when the acknowledgements for the last data packets in a transaction are dropped. When the sender or flow control module 220 receives the acknowledgements for these additional data packets, the sender can determine from these additional acknowledgements whether the last data packets have been received or need to be retransmitted, thus avoiding a timeout. In one embodiment, if the last packets have been received but their acknowledgements were dropped, a flow control module 220 generates an acknowledgement for the data packets and sends the acknowledgement to the sender, thus communicating to the sender that the data packets have been delivered. In another embodiment, if the last packets have not been received, a flow control module 200 sends a packet to the sender to cause the sender to retransmit the dropped data packets.

### Repacketization

In yet another embodiment, the appliance 200 or flow controller 220 applies a repacketization technique for improving the flow of transport layer network traffic. In some embodiments, performance of TCP is proportional to packet size. Thus increasing packet

sizes improves performance unless it causes substantially increased packet loss rates or other nonlinear effects, like IP fragmentation. In general, wired media (such as copper or fibre optics) have extremely low bit-error rates, low enough that these can be ignored. For these media, it is advantageous for the packet size to be the maximum possible before

5 fragmentation occurs (the maximum packet size is limited by the protocols of the underlying transmission media). Whereas for transmission media with higher loss rates (e.g., wireless technologies such as WiFi, etc., or high-loss environments such as power-line networking, etc.), increasing the packet size may lead to lower transmission rates, as media-induced errors cause an entire packet to be dropped (i.e., media-induced errors beyond the capability of the  
10 standard error correcting code for that media), increasing the packet loss rate. A sufficiently large increase in the packet loss rate will actually negate any performance benefit of increasing packet size. In some cases, it may be difficult for a TCP endpoint to choose an optimal packet size. For example, the optimal packet size may vary across the transmission path, depending on the nature of each link.

15 By inserting an appliance 200 or flow control module 220 into the transmission path, the flow controller 220 monitors characteristics of the link and repacketizes according to determined link characteristics. In one embodiment, an appliance 200 or flow controller 220 repacketizes packets with sequential data into a smaller number of larger packets. In another embodiment, an appliance 200 or flow controller 220 repacketizes packets by breaking part a  
20 sequence of large packets into a larger number of smaller packets. In other embodiments, an appliance 200 or flow controller 220 monitors the link characteristics and adjusts the packet sizes through recombination to improve throughput.

### QoS

25 Still referring to FIG. 2A, the flow controller 220, in some embodiments, may include a QoS Engine 236, also referred to as a QoS controller. In another embodiment, the appliance 200 and/or network optimization engine 250 includes the QoS engine 236, for example, separately but in communication with the flow controller 220. The QoS Engine 236 includes any logic, business rules, function or operations for performing one or more Quality  
30 of Service (QoS) techniques improving the performance, operation or quality of service of any of the network connections. In some embodiments, the QoS engine 236 includes network traffic control and management mechanisms that provide different priorities to different users, applications, data flows or connections. In other embodiments, the QoS engine 236 controls, maintains, or assures a certain level of performance to a user,

application, data flow or connection. In one embodiment, the QoS engine 236 controls, maintains or assures a certain portion of bandwidth or network capacity for a user, application, data flow or connection. In some embodiments, the QoS engine 236 monitors the achieved level of performance or the quality of service corresponding to a user, application, data flow or connection, for example, the data rate and delay. In response to monitoring, the QoS engine 236 dynamically controls or adjusts scheduling priorities of network packets to achieve the desired level of performance or quality of service.

In some embodiments, the QoS engine 236 prioritizes, schedules and transmits network packets according to one or more classes or levels of services. In some embodiments, the class or level service may include: 1) best efforts, 2) controlled load, 3) guaranteed or 4) qualitative. For a best efforts class of service, the appliance 200 makes reasonable effort to deliver packets (a standard service level). For a controlled load class of service, the appliance 200 or QoS engine 236 approximates the standard packet error loss of the transmission medium or approximates the behavior of best-effort service in lightly loaded network conditions. For a guaranteed class of service, the appliance 200 or QoS engine 236 guarantees the ability to transmit data at a determined rate for the duration of the connection. For a qualitative class of service, the appliance 200 or QoS engine 236 the qualitative service class is used for applications, users, data flows or connection that require or desire prioritized traffic but cannot quantify resource needs or level of service. In these cases, the appliance 200 or QoS engine 236 determines the class of service or prioritization based on any logic or configuration of the QoS engine 236 or based on business rules or policies. For example, in one embodiment, the QoS engine 236 prioritizes, schedules and transmits network packets according to one or more policies as specified by the policy engine 295, 295'.

## Protocol Acceleration

The protocol accelerator 234 includes any logic, business rules, function or operations for optimizing, accelerating, or otherwise improving the performance, operation or quality of service of one or more protocols. In one embodiment, the protocol accelerator 234 accelerates any application layer protocol or protocols at layers 5-7 of the network stack. In other embodiments, the protocol accelerator 234 accelerates a transport layer or a layer 4 protocol. In one embodiment, the protocol accelerator 234 accelerates layer 2 or layer 3 protocols. In some embodiments, the protocol accelerator 234 is configured, constructed or designed to optimize or accelerate each of one or more protocols according to the type of data, characteristics and/or behavior of the protocol. In another embodiment, the protocol

accelerator 234 is configured, constructed or designed to improve a user experience, response times, network or computer load, and/or network or bandwidth utilization with respect to a protocol.

5 In one embodiment, the protocol accelerator 234 is configured, constructed or designed to minimize the effect of WAN latency on file system access. In some embodiments, the protocol accelerator 234 optimizes or accelerates the use of the CIFS (Common Internet File System) protocol to improve file system access times or access times to data and files. In some embodiments, the protocol accelerator 234 optimizes or accelerates the use of the NFS (Network File System) protocol. In another embodiment, the protocol  
10 accelerator 234 optimizes or accelerates the use of the File Transfer protocol (FTP).

In one embodiment, the protocol accelerator 234 is configured, constructed or designed to optimize or accelerate a protocol carrying as a payload or using any type and form of markup language. In other embodiments, the protocol accelerator 234 is configured, constructed or designed to optimize or accelerate a HyperText Transfer Protocol (HTTP). In  
15 another embodiment, the protocol accelerator 234 is configured, constructed or designed to optimize or accelerate a protocol carrying as a payload or otherwise using XML (eXtensible Markup Language).

#### Transparency and Multiple Deployment Configuration

20 In some embodiments, the appliance 200 and/or network optimization engine 250 is transparent to any data flowing across a network connection or link, such as a WAN link. In one embodiment, the appliance 200 and/or network optimization engine 250 operates in such a manner that the data flow across the WAN is recognizable by any network monitoring, QOS management or network analysis tools. In some embodiments, the appliance 200 and/or  
25 network optimization engine 250 does not create any tunnels or streams for transmitting data that may hide, obscure or otherwise make the network traffic not transparent. In other embodiments, the appliance 200 operates transparently in that the appliance does not change any of the source and/or destination address information or port information of a network packet, such as internet protocol addresses or port numbers. In other embodiments, the  
30 appliance 200 and/or network optimization engine 250 is considered to operate or behave transparently to the network, an application, client, server or other appliances or computing device in the network infrastructure. That is, in some embodiments, the appliance is transparent in that network related configuration of any device or appliance on the network does not need to be modified to support the appliance 200.

The appliance 200 may be deployed in any of the following deployment configurations: 1) in-line of traffic, 2) in proxy mode, or 3) in a virtual in-line mode. In some embodiments, the appliance 200 may be deployed inline to one or more of the following: a router, a client, a server or another network device or appliance. In other embodiments, the appliance 200 may be deployed in parallel to one or more of the following: : a router, a client, a server or another network device or appliance. In parallel deployments, a client, server, router or other network appliance may be configured to forward, transfer or transit networks to or via the appliance 200.

In the embodiment of in-line, the appliance 200 is deployed inline with a WAN link of a router. In this way, all traffic from the WAN passes through the appliance before arriving at a destination of a LAN.

In the embodiment of a proxy mode, the appliance 200 is deployed as a proxy device between a client and a server. In some embodiments, the appliance 200 allows clients to make indirect connections to a resource on a network. For example, a client connects to a resource via the appliance 200, and the appliance provides the resource either by connecting to the resource, a different resource, or by serving the resource from a cache. In some cases, the appliance may alter the client's request or the server's response for various purposes, such as for any of the optimization techniques discussed herein. In other embodiments, the appliance 200 behaves as a transparent proxy, by intercepting and forwarding requests and responses transparently to a client and/or server. Without client-side configuration, the appliance 200 may redirect client requests to different servers or networks. In some embodiments, the appliance 200 may perform any type and form of network address translation, referred to as NAT, on any network traffic traversing the appliance.

In some embodiments, the appliance 200 is deployed in a virtual in-line mode configuration. In this embodiment, a router or a network device with routing or switching functionality is configured to forward, reroute or otherwise provide network packets destined to a network to the appliance 200. The appliance 200 then performs any desired processing on the network packets, such as any of the WAN optimization techniques discussed herein. Upon completion of processing, the appliance 200 forwards the processed network packet to the router to transmit to the destination on the network. In this way, the appliance 200 can be coupled to the router in parallel but still operate as it if the appliance 200 were inline. This deployment mode also provides transparency in that the source and destination addresses and port information are preserved as the packet is processed and transmitted via the appliance through the network.

### End Node Deployment

Although the network optimization engine 250 is generally described above in conjunction with an appliance 200, the network optimization engine 250, or any portion thereof, may be deployed, distributed or otherwise operated on any end node, such as a client 102 and/or server 106. As such, a client or server may provide any of the systems and methods of the network optimization engine 250 described herein in conjunction with one or more appliances 200 or without an appliance 200.

Referring now to FIG. 2B, an example embodiment of the network optimization engine 250 deployed on one or more end nodes is depicted. In brief overview, the client 102 may include a first network optimization engine 250' and the server 106 may include a second network optimization engine 250''. The client 102 and server 106 may establish a transport layer connection and exchange communications with or without traversing an appliance 200.

In one embodiment, the network optimization engine 250' of the client 102 performs the techniques described herein to optimize, accelerate or otherwise improve the performance, operation or quality of service of network traffic communicated with the server 106. In another embodiment, the network optimization engine 250'' of the server 106 performs the techniques described herein to optimize, accelerate or otherwise improve the performance, operation or quality of service of network traffic communicated with the client 102. In some embodiments, the network optimization engine 250' of the client 102 and the network optimization engine 250'' of the server 106 perform the techniques described herein to optimize, accelerate or otherwise improve the performance, operation or quality of service of network traffic communicated between the client 102 and the server 106. In yet another embodiment, the network optimization engine 250' of the client 102 performs the techniques described herein in conjunction with an appliance 200 to optimize, accelerate or otherwise improve the performance, operation or quality of service of network traffic communicated with the client 102. In still another embodiment, the network optimization engine 250'' of the server 106 performs the techniques described herein in conjunction with an appliance 200 to optimize, accelerate or otherwise improve the performance, operation or quality of service of network traffic communicated with the server 106.

### C. Client Agent

Referring now to FIG. 3, an embodiment of a client agent 120 is depicted. The client 102 has a client agent 120 for establishing, exchanging, managing or controlling communications with the appliance 200, appliance 205 and/or server 106 via a network 104. In some embodiments, the client agent 120, which may also be referred to as a WAN client, accelerates WAN network communications and/or is used to communicate via appliance 200 on a network. In brief overview, the client 102 operates on computing device 100 having an operating system with a kernel mode 302 and a user mode 303, and a network stack 267 with one or more layers 310a-310b. The client 102 may have installed and/or execute one or more applications. In some embodiments, one or more applications may communicate via the network stack 267 to a network 104. One of the applications, such as a web browser, may also include a first program 322. For example, the first program 322 may be used in some embodiments to install and/or execute the client agent 120, or any portion thereof. The client agent 120 includes an interception mechanism, or interceptor 350, for intercepting network communications from the network stack 267 from the one or more applications.

As with the appliance 200, the client has a network stack 267 including any type and form of software, hardware, or any combinations thereof, for providing connectivity to and communications with a network 104. The network stack 267 of the client 102 includes any of the network stack embodiments described above in conjunction with the appliance 200. In some embodiments, the client agent 120, or any portion thereof, is designed and constructed to operate with or work in conjunction with the network stack 267 installed or otherwise provided by the operating system of the client 102.

In further details, the network stack 267 of the client 102 or appliance 200 (or 205) may include any type and form of interfaces for receiving, obtaining, providing or otherwise accessing any information and data related to network communications of the client 102. In one embodiment, an interface to the network stack 267 includes an application programming interface (API). The interface may also have any function call, hooking or filtering mechanism, event or call back mechanism, or any type of interfacing technique. The network stack 267 via the interface may receive or provide any type and form of data structure, such as an object, related to functionality or operation of the network stack 267. For example, the data structure may include information and data related to a network packet or one or more network packets. In some embodiments, the data structure includes, references or identifies a portion of the network packet processed at a protocol layer of the network stack 267, such as a network packet of the transport layer. In some embodiments, the data structure 325 is a kernel-level data structure, while in other embodiments, the data structure 325 is a user-mode

data structure. A kernel-level data structure may have a data structure obtained or related to a portion of the network stack 267 operating in kernel-mode 302, or a network driver or other software running in kernel-mode 302, or any data structure obtained or received by a service, process, task, thread or other executable instructions running or operating in kernel-mode of the operating system.

Additionally, some portions of the network stack 267 may execute or operate in kernel-mode 302, for example, the data link or network layer, while other portions execute or operate in user-mode 303, such as an application layer of the network stack 267. For example, a first portion 310a of the network stack may provide user-mode access to the network stack 267 to an application while a second portion 310a of the network stack 267 provides access to a network. In some embodiments, a first portion 310a of the network stack has one or more upper layers of the network stack 267, such as any of layers 5-7. In other embodiments, a second portion 310b of the network stack 267 includes one or more lower layers, such as any of layers 1-4. Each of the first portion 310a and second portion 310b of the network stack 267 may include any portion of the network stack 267, at any one or more network layers, in user-mode 303, kernel-mode, 302, or combinations thereof, or at any portion of a network layer or interface point to a network layer or any portion of or interface point to the user-mode 302 and kernel-mode 203.

The interceptor 350 may include software, hardware, or any combination of software and hardware. In one embodiment, the interceptor 350 intercepts or otherwise receives a network communication at any point in the network stack 267, and redirects or transmits the network communication to a destination desired, managed or controlled by the interceptor 350 or client agent 120. For example, the interceptor 350 may intercept a network communication of a network stack 267 of a first network and transmit the network communication to the appliance 200 for transmission on a second network 104. In some embodiments, the interceptor 350 includes or is a driver, such as a network driver constructed and designed to interface and work with the network stack 267. In some embodiments, the client agent 120 and/or interceptor 350 operates at one or more layers of the network stack 267, such as at the transport layer. In one embodiment, the interceptor 350 includes a filter driver, hooking mechanism, or any form and type of suitable network driver interface that interfaces to the transport layer of the network stack, such as via the transport driver interface (TDI). In some embodiments, the interceptor 350 interfaces to a first protocol layer, such as the transport layer and another protocol layer, such as any layer above the transport protocol layer, for example, an application protocol layer. In one embodiment, the interceptor 350

includes a driver complying with the Network Driver Interface Specification (NDIS), or a NDIS driver. In another embodiment, the interceptor 350 may be a min-filter or a mini-port driver. In one embodiment, the interceptor 350, or portion thereof, operates in kernel-mode 202. In another embodiment, the interceptor 350, or portion thereof, operates in user-mode 203. In some embodiments, a portion of the interceptor 350 operates in kernel-mode 202 while another portion of the interceptor 350 operates in user-mode 203. In other embodiments, the client agent 120 operates in user-mode 203 but interfaces via the interceptor 350 to a kernel-mode driver, process, service, task or portion of the operating system, such as to obtain a kernel-level data structure 225. In further embodiments, the interceptor 350 is a user-mode application or program, such as application.

In one embodiment, the interceptor 350 intercepts or receives any transport layer connection requests. In these embodiments, the interceptor 350 executes transport layer application programming interface (API) calls to set the destination information, such as destination IP address and/or port to a desired location for the location. In this manner, the interceptor 350 intercepts and redirects the transport layer connection to an IP address and port controlled or managed by the interceptor 350 or client agent 120. In one embodiment, the interceptor 350 sets the destination information for the connection to a local IP address and port of the client 102 on which the client agent 120 is listening. For example, the client agent 120 may comprise a proxy service listening on a local IP address and port for redirected transport layer communications. In some embodiments, the client agent 120 then communicates the redirected transport layer communication to the appliance 200.

In some embodiments, the interceptor 350 intercepts a Domain Name Service (DNS) request. In one embodiment, the client agent 120 and/or interceptor 350 resolves the DNS request. In another embodiment, the interceptor transmits the intercepted DNS request to the appliance 200 for DNS resolution. In one embodiment, the appliance 200 resolves the DNS request and communicates the DNS response to the client agent 120. In some embodiments, the appliance 200 resolves the DNS request via another appliance 200' or a DNS server 106.

In yet another embodiment, the client agent 120 may include two agents 120 and 120'. In one embodiment, a first agent 120 may include an interceptor 350 operating at the network layer of the network stack 267. In some embodiments, the first agent 120 intercepts network layer requests such as Internet Control Message Protocol (ICMP) requests (e.g., ping and traceroute). In other embodiments, the second agent 120' may operate at the transport layer and intercept transport layer communications. In some embodiments, the first agent

120 intercepts communications at one layer of the network stack 210 and interfaces with or communicates the intercepted communication to the second agent 120'.

The client agent 120 and/or interceptor 350 may operate at or interface with a protocol layer in a manner transparent to any other protocol layer of the network stack 267. For example, in one embodiment, the interceptor 350 operates or interfaces with the transport layer of the network stack 267 transparently to any protocol layer below the transport layer, such as the network layer, and any protocol layer above the transport layer, such as the session, presentation or application layer protocols. This allows the other protocol layers of the network stack 267 to operate as desired and without modification for using the interceptor 350. As such, the client agent 120 and/or interceptor 350 can interface with the transport layer to secure, optimize, accelerate, route or load-balance any communications provided via any protocol carried by the transport layer, such as any application layer protocol over TCP/IP.

Furthermore, the client agent 120 and/or interceptor 350 may operate at or interface with the network stack 267 in a manner transparent to any application, a user of the client 102, the client 102 and/or any other computing device 100, such as a server or appliance 200, 206, in communications with the client 102. The client agent 120, or any portion thereof, may be installed and/or executed on the client 102 in a manner without modification of an application. In one embodiment, the client agent 120, or any portion thereof, is installed and/or executed in a manner transparent to any network configuration of the client 102, appliance 200, 205 or server 106. In some embodiments, the client agent 120, or any portion thereof, is installed and/or executed with modification to any network configuration of the client 102, appliance 200, 205 or server 106. In one embodiment, the user of the client 102 or a computing device in communications with the client 102 are not aware of the existence, execution or operation of the client agent 12, or any portion thereof. As such, in some embodiments, the client agent 120 and/or interceptor 350 is installed, executed, and/or operated transparently to an application, user of the client 102, the client 102, another computing device, such as a server or appliance 200, 2005, or any of the protocol layers above and/or below the protocol layer interfaced to by the interceptor 350.

The client agent 120 includes a streaming client 306, a collection agent 304, SSL VPN agent 308, a network optimization engine 250, and/or acceleration program 302. In one embodiment, the client agent 120 is an Independent Computing Architecture (ICA) client, or any portion thereof, developed by Citrix Systems, Inc. of Fort Lauderdale, Florida, and is also referred to as an ICA client. In some embodiments, the client agent 120 has an application

streaming client 306 for streaming an application from a server 106 to a client 102. In another embodiment, the client agent 120 includes a collection agent 304 for performing end-point detection/scanning and collecting end-point information for the appliance 200 and/or server 106. In some embodiments, the client agent 120 has one or more network accelerating or optimizing programs or agents, such as a network optimization engine 250 and an acceleration program 302. In one embodiment, the acceleration program 302 accelerates communications between client 102 and server 106 via appliance 205'. In some embodiments, the network optimization engine 250 provides WAN optimization techniques as discussed herein.

The streaming client 306 is an application, program, process, service, task or set of executable instructions for receiving and executing a streamed application from a server 106. A server 106 may stream one or more application data files to the streaming client 306 for playing, executing or otherwise causing to be executed the application on the client 102. In some embodiments, the server 106 transmits a set of compressed or packaged application data files to the streaming client 306. In some embodiments, the plurality of application files are compressed and stored on a file server within an archive file such as a CAB, ZIP, SIT, TAR, JAR or other archive. In one embodiment, the server 106 decompresses, unpackages or unarchives the application files and transmits the files to the client 102. In another embodiment, the client 102 decompresses, unpackages or unarchives the application files. The streaming client 306 dynamically installs the application, or portion thereof, and executes the application. In one embodiment, the streaming client 306 may be an executable program. In some embodiments, the streaming client 306 may be able to launch another executable program.

The collection agent 304 is an application, program, process, service, task or set of executable instructions for identifying, obtaining and/or collecting information about the client 102. In some embodiments, the appliance 200 transmits the collection agent 304 to the client 102 or client agent 120. The collection agent 304 may be configured according to one or more policies of the policy engine 236 of the appliance. In other embodiments, the collection agent 304 transmits collected information on the client 102 to the appliance 200. In one embodiment, the policy engine 236 of the appliance 200 uses the collected information to determine and provide access, authentication and authorization control of the client's connection to a network 104.

In one embodiment, the collection agent 304 is an end-point detection and scanning program, which identifies and determines one or more attributes or characteristics of the

client. For example, the collection agent 304 may identify and determine any one or more of the following client-side attributes: 1) the operating system an/or a version of an operating system, 2) a service pack of the operating system, 3) a running service, 4) a running process, and 5) a file. The collection agent 304 may also identify and determine the presence or  
5 version of any one or more of the following on the client: 1) antivirus software, 2) personal firewall software, 3) anti-spam software, and 4) internet security software. The policy engine 236 may have one or more policies based on any one or more of the attributes or characteristics of the client or client-side attributes.

The SSL VPN agent 308 is an application, program, process, service, task or set of  
10 executable instructions for establishing a Secure Socket Layer (SSL) virtual private network (VPN) connection from a first network 104 to a second network 104', 104'', or a SSL VPN connection from a client 102 to a server 106. In one embodiment, the SSL VPN agent 308 establishes a SSL VPN connection from a public network 104 to a private network 104' or 104''. In some embodiments, the SSL VPN agent 308 works in conjunction with appliance  
15 205 to provide the SSL VPN connection. In one embodiment, the SSL VPN agent 308 establishes a first transport layer connection with appliance 205. In some embodiment, the appliance 205 establishes a second transport layer connection with a server 106. In another embodiment, the SSL VPN agent 308 establishes a first transport layer connection with an application on the client, and a second transport layer connection with the appliance 205. In  
20 other embodiments, the SSL VPN agent 308 works in conjunction with WAN optimization appliance 200 to provide SSL VPN connectivity.

In some embodiments, the acceleration program 302 is a client-side acceleration program for performing one or more acceleration techniques to accelerate, enhance or otherwise improve a client's communications with and/or access to a server 106, such as  
25 accessing an application provided by a server 106. The logic, functions, and/or operations of the executable instructions of the acceleration program 302 may perform one or more of the following acceleration techniques: 1) multi-protocol compression, 2) transport control protocol pooling, 3) transport control protocol multiplexing, 4) transport control protocol buffering, and 5) caching via a cache manager. Additionally, the acceleration program 302  
30 may perform encryption and/or decryption of any communications received and/or transmitted by the client 102. In some embodiments, the acceleration program 302 performs one or more of the acceleration techniques in an integrated manner or fashion. Additionally, the acceleration program 302 can perform compression on any of the protocols, or multiple-protocols, carried as a payload of a network packet of the transport layer protocol.

In one embodiment, the acceleration program 302 is designed, constructed or configured to work with appliance 205 to provide LAN side acceleration or to provide acceleration techniques provided via appliance 205. For example, in one embodiment of a NetScaler appliance 205 manufactured by Citrix Systems, Inc., the acceleration program 302 includes a NetScaler client. In some embodiments, the acceleration program 302 provides NetScaler acceleration techniques stand-alone in a remote device, such as in a branch office. In other embodiments, the acceleration program 302 works in conjunction with one or more NetScaler appliances 205. In one embodiment, the acceleration program 302 provides LAN-side or LAN based acceleration or optimization of network traffic.

In some embodiments, the network optimization engine 250 may be designed, constructed or configured to work with WAN optimization appliance 200. In other embodiments, network optimization engine 250 may be designed, constructed or configured to provide the WAN optimization techniques of appliance 200, with or without an appliance 200. For example, in one embodiment of a WANScaler appliance 200 manufactured by Citrix Systems, Inc. the network optimization engine 250 includes the WANScaler client. In some embodiments, the network optimization engine 250 provides WANScaler acceleration techniques stand-alone in a remote location, such as a branch office. In other embodiments, the network optimization engine 250 works in conjunction with one or more WANScaler appliances 200.

In another embodiment, the network optimization engine 250 includes the acceleration program 302, or the function, operations and logic of the acceleration program 302. In some embodiments, the acceleration program 302 includes the network optimization engine 250 or the function, operations and logic of the network optimization engine 250. In yet another embodiment, the network optimization engine 250 is provided or installed as a separate program or set of executable instructions from the acceleration program 302. In other embodiments, the network optimization engine 250 and acceleration program 302 are included in the same program or same set of executable instructions.

In some embodiments and still referring to FIG. 3, a first program 322 may be used to install and/or execute the client agent 120, or any portion thereof, automatically, silently, transparently, or otherwise. In one embodiment, the first program 322 is a plugin component, such an ActiveX control or Java control or script that is loaded into and executed by an application. For example, the first program comprises an ActiveX control loaded and run by a web browser application, such as in the memory space or context of the application. In another embodiment, the first program 322 comprises a set of executable instructions loaded

into and run by the application, such as a browser. In one embodiment, the first program 322 is designed and constructed program to install the client agent 120. In some embodiments, the first program 322 obtains, downloads, or receives the client agent 120 via the network from another computing device. In another embodiment, the first program 322 is an installer program or a plug and play manager for installing programs, such as network drivers and the client agent 120, or any portion thereof, on the operating system of the client 102.

In some embodiments, each or any of the portions of the client agent 120 - a streaming client 306, a collection agent 304, SSL VPN agent 308, a network optimization engine 250, acceleration program 302, and interceptor 350 - may be installed, executed, configured or operated as a separate application, program, process, service, task or set of executable instructions. In other embodiments, each or any of the portions of the client agent 120 may be installed, executed, configured or operated together as a single client agent 120.

#### D. Systems and Methods for QoS for Single Connection having Different Priority Packets

Referring now to FIGs. 4A and 4B, embodiments of the present solution for providing quality of service (QoS) for connections with multiple prioritized data streams are depicted.

Embodiments of the present solution provide connection based priority by performing filtering or averaging of a window of protocol data units (PDUs) or stream of bytes traversing the connection. Based on the filtering or averaging, a priority is set for the connection. Although the underlying PDUs may have a different priority than the one set for the connection, the connection's priority is set to the average priority of such PDUs. The system provides connection based priority to each of the connections traversing an intermediary. As the priority of PDUs or bytes being communicates via a connection change, the system may dynamically change the connection's priority. The system may change the connections' priority slowly or quickly or in any rate controlled manner. The system may change the connections' priority dynamically or adaptively based on other factors, such as link utilization.

Referring now to FIG. 4A, an embodiment of a system is depicted for providing QoS for connections which may carry different priority data streams within the connection. In brief overview, an intermediary such as an intermediary device (e.g. an appliance 200) or

agent 120 or on a client, may receive or intercept data streams via a plurality of connections 410A-410N (generally referred to as connection 410). The connection may be any type and form of transport layer connection, such as a TCP connection or UDP connection. The connection may traverse one or more networks 104. The connection may comprise a plurality of connections: a first connection between the intermediary and one end point and a second connection between the intermediary and another end point. The transport layer connection may carry one or more protocols above the transport layer, such as any application layer and/or session protocol. For example, the transport layer connection may carry a remote display protocol, such as ICA or RDP. Each protocol layer may comprise data and/or fields or protocol structure which may at times be referred to as a protocol data unit (PDU). In some embodiments, the sender, receiver and/or intermediary may encrypt these PDUs or tunnel the PDUs over the transport layer connection.

Any protocol carried by the transport connection 410 may identify a priority for the PDUs or stream of data in the protocol. In some embodiments, the sender of the protocol data may mark, set or insert a value into a field of the PDU to identify a priority. The protocol of the PDU may use channels to communicate data having different priorities over the same connection. Each channel have a different predetermined priority or may otherwise be set to or assigned a different priority. A structure of the protocol, such as a header or field, may be used to identify the channel and/or priority. In some embodiments, each PDU for the same protocol may have the same or different priority than another PDU for the same protocol over the same connection. As such, in some embodiments, a stream of PDUs communicated over the same connection have varying priorities or priorities that change over time. Each connection 410A-410N may carry packets having PDUs with the same or different priorities. A single connection 410N for example may carry a multi-channel protocol, such as ICA, in which each channel carries data having different priorities from another channel.

The intermediary includes a network optimization engine 250 with a QoS engine 236. The filter 420 of the QoS engine applies a filter 420A-420N based on a filter window 422 to the priorities of the data streams received via each of the connections 410A-410N. The QoS engine sets a connection priority 430A-430N to each of the connections based on the filtering which may produce an average value over the priorities in the window. In setting the connection priority, the QoS engine may set any fields or protocol structure at any layer in the network traffic of the connection to identify the set priority, such as via any ToC/Diff bits in headers of the packets. The QoS engine may adaptively change a connection's priority

based on changes in the filtered priorities and do so adaptively to link utilization 424. The QoS may apply any rate limit 425 controls to the rate of change of any connection's priority.

In further details, the intermediary may comprise any embodiments of the appliance 200, 205' described herein. In some embodiments, the intermediary is a WAN optimization device. In some embodiments, the intermediary may comprise any embodiments of the client agent 120 described herein. In some embodiments, the agent may execute on a sender. In some embodiments, the agent may execute on a receiver. In some embodiments, the intermediary may comprise any embodiments of the network optimization 250 described herein. In some embodiments, the intermediary may proxy transport layer connections. In some embodiments, the intermediary may terminate a transport layer connection from one end point such as a sender and establish a second layer connection to the receiver. In these embodiments, the intermediary may split the sender to receiver's transport layer connection into multiple transport layer connections. The intermediary may split these connections transparently to the sender and/or receiver, such as a client and a server.

The intermediary may establish, proxy or otherwise process a plurality of transport layer connections 410A-410N. Within any one connection, the data stream may comprise data or PDUs of different priorities. In some embodiments, one or more packets, or any data or PDU thereof, traversing the connection may have or be set to priority A, while one or more other packets, or any data or PDU thereof, traversing the same connection may have or be set to priority B, and so on. In some embodiments, protocol data units carried via packets within the same connection may have differing priorities. In some embodiments, a first stream or group of bytes within the connection may have one priority while another or subsequent stream or group of bytes within the same connection has a different priority. In some embodiments, the protocol carried by the connection may comprises a plurality of virtual streams or channels. In some embodiments, the protocol is ICA. Each channel of the protocol may be assigned or use one or more priorities. Each channel within the connection may have a different priority than another channel. The priority for any channel, PDU or stream or group of bytes may be set or identified via any tag, flag, field or value within a header or field of a packet. The priority for any channel, PDU or stream/group of bytes may be set or identified via any tag, flag, field or value within the body or protocol of a packet. In some embodiments, the priority may be negotiated between a sender and receiver. In some embodiments, the priority may be negotiated between a sender and the intermediary and/or between the receiver and the intermediary.

The priority range used for any channel, PDU or stream/group of bytes may comprise any range. In some embodiments, the priority may be numerical based, such as a range of 0-X, 1-N, etc. In some embodiments, the priority may be string based, such as Low, Medium or High. The priority may be coarse with a couple or more priority levels or may be granular with many priority levels. The priority may correspond to the type or name of a channel. In some embodiments, the priority may be predetermined based on the type or name of the channel.

Multiple priority data streams, either via a channel, PDUs or bytes, may be time and/or space multiplexed within a single transport layer connection. In some embodiments, the priority of any channel, PDUs or bytes change over time in any degree or granularity. In some embodiments, the priority of any channel, PDUs or bytes may remain constant over a period of time. In some embodiments, the priority of any channel, PDUs or bytes may change but remain consistent over a period of time. In some embodiments, the priority may change after each PDU or payload of a packet. In some embodiments, the priority may change after a series of PDUs, payload of packets or packets. In some embodiments, the priority may not change for after each PDU or payload of a packet. In some embodiments, the priority may not change after a series of PDUs, payload of packets or packets.

The optimization engine 250 of the intermediary may perform any optimizations described herein on the network traffic of connections 410A-410N. In some embodiments, the optimization engine includes a QoS engine 236. In some embodiments, the QoS engine includes any embodiments described in connection with FIG. 2A. The QoS engine may include executable instructions to perform on the device of the intermediary any of the filtering and connection based prioritizing described herein. The QoS engine may filter the differing priorities of data streams within each connection to determine a priority 430A-430N for the connection. The intermediary may store or buffer the data stream and then transmit the buffered data stream in accordance with the priority 430A-430N for the connection. The QoS engine may use a defined window of space and/or time over which to filter priorities. As the priorities within each connection change over time and space, the QoS may adapt or change the priority 430A-430N established for each connection.

The QoS engine may identify, extract or determine a priority for each packet, PDU or group of bytes. The QoS may store the priority in a storage or memory element of the device, such as a buffer. The QoS may store the priorities for any packet, PDU or group of bytes received by the intermediary. The QoS may store the priorities for any packet, PDU or group of bytes currently in flight between the sender and receiver. The QoS may store the priorities

for any packet, PDU or group of bytes waiting to be transmitted by the intermediary. The QoS may store the priorities for any packet, PDU or group of bytes within a current window. The QoS may store the priorities for any packet, PDU or group of bytes within a previous window. The QoS may store the priorities for any packet, PDU or group of bytes for a plurality of windows. The QoS may store the priorities for a predetermined amount of time. The QoS may store the priorities for a predetermined number of packets, PDUs or bytes. The QoS engine may store the priority as a state for a PDU, packet or group of bytes. In some embodiments, the system does not store PDUs, packets or group of bytes but just their priorities. In some embodiments, for example one with 8 different priorities, a 3 bit state per PDU may be used.

The QoS engine may comprise any type and form of filter 420. The filter may be channel, PDU or byte based. The filter may comprise logic, operations or functions to perform a mathematical or statistical operation on priorities, such as priorities identified and stored by the intermediary. The filter may take an average of priority values of a plurality of PDUs for a connection. The filter may take an average of priority values of a stream of bytes for a connection. The filter may take an average of priority values across a plurality of channels traversing the connection. The filter may take an average of priority values of a plurality of PDUs, channels or bytes within a window 422. The filter may take a moving average of priority values. The filter may take a weighed average of priority values. In some embodiments, the weightings may be based on number of bytes or PDUs with the assigned priority. In some embodiments, the weightings may be based on number of packets carrying the assigned priority within the window. In some embodiments, the weightings may be based on number of PDUs for a specific channel are traversing the connection, such as within the window.

The filter may comprise any type and form of digital filter. The filter comprise any type and form of analog filter. In some embodiments, the filter comprises a Finite Impulse Response (FIR) filter. In some embodiments, the filter comprises a box car filter. In some embodiments, the filter comprises an Infinite Impulse Response (IIR) filter. In some embodiments, the filter comprises executable instructions designed and constructed to perform an averaging function or filter to an input or inputs, such as a plurality of priorities, and provides as output an average value. In some embodiments, the filter may comprise a hardware device or electronic circuit that performs an averaging function or filter to an input or inputs, such as a plurality of priorities, and provides as output an average value in

accordance with the type of filtering or averaging the filter is designed and/or constructed to perform.

The window 422, sometimes referred to as an averaging window or average window, may comprise a window size based on any measurement unit, such as size of bytes or time.

- 5 The size of the window may be based on time, for example priorities over a predetermined time period or frequency . The size of window may be based on space, such as number of PDUs or bytes received or buffered. The size of the window may be based on time and space, such as a number of PDUs or bytes received over a period of time. The size of the window may be based on changes to priorities or based on rate of change of priorities. The
- 10 window may be a rolling window which may overlap or not overlap with previous or subsequent windows. The window may be based on sampling or a sample plan, such as sampling a number of priorities over the window. The sampling can be time based and/or space based (e.g., every nth priority within a predetermined time period).

- Each connection may have the same window or a different window. The QoS engine
- 15 may assign windows to each connection based on link utilization of that connection. The QoS engine may assign windows to each connection based on link utilization of the intermediary. The QoS engine may assign windows to each connection based on historical operational and/or performance statistics for that connection. The QoS engine may change the window for each connection based on link utilization or historical statistics. The QoS
- 20 engine may change the window for each connection at a predetermined frequency.

- The QoS engine may determine or establish a priority for the connection, or a connection based priority, based on a result of applying the filter to the window of priorities within the connection. The QoS engine may determine or establish a priority for the connection for each new or subsequent window. The QoS may determine or establish the
- 25 priority for the connection based on a predetermined frequency, which may be temporal or window based. The QoS may determine or establish a single priority for the connection for transmitting the packets, PDUs or bytes within the connection. The single priority may be the same or different from the priorities of the underlying packets, PDUs or bytes within the connection. The QoS may determine or establish a priority for the connection for the
- 30 packets, PDUs or bytes to be transmitted by the intermediary via the connection to the receiver. The QoS may determine or establish a priority for the connection for transmitting the packets, PDUs or bytes of the window in which the priority was determined.

The QoS engine or network optimization engine 250 may set any QoS fields or bits in the network packets to be transmitted by the intermediary. The QoS engine or network

optimization engine 250 may set any QoS fields or bits in any header of any layer of the network stack. In some embodiments, the QoS engine or network optimization engine 250 may set any priority header or field in the network or IP layer. In some embodiments, the QoS engine or network optimization engine 250 may set any priority header or field in the transport layer. In some embodiments, the QoS engine or network optimization engine 250 may set any priority header or field in any layer below the layer of the PDU, data or channels from which the average priorities were determined. In some embodiments, the QoS engine or network optimization engine 250 may set any priority header or field in any layer above the layer of the PDU, data or channels from which the average priorities were determined. The QoS engine or network optimization engine 250 may set any QoS fields or bits in any protocol carried by the packets. The QoS engine or network optimization engine 250 may set any QoS fields or bits in any payload carried by the packets. In some embodiments, the QoS engine may set any DIFFserv (Differentiated Service Model) bits of the packets. In some embodiments, the QoS engine may set any Type of Service (ToS) bits of the packets.

The QoS engine may provide adaptive or dynamic connection based priorities that take into account other factors besides the priorities of PDUs, channels or bytes within each connection. The QoS engine may monitor and measure utilization of the link by the intermediary to the one or more receivers. The QoS engine may monitor and measure utilization of the link by each of the connections. The QoS engine may monitor and measure bandwidth utilization by the intermediary and/or the connections. For example the QoS engine may determine the number of bytes being sent via the intermediary. The QoS engine may determine the transmission rate of bytes or PDUs being sent via the intermediary. The QoS engine may determine the number of bytes being sent via a connection. The QoS engine may determine the transmission rate of PDUs being sent via a connection. The QoS engine may monitor and measure throughput utilization by the intermediary and/or the connections.

Based on utilization, the QoS engine may adjust the priority assigned or to be assigned to any connection. The QoS engine may interface, integrate or communicate with or be responsive to the filter. In some embodiments, the filter is responsive to the monitoring and link utilization determines. In some embodiments, the QoS engine may provide as input to the filter a weighting or other factor based on the utilization. In some embodiments, the QoS engine may up weight a priority of a connection determined by the filter. In some embodiments, the QoS engine may down weight a priority of a connection determined by the filter. In some embodiments, the QoS engine independently of the filter adjust the priority of the connection based on utilization. For example, the QoS engine may assign the priority

from the filter to the connection and as the utilization changes, the QoS engine changes the priority of the connection accordingly. The filter may continually change the priority to be assigned or that are assigned to the connection based on priorities of the PDUs, channel or byte stream within the same connection while the QoS engine changes the assigned priority of the same connection based on utilization.

The QoS engine may provide rate control over changes to the connection based on priority. The QoS engine may control the rate of change of a priority. The QoS engine may control how quickly or slowly the priority of a connection changes. The QoS engine may allow the priority of a connection to change only within a predetermined amount, for example, by a percentage at a time. The QoS engine may allow the priority of a connection to change after a predetermined number of windows. The QoS engine may allow the priority of a connection to change after a predetermined number of bytes or PDUs. The QoS engine may allow the priority of a connection to change after a predetermined time period. The QoS engine may allow the priority of a connection to change based on utilization or changes thereto. For example, the engine QoS engine may allow the priority of a connection to change if the utilization is within a predetermined threshold.

The QoS engine may dynamically adapt the predetermined amount responsive to the rate of change. The QoS engine may dynamically adapt the predetermined amount responsive to the values of the priorities assigned to a connection. For example, the QoS engine may allow slower or quicker changes for lower priorities. In some embodiments, the QoS engine may allow slower or quicker changes for higher priorities. For example, the QoS engine may quickly lower priority of the connection when encountering low priority PDU(s), and then slowly raise the priority of the connection based on time and number of PDUs with specific priorities. The QoS engine may only allow the priority of a connection to change by a predetermined level, number or granularity of priority. The QoS engine may throttle the changes to the priority by a predetermined level, number or granularity of priority. The QoS engine may only allow the priority of a connection to change over time by a predetermined rate. The QoS engine may throttle the rate of change to the priority by a predetermined rate.

The intermediary, such as via the QoS or network optimization engine may receive and buffer a first set of packets as they are received, perform priority filtering on such packets, assign the connection priority to the connection and transmit those buffered packets via the connection while the connection priority determined based on those packets is applied or in effect for the connection. The intermediary, such as via the QoS or network optimization engine may receive and buffer a next or subsequent set of packets as they are

received, perform priority filtering on such packets, assign the connection priority to the connection and transmit those buffered packets of the next or subsequent set of packets via the connection while the connection priority determined based on those packets is applied or in effect for the connection. As such, in some embodiments, the intermediary may transmit  
5 those packets in which the connection priority is based on thru the connection with the connection is assigned the connection priority. For the example, the intermediary may assign the connection priority for a window and transmit those packets in that window while the connection priority for that window is assigned to the connection.

The QoS engine may transmit packets from the intermediary based on the priorities  
10 assigned to each connection. For example, the QoS Engine may first transmit one or more packets of a first connection with a connection priority higher than a second connection. The QoS engine may use any type and form of prioritization scheme to transmit packets from the intermediary based on the priorities assigned to each connection. In some embodiments, the QoS engine sets the priority in the packets or PDUs to be transmitted. In some embodiments,  
15 the intermediary transmits the packets to a QoS device to perform QoS for the connections using the connection based priority.

Referring now to FIG. 4B, an embodiment of a method for connection based prioritization for connections carrying multiple priority data streams is depicted. In brief overview, at step 450, the intermediary identifies for each connection priorities of virtual  
20 channels, PDUs or bytes received within the connection. At step 455, the intermediary applies a filter to a count/sample of PDU/bytes. At step 460, based on filtering, the intermediary assigns a priority to connection. So although the PDUs or bytes within the connection have differing priorities, the connection has a single priority based on filtering of these differing priorities. The intermediary may continue identifying and filtering PDU/byte  
25 priorities to assign a priority for the connection. The intermediary may adjust a connection's priority based on link utilization and/or rate control. At step 465, the intermediary transmits PDUs/bytes of the connections based on the priorities assigned to the connections.

At step 450, the intermediary identifies priorities within connections received by or traversing the intermediary. The intermediary may receive and buffer packets for a  
30 connection for which the priorities are to be identified. The intermediary may receive and buffer PDUs, data or bytes of a payload of a transport layer packet for a connection for which the priorities are to be identified. The intermediary may receive for each single transport layer connection a plurality of packets carrying a plurality of protocol data units. Each protocol data unit of the plurality of protocol data units may identify the same or different

priority of another protocol data unit in the plurality of packets. In some embodiments, the intermediary may receive for each single transport layer connection the plurality of packets carrying the plurality of protocol data units of a remote display protocol, such as ICA or RDP. Each protocol data unit may identify or correspond to a virtual channel, which may have a priority.

The intermediary may identify the priority of a virtual channel, PDU or group of bytes from any header or field of any protocol at any layer of the network stack carried via the connection. The intermediary may identify the priority of a virtual channel, PDU or group of bytes from any header or field of the protocol carrying the virtual channel, PDU or group of bytes. The intermediary may identify the priority of a virtual channel, PDU or group of bytes from any header or field of a protocol above the transport layer of the connection. The intermediary may identify the priority of a virtual channel, PDU or group of bytes from the payload carried by the transport layer packet(s) of the connection.

The intermediary may identify priorities from each instance of a priority identified from a PDU. The intermediary may identify priorities from each instance of a priority identified within the data or byte stream. The intermediary may identify priorities within a predetermined window for each connection. The intermediary may sample priorities based on virtual channels, PDUs or group of bytes. The intermediary may store the priorities to memory or storage element of the intermediary. The intermediary may store a group of priorities from a window. The intermediary may use a data structure, array or object to hold in memory identified priorities. The intermediary may have a data structure, array or object store or hold identified priorities on a connection by connection basis.

At step 455, the intermediary applies a filter to a group or set of priorities for a connection. The filter of the intermediary may determine an average priority based on finite impulse response (FIR) filtering. The filter of the intermediary may determine an average priority based on infinite impulse response (IIR) filtering. The filter of the intermediary may specify, configure or identify a predetermined window for filtering in terms of or as a predetermined number of protocol data units. The intermediary may determine an average priority for a predetermined window of protocol data units of the plurality of protocol data units received by the intermediary. The filter of the intermediary may specify, configure or identify a predetermined window for filtering in terms of or as a predetermined time and/or space conditions.

The intermediary applies a filter to a group or set of priorities each of the connections traversing or managed by the intermediary. The intermediary may apply the filter to

generate, calculate or provide a value for a priority for the connection based on the differing priorities of data traversing thru or carried by the connection. The intermediary may apply the filter to one or more windows of priorities. The intermediary may apply the filter to determine a type of average to a group of priorities, such as a window of priorities. The intermediary may apply the filter to take a weighted average of a group of priorities, such as a window of priorities. The intermediary may take as input to the filter other factors, such as a link utilization and rate control policies, in applying the filter to a group of priorities. The intermediary may apply the filter to determine a type of average of the priorities carried in a protocol or protocol layer above the transport layer of the connection.

At step 460, the intermediary may assign a priority to a connection based on filtering. The intermediary may assign the average priority from the filter as a connection priority to the corresponding transport layer connection. The intermediary may assign an initial priority to a connection based on applying filtering to a first window of priorities received via the connection. The intermediary may reassign, adjust or otherwise change the priority of the connection based on applying filtering to subsequent windows of priorities. The intermediary may adjust the priority for any one or more connections based on the intermediary's link utilization. The intermediary may adjust the priority for any one or more connections based on the connection's link utilization. The intermediary may continually identify priorities within a connection, apply filtering to the priorities and adjust the priority of the connection. The intermediary may control the rate at which the priority of the connection is adjusted. In some embodiments, although the filtering of priorities may produce a higher change in rate, the intermediary may throttle or adjust the produced priority to a lower priority. In some embodiments, although the filtering of priorities may produce a lower change in rate, the intermediary may increase or adjust the produced priority to a higher priority. In some embodiments, the intermediary may set the priority for the connection by setting any QoS bits or fields of the packets or PDUs to be transmitted, such as ToS or Diffserv bits. As such, in some embodiments, the intermediary sets a Differentiated Services Code Point (DSCP) field in the plurality of packets (each of these packets may correspond to the filtered window) to correspond to the connection priority assigned to the transport layer connection. The intermediary may transmit one or more data packets with the connection priority identified via one of an internet protocol (IP layer) or transport layer and the one or more packets carrying a first protocol data unit having the priority identified via a protocol above the transport layer.

At step 465, the intermediary may transmit packets, PDUs or group of bytes of a connection based on the connection's currently assigned priority. The intermediary may transmit via the transport layer connection, the plurality of packets carrying those protocol data units within the predetermined window of protocol data units while the connection priority of the transport layer connection is assigned the average priority for the predetermined window of protocol data units. The intermediary may transmit those packets buffered for the window in which the connection priority was determined upon setting the connection priority of the connection. The intermediary may transmit those packets buffered for the window in which the connection priority was determined while the connection priority of the connection is set to such a connection priority. The intermediary may buffer the next window of packets until the connection priority for the next window of packets is determined and assigned to the connection. Upon setting the connection priority of the connection for the next window, the intermediary may transmit the buffered packets.

The intermediary may apply any quality of service scheme or algorithm among multiple connections for which the intermediary may assign priorities. The intermediary may take into account the current priority of each of the connections for transmitting packets, PDUs or group of bytes via the intermediary. A higher priority connection may receive more transmission time and/or transmission of a greater amount of data than a lower priority connection. In some embodiments, the intermediary transmits packets, PDUs or group of bytes in accordance with any QoS setting of the packets, PDUs or group of bytes, which has been set based on the connection's priority. In some embodiments, the intermediary transmits packets, PDUs or group of bytes in accordance with a ToS or Diffserv setting has been set based on the connection's priority. In some embodiments, the intermediary transmits packets, PDUs or group of bytes with a QoS setting, such as ToS or Differv setting, based on the connection's priority to another QoS system or device for QoS processing and transmission.

Any of the above steps may be continually performed for a stream of packets received and processed by the intermediary device. For example, the intermediary may identify a second, next or subsequent average priority for a second, next or subsequent predetermined window of protocol data units of a second, next or subsequent plurality of protocol data units and assign the determines average priority as the connection priority of the transport layer connection. The intermediary may transmit via the single transport layer connection, a second, next or subsequent plurality of packets carrying those protocol data units within the second, next or subsequent predetermined window of protocol data units while the

connection priority of the transport layer connection is assigned the second, next or subsequent average priority for the second, next or subsequent predetermined window of protocol data units.

What is claimed:

1. A method for providing quality of service (QoS) for a stream of protocol data units via a single transport layer connection, the method comprising:

(a) receiving, by a device via a single transport layer connection, a plurality of packets carrying a plurality of protocol data units, each protocol data unit of the plurality of protocol data units identifying a priority, at least one or more of the plurality of protocol data units comprising a first priority different than a second priority of other protocol data units of the plurality of protocol data units;

(b) determining, by the device, an average priority for a predetermined window of protocol data units of the plurality of protocol data units;

(c) assigning, by the device, the average priority as a connection priority of the single transport layer connection; and

(d) transmitting, by the device via the single transport layer connection, the plurality of packets carrying those protocol data units within the predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the average priority for the predetermined window of protocol data units.

2. The method of claim 1, wherein step (a) further comprises receiving, by the device via the single transport layer connection the plurality of packets carrying the plurality of protocol data units of a remote display protocol, each protocol data unit corresponding to a virtual channel, the virtual channel assigned the priority.

3. The method of claim 1, wherein step (a) further comprises buffering, by the device, the plurality of packets until the connection priority for the plurality of packets is assigned to the transport layer connection.

4. The method of claim 1, wherein step (b) further comprises determining, by a filter of the device, the average priority based on finite impulse response (FIR) filtering.

5. The method of claim 1, wherein step (b) further comprises determining, by a filter of the device, the average priority based on infinite impulse response (IIR) filtering.

6. The method of claim 1, wherein step (b) further comprise identifying, by the device, the predetermined window as comprising a predetermined number of protocol data units.

7. The method of claim 1, wherein step (c) further comprises setting, by the device, a Differentiated Services Code Point (DSCP) field in the plurality of packets to correspond to the connection priority assigned to the single transport layer connection.

8. The method of claim 1, wherein step (d) further comprising transmitting, by the device, a first packet of the plurality of data packets with the connection priority identified via one of an internet protocol (IP layer) or transport layer and the first packet carrying a first protocol data unit having the priority identified via a protocol above the transport layer.
9. The method of claim 1, further comprising determining, by the device, a second average priority for a second predetermined window of protocol data units of a second plurality of protocol data units and assigning, by the device, the second average priority as the connection priority of the single transport layer connection.
10. The method of claim 9, further comprising transmitting, by the device via the single transport layer connection, a second plurality of packets carrying those protocol data units within the second predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the second average priority for the second predetermined window of protocol data units.
11. A system for providing quality of service (QoS) for a stream of protocol data units via a single transport layer connection, the system comprising:
- a device intermediary to a sender and receiver that receives via a single transport layer connection a plurality of packets carrying a plurality of protocol data units, each protocol data unit of the plurality of protocol data units identifying a priority, at least one or more of the plurality of protocol data units comprising a first priority different than a second priority of other protocol data units of the plurality of protocol data units;
  - a filter of the device determining an average priority for a predetermined window of protocol data units of the plurality of protocol data units;
  - an engine of the device assigning the average priority as a connection priority of the single transport layer connection; and
  - wherein the device transmits via the single transport layer connection the plurality of packets carrying those protocol data units within the predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the average priority for the predetermined window of protocol data units.
12. The system of claim 11, wherein the device receives via the single transport layer connection the plurality of packets carrying the plurality of protocol data units of a

remote display protocol, each protocol data unit corresponding to a virtual channel, the virtual channel assigned the priority.

13. The system of claim 11, wherein the engine buffers the plurality of packets until the connection priority for the plurality of packets is assigned to the transport layer

connection.

14. The system of claim 11, wherein the filter determines the average priority based on finite impulse response (FIR) filtering.

15. The system of claim 11, wherein the filter determines the average priority based on infinite impulse response (IIR) filtering.

16. The system of claim 11, wherein engine identifies the predetermined window as comprising a predetermined number of protocol data units.

17. The system of claim 11, wherein the engine sets the Differentiated Services Code Point (DSCP) field in the plurality of packets to correspond to the connection priority assigned to the single transport layer connection.

18. The system of claim 11, wherein the device transmits a first packet of the plurality of data packets with the connection priority identified via one of an internet protocol (IP layer) or transport layer and the first packet carrying a first protocol data unit having the priority identified via a protocol above the transport layer.

19. The system of claim 11, wherein the filter determines a second average priority for a second predetermined window of protocol data units of a second plurality of protocol data units and the engine assigns the second average priority as the connection priority of the single transport layer connection.

20. The system of claim 19, wherein device transmits via the single transport layer connection, a second plurality of packets carrying those protocol data units within the second predetermined window of protocol data units while the connection priority of the single transport layer connection is assigned the second average priority for the second predetermined window of protocol data units.




FIG. 1A




FIG. 1B






FIG. 1C

*Fig. 1D*

*Fig. 1E*




FIG. 2A





FIG. 2B

Client 102

100

*Fig. 3*





FIG. 4A

*Fig. 4B*

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2012/027906

A. CLASSIFICATION OF SUBJECT MATTER  
INV. H04L12/56  
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
H04L

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

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

EP0-Internal, WPI Data

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                  | Relevant to claim No.  |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| Y         | US 7 046 631 B1 (GIROUX NATALIE [CA] ET AL) 16 May 2006 (2006-05-16)                                                                | 1-3,<br>6-13,<br>16-20 |
| A         | abstract<br>column 5, line 15 - line 34<br>column 6, line 15 - column 7, line 67<br>column 9, line 38 - column 12, line 43<br>----- | 4,5,14,<br>15          |
| Y         | US 2005/025159 A1 (LEE DUAN-SHIN [TW] ET AL) 3 February 2005 (2005-02-03)                                                           | 1-3,<br>6-13,<br>16-20 |
| A         | paragraph [0005] - paragraph [0011]<br>paragraph [0024]<br>paragraph [0050] - paragraph [0061]<br>-----<br>-/--                     | 4,5,14,<br>15          |



Further documents are listed in the continuation of Box C.



See patent family annex.

\* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

11 May 2012

Date of mailing of the international search report

21/05/2012

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040,  
Fax: (+31-70) 340-3016

Authorized officer

Clemente Lafuente, G

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2012/027906

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                              | Relevant to claim No.                   |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| Y<br>A    | US 2004/221053 A1 (CODELLA CHRISTOPHER<br>FRANK [US] ET AL)<br>4 November 2004 (2004-11-04)<br>paragraph [0033] - paragraph [0037]<br><br>----- | 1-3,<br>6-13,<br>16-20<br>4,5,14,<br>15 |

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2012/027906

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date     |
|-------------------------------------------|---------------------|----------------------------|-------------------------|
| US 7046631                                | B1                  | 16-05-2006                 | NONE                    |
| -----                                     |                     |                            |                         |
| US 2005025159                             | A1                  | 03-02-2005                 | TW I228008 B 11-02-2005 |
|                                           |                     | US 2005025159 A1           | 03-02-2005              |
| -----                                     |                     |                            |                         |
| US 2004221053                             | A1                  | 04-11-2004                 | NONE                    |
| -----                                     |                     |                            |                         |