



(51) International Patent Classification:

H04L 29/08 (2006.01) *G06F 9/50* (2006.01)
G06F 9/455 (2006.01)

(21) International Application Number:

PCT/US2012/059586

(22) International Filing Date:

10 October 2012 (10.10.2012)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/547,493 14 October 2011 (14.10.2011) US

(71) Applicant: CITRIX SYSTEMS, INC. [US/US]; 851 West Cypress Creek Road, Fort Lauderdale, Florida 33309 (US).

(72) Inventor: MURGIA, Marco; c/o Citrix Systems, Inc., 4988 Great America Parkway, Santa Clara, California 95054 (US).

(74) Agent: MCKENNA, Christopher, J.; Foley & Lardner LLP, 3000 K Street N.W., Suite 600, Washington D.C., District of Columbia 20007-5109 (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, BN, 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, PA, 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, 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).

Declarations under Rule 4.17:


— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

— as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

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

(54) Title: SYSTEMS AND METHODS FOR DYNAMIC ADAPTATION OF NETWORK ACCELERATORS



Appliance 200

FIG. 6C

(57) Abstract: Systems and methods of the present solution provide a more optimal solution by dynamically and automatically reacting to changing network workload. A system that starts slowly, either by just examining traffic passively or by doing sub-optimal acceleration can learn over time, how many peer WAN optimizers are being serviced by an appliance, how much traffic is coming from each peer WAN optimizers, and the type of traffic being seen. Knowledge from this learning can serve to provide a better or improved baseline for the configuration of an appliance. In some embodiments, based on resources (e.g., CPU, Memory, Disk), the system from this knowledge may determine how many WAN optimization instances should be used and of what size, and how the load should be distributed across the instances of the WAN optimizer.

SYSTEMS AND METHODS FOR DYNAMIC ADAPTATION OF NETWORK ACCELERATORS

Related Application

5 This application claims the benefit of and priority to U.S. Provisional Application No. 61/547,493, entitled “Systems and Methods For Dynamic Adaptation of Network Accelerators” and filed on October 14, 2011, which is incorporated herein by reference for all purposes.

Field of the Invention

10 The present disclosure generally relates to data communication networks. In particular, the present disclosure relates to systems and methods for the dynamic adaptation of network accelerators on a platform.

Background of the Invention

15 Traditional network elements have been developed and deployed as discrete network appliances and functions. As such, they have been deployed in a particular way and sized to accommodate some sort of network traffic model. The model can be as simple as “I need gigabit speed for this link” to more sophisticated models where traffic analysis and peak loading are measured and placed in a model that determines network element sizing.

20

Brief Summary of the Invention

As computing in general and networking in particular moves to more virtualized environments, there exist several problems with the prior models. As the workload become more mobile and dynamic, traditional network engineering becomes virtually impossible and
25 new mechanisms must be used to “adapt” the network infrastructure to the offered load.

In some deployments, a load balancer may be used to load balance multiple WAN optimizers. In further deployments, a virtualized load balancer may be used to load balance multiple virtualized WAN optimizers. The load balancer may be configured to use or apply any of a plurality of load balancing methods, such as but not limited to a least connection
30 method, least connection with agent id persistence, a static configuration, least accumulated load using receive bandwidth, least accumulated load using receive bandwidth with agent id

persistence and source internet protocol (IP) hashing. Each of these methods may provide varying load balance effectiveness, compression history performance, bandwidth management and simplicity with respect to load balancing WAN optimizers. There may be no one-size-fits-all solution for all deployments. There may be a trade off between ease of deployment and WAN optimization.

Systems and methods of the present solution provide a more optimal solution by dynamically and automatically reacting to changing network workload. A system that starts slowly, either by just examining traffic passively or by doing sub-optimal acceleration can learn over time, how many peer WAN optimizers are being serviced by an appliance, how much traffic is coming from each peer WAN optimizers, and the type of traffic being seen. Knowledge from this learning can serve to provide a better or improved baseline for the configuration of an appliance. In some embodiments, based on resources (e.g., CPU, Memory, Disk), the system from this knowledge may determine how many WAN optimization instances should be used and of what size, and how the load should be distributed across the instances of the WAN optimizer. Some example rules are as follows:

1. If a small number of peer WAN optimizers exist, a smaller number of large WAN optimizer instances should be provisioned on an appliance because compression histories will be less fragmented and compression ratios higher (better)
2. When peers WAN optimizers are of significantly different sizes, they should be distributed unevenly across the WAN optimizers instances (perhaps using WAN optimizers instances of different sizes)

Over time the system may continue to monitor the workload for changes and change the load distribution as needed, for example, adding additional WAN optimizers instances because first pass compression (high CPU utilization) is being used extensively and compression history fragmentation effects would be minimal. Or traffic from a remote site has grown significantly and a larger WAN optimizers instance is required to keep compression history from fragmenting.

In one aspects, the present solution is directed to a method for managing a plurality of instances of a Wide Area Network (WAN) optimizer executing on an intermediary device. The method includes establishing, on a device intermediary to a plurality of clients and plurality of servers, a plurality of instances of a Wide Area Network (WAN) optimizer to accelerate WAN communications between the plurality of clients and the plurality of servers.

The method includes monitoring, by the device, network traffic traversing the device for each of the plurality of instances of the WAN optimizer and selecting, by a manager executing on the device responsive to the monitoring, a change of a load balancing scheme to load balance the plurality of instances of the WAN optimizer.

5 The method may automatically establish, by the device, a configuration of a size of each of the plurality of instances of the WAN optimizer based on data stored from monitoring of previous execution of the plurality of instances of the WAN optimizer. Each of the plurality of instances of the WAN optimizer may execute as a virtual machine in a virtualized environment.

10 In some embodiments, the method includes comprises monitoring, by the device, compression history allocation, compression fragmentation and compression ratios of each of the plurality of instances of the WAN optimizer. In some embodiments, the method includes monitoring, by the device, or more of the following of each of the plurality of instances of the WAN optimizer: resource utilization, number of connections, number of claims and
15 bandwidth usage.

 In some embodiments, the method includes determining, by the device, that a metric computed from monitoring network traffic has exceeded a threshold and responsive to the determination, automatically selecting by the device a second load balancing scheme to load balance the plurality of instances of the WAN optimizer. In some embodiments, the method
20 includes automatically switching, by the device, from the load balancing scheme to the selected load balancing scheme while executing the plurality of instances of the WAN optimizer. In some embodiments, the method includes automatically changing, by the device responsive to the monitoring, the number of instances of the WAN optimizer executing on the device. In some embodiments, the method includes automatically adjusting, by the
25 device responsive to the monitoring, a size of resource usage used by one or more of the plurality of instances of the WAN optimizer. The method may also include applying, by the device one or more rules to data collected from monitoring, to determine to change one or more of the following: a number of instances of the WAN optimizer, a size of one or more WAN optimizers and the load balancing scheme.

30 In some aspects, the present solution is directed to a system for managing a plurality of instances of a Wide Area Network (WAN) optimizer executing on an intermediary device. The system includes a device intermediary to a plurality of clients and plurality of servers and a plurality of instances of a Wide Area Network (WAN) optimizer executing on the device to accelerate WAN communications between the plurality of clients and the plurality of servers;

The system includes a monitor that monitors network traffic traversing the device for each of the plurality of instances of the WAN optimizer; and a manager executing on the device that, responsive to the monitor, selects a change of a load balancing scheme to load balance the plurality of instances of the WAN optimizer.

5 In some embodiments, the manager automatically establishes a configuration of a size of each of the plurality of instances of the WAN optimizer based on data stored from monitoring of previous execution of the plurality of instances of the WAN optimizer. In some embodiments, each of the plurality of instances of the WAN optimizer execute as a virtual machine in a virtualized environment.

10 In some embodiments, the monitor monitors compression history allocation, compression fragmentation and compression ratios of each of the plurality of instances of the WAN optimizer. In some embodiments, the monitor monitors one or more of the following of each of the plurality of instances of the WAN optimizer: resource utilization, number of connections, number of claims and bandwidth usage.

15 In some embodiments, manager determines that a metric computed from monitoring network traffic has exceeded a threshold and responsive to the determination, automatically selects a second load balancing scheme to load balance the plurality of instances of the WAN optimizer. In some embodiments, the manager automatically switches from a current load balancing scheme to the selected load balancing scheme while executing the plurality of
20 instances of the WAN optimizer. In some embodiments, the manager, responsive to the monitor, automatically changes the number of instances of the WAN optimizer executing on the device. In some embodiments, the manager automatically adjusts, responsive to the monitor, a size of resource usage used by one or more of the plurality of instances of the WAN optimizer. In some embodiments, the manager applies one or more rules to data
25 collected from monitoring, to determine to change one or more of the following: a number of instances of the WAN optimizer, a size of one or more WAN optimizers and the load balancing scheme.

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 taken in conjunction with the accompanying drawings, in which:

5 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 other network appliances;

10 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;

 FIG. 2A is a block diagram of an embodiment of an appliance for processing
15 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 an appliance;

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

20 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 virtualization environment;

 FIG. 4B is a block diagram of another embodiment of a virtualization environment;

 FIG. 4C is a block diagram of an embodiment of a virtualized network optimization
25 engine;

 FIG. 4D is a block diagram of an embodiment of a virtualized application delivery controller;

 FIG. 5A are block diagrams of embodiments of approaches to implementing parallelism in a multi-core system;

30 FIG. 5B is a block diagram of an embodiment of a system utilizing a multi-core system;

FIG. 5C is a block diagram of another embodiment of an aspect of a multi-core system;

FIG. 6A is a block diagram of an embodiment of a virtualization platform for dynamic adaptation of virtual appliances;

5 FIG. 6B is a diagram of an embodiment of load balancing methods for virtual appliances;

FIG. 6C is a block diagram of an embodiment of a system for dynamic adaptation of virtual appliances;

10 FIG. 6D is a block diagram of an embodiment of a method for dynamic adaptation of virtual appliances.

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
15 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
20 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;
- 25 - Section C describes embodiments of a client agent for accelerating communications between a client and a server;
- Section D describes embodiments of systems and methods for virtualizing a network optimization engine;
- Section E describes embodiments of systems and methods for providing a
30 multi-core architecture and environment; and
- Section F describes embodiments of systems and methods for dynamic adaptation of virtual appliances.

A. Network and Computing Environment

5 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 205 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 on 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.

5 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
10 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.
15 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.
20 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 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
25 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 or to act or simulate as terminating the connection.

In some embodiments, the second appliance 205 is a performance enhancing proxy.
30 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 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 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 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 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'.

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 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.

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 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 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. 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 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 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 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
5 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 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
10 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 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)
15 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 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
20 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,
25 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
30 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)

5 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.

10 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
15 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
20 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
25 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
30 205. 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.

5 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
10 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
15 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
20 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
25 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
30 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 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 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, 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 computer, laptop or notebook computer, server, handheld computer, mobile telephone, any
5 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.

10 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 delivering a computing environment or an application and/or data file to one or more clients
15 102. 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
20 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
25 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
30 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 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 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 (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 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 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 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.

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. 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.

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® 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.

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, 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.

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 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 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 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 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 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

5 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
10 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
15 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
20 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,
25 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,
30 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).

5 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.

5

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 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 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 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 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, 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 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
5 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
10 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
15 (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.
20 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 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
25 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 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
30 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 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 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 "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 (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 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 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 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 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 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 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.

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 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

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 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).

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 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.

5 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) 10 window virtualization, 3) recongestion 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 15 100 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 20 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, 25 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 30 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."

5 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
10 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.

15 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
20 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
25 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
30 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. 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 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 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 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 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 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 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 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 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 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 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 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 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 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 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 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
5 200 monitors each TCP connection, characterizing it with respect to the different dimensions, selecting a recongestion algorithm that is appropriate for the current characterization.

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
10 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 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
15 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
20 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
25 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
30 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) 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 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 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. 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 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.

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.

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 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 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 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 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 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 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 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 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 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.

5 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 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
10 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 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
15 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
20 (using the transmit numbers) declares a packet lost, the sender considers the packet to be only 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
25 presumed lost list. Packets that contain data known to be missing are declared actually lost 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

30 In some embodiments, the appliance 200 or flow controller 220 applies a technique 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 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 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 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

5 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
10 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
15 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 fiber optics) have extremely low bit-error rates, low enough that these can be ignored. For these
20 media, it is advantageous for the packet size to be the maximum possible before 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
25 cause an entire packet to be dropped (i.e., media-induced errors beyond the capability of the 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
30 path, depending on the nature of each link.

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 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.

5

QoS

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
10 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 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
15 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
20 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
25 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
30 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.

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 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 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

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 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 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 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.

Referring now to FIG. 2B, another embodiment of the appliance 205 is depicted. In brief overview, the appliance 205 provides one or more of the following services, functionality or operations: SSL VPN connectivity 280, switching/load balancing 284, Domain Name Service resolution 286, acceleration 288 and an application firewall 290 for communications between one or more clients 102 and one or more servers 106. Each of the servers 106 may provide one or more network related services 270a-270n (referred to as services 270). For example, a server 106 may provide an http service 270. The appliance 205 comprises one or more virtual servers or virtual internet protocol servers, referred to as a vServer, VIP server, or just VIP 275a-275n (also referred herein as vServer 275). The vServer 275 receives, intercepts or otherwise processes communications between a client 102 and a server 106 in accordance with the configuration and operations of the appliance 205.

The vServer 275 may comprise software, hardware or any combination of software and hardware. The vServer 275 may comprise any type and form of program, service, task, process or executable instructions operating in user mode 202, kernel mode 204 or any combination thereof in the appliance 205. The vServer 275 includes any logic, functions, rules, or operations to perform any embodiments of the techniques described herein, such as SSL VPN 280, switching/load balancing 284, Domain Name Service resolution 286, acceleration 288 and an application firewall 290. In some embodiments, the vServer 275 establishes a connection to a service 270 of a server 106. The service 275 may comprise any program, application, process, task or set of executable instructions capable of connecting to and communicating to the appliance 205, client 102 or vServer 275. For example, the service

275 may comprise a web server, http server, ftp, email or database server. In some embodiments, the service 270 is a daemon process or network driver for listening, receiving and/or sending communications for an application, such as email, database or an enterprise application. In some embodiments, the service 270 may communicate on a specific IP address, or IP address and port.

In some embodiments, the vServer 275 applies one or more policies of the policy engine 236 to network communications between the client 102 and server 106. In one embodiment, the policies are associated with a vServer 275. In another embodiment, the policies are based on a user, or a group of users. In yet another embodiment, a policy is global and applies to one or more vServers 275a-275n, and any user or group of users communicating via the appliance 205. In some embodiments, the policies of the policy engine have conditions upon which the policy is applied based on any content of the communication, such as internet protocol address, port, protocol type, header or fields in a packet, or the context of the communication, such as user, group of the user, vServer 275, transport layer connection, and/or identification or attributes of the client 102 or server 106.

In other embodiments, the appliance 205 communicates or interfaces with the policy engine 236 to determine authentication and/or authorization of a remote user or a remote client 102 to access the computing environment 15, application, and/or data file from a server 106. In another embodiment, the appliance 205 communicates or interfaces with the policy engine 236 to determine authentication and/or authorization of a remote user or a remote client 102 to have the application delivery system 190 deliver one or more of the computing environment 15, application, and/or data file. In yet another embodiment, the appliance 205 establishes a VPN or SSL VPN connection based on the policy engine's 236 authentication and/or authorization of a remote user or a remote client 102. In one embodiment, the appliance 205 controls the flow of network traffic and communication sessions based on policies of the policy engine 236. For example, the appliance 205 may control the access to a computing environment 15, application or data file based on the policy engine 236.

In some embodiments, the vServer 275 establishes a transport layer connection, such as a TCP or UDP connection with a client 102 via the client agent 120. In one embodiment, the vServer 275 listens for and receives communications from the client 102. In other embodiments, the vServer 275 establishes a transport layer connection, such as a TCP or UDP connection with a client server 106. In one embodiment, the vServer 275 establishes the transport layer connection to an internet protocol address and port of a server 270 running

on the server 106. In another embodiment, the vServer 275 associates a first transport layer connection to a client 102 with a second transport layer connection to the server 106. In some embodiments, a vServer 275 establishes a pool of transport layer connections to a server 106 and multiplexes client requests via the pooled transport layer connections.

5 In some embodiments, the appliance 205 provides a SSL VPN connection 280 between a client 102 and a server 106. For example, a client 102 on a first network 102 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
10 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, the client agent 120 intercepts communications of the 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
15 network 104. The appliance 205 receives the intercepted communication from the client agent 102, 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. As such, the appliance 205 provides an end-to-end secure transport layer connection for the client 102 between the two
20 networks 104, 104'.

In one embodiment, the appliance 205 hosts an intranet internet protocol or IntranetIP 282 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
25 appliance 205 establishes, assigns or otherwise provides an IntranetIP address 282, which is a network 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 282. In one embodiment, the appliance 205 acts as or on behalf of the client 102 on
30 the second private network 104. For example, in another embodiment, a vServer 275 listens for and responds to communications to the IntranetIP 282 of the client 102. In some embodiments, if a computing device 100 on the second network 104' transmits a request, the appliance 205 processes the request as if it were the client 102. For example, the appliance 205 may respond to a ping to the client's IntranetIP 282. In another example, the appliance

may establish a connection, such as a TCP or UDP connection, with computing device 100 on the second network 104 requesting a connection with the client's IntranetIP 282.

In some embodiments, the appliance 205 provides one or more of the following acceleration techniques 288 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. In one embodiment, the appliance 205 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 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 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 numbers expected by the client 102 on the appliance's 200 transport layer connection to the client 102. In some embodiments, the packet checksum of the transport layer protocol is recalculated to account for these translations.

In another embodiment, the appliance 205 provides switching or load-balancing functionality 284 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 or application-layer request data. In one embodiment, although the 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, the health monitoring programs 216 of the appliance 205 monitor the health of servers to determine the

server 106 for which to distribute a client's request. In some embodiments, if the 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
5 embodiments, the appliance intercepts a DNS request transmitted by the client 102. In one 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 205. In another embodiment, the
10 appliance 205 responds to a client's DNS request with an IP address of or hosted by a second appliance 205'. 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 205.

In yet another embodiment, the appliance 205 provides application firewall functionality 290 for communications between the client 102 and server 106. In one
15 embodiment, the policy engine 236 provides rules for detecting and blocking illegitimate requests. In some embodiments, the application firewall 290 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 236 comprises one or more application firewall or security control
20 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
25 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 290 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
30 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 290 ensures cookies are not modified. In other embodiments, the application firewall 290 protects against forceful browsing by enforcing legal URLs.

In still yet other embodiments, the application firewall 290 protects any confidential information contained in the network communication. The application firewall 290 may inspect or analyze any network communication in accordance with the rules or policies of the engine 236 to identify any confidential information in any field of the network packet. In some embodiments, the application firewall 290 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 comprise these occurrences or the confidential information. Based on these occurrences, in one embodiment, the application firewall 290 may take a policy action on the network communication, such as prevent transmission of the network communication. In another embodiment, the application firewall 290 may rewrite, remove or otherwise mask such identified occurrence or confidential information.

Referring now to FIG. 4D, a diagram of an embodiment of a virtual appliance 460 operating on a hypervisor 401 of a server 106 is depicted. As with the appliance 205 of FIGs. 2A and 2B, the virtual appliance 460 may provide functionality for availability, performance, offload and security. For availability, the virtual appliance may perform load balancing between layers 4 and 7 of the network and may also perform intelligent service health monitoring. For performance increases via network traffic acceleration, the virtual appliance may perform caching and compression. To offload processing of any servers, the virtual appliance may perform connection multiplexing and pooling and/or SSL processing. For security, the virtual appliance may perform any of the application firewall functionality and SSL VPN function of appliance 200.

Any of the modules of the appliance 200 as described in connection with FIGs. 2A may be packaged, combined, designed or constructed in a form of the virtualized appliance delivery controller 460 deployable as one or more software modules or components executable in a virtualized environment 400 or non-virtualized environment on any server, such as an off the shelf server. For example, the virtual appliance may be provided in the form of an installation package to install on a computing device. Any of the components or functionality of the appliance described in FIG. 2C may be designed and constructed as a software component or module to run on any operating system of a computing device and/or of a virtualized environment 300.

Still referring to FIG. 4D, and in brief overview, any one or more vServers 275A-275N may be in operation or executed in a virtualized environment 400 of any type of computing device 100, such as any server 106. Any of the modules or functionality of the

appliance 200 described in connection with FIG. 2C may be designed and constructed to operate in either a virtualized or non-virtualized environment of a server. Any of the vServer 275, SSL VPN 280, Intranet UP 282, Switching 284, DNS 286, acceleration 288, App FW 280 and monitoring agent may be packaged, combined, designed or constructed in a form of application delivery controller 460 deployable as one or more software modules or components executable on a device and/or virtualized environment 400.

In some embodiments, a server may execute multiple virtual machines 406a-406n in the virtualization environment with each virtual machine running the same or different embodiments of the virtual application delivery controller 460. In some embodiments, the server may execute one or more virtual appliances 460 on one or more virtual machines on a core of a multi-core processing system. In some embodiments, the server may execute one or more virtual appliances 460 on one or more virtual machines on each processor of a multiple processor device.

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.

5 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
10 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
15 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,
20 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
25 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
30 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 an 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.

5 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
10 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
15 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 and/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
20 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
25 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
30 205 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

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 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 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 Providing Virtualized Appliances

Referring now to FIG. 4A, a block diagram depicts one embodiment of a virtualization environment 400. In brief overview, a computing device 100 includes a hypervisor layer, a virtualization layer, and a hardware layer. The hypervisor layer includes a

hypervisor 401 (also referred to as a virtualization manager) that allocates and manages access to a number of physical resources in the hardware layer (e.g., the processor(s) 421, and disk(s) 428) by at least one virtual machine executing in the virtualization layer. The virtualization layer includes at least one operating system 410 and a plurality of virtual resources allocated to the at least one operating system 410. Virtual resources may include, without limitation, a plurality of virtual processors 432a, 432b, 432c (generally 432), and virtual disks 442a, 442b, 442c (generally 442), as well as virtual resources such as virtual memory and virtual network interfaces. The plurality of virtual resources and the operating system 410 may be referred to as a virtual machine 406. A virtual machine 406 may include a control operating system 405 in communication with the hypervisor 401 and used to execute applications for managing and configuring other virtual machines on the computing device 100.

In greater detail, a hypervisor 401 may provide virtual resources to an operating system in any manner which simulates the operating system having access to a physical device. A hypervisor 401 may provide virtual resources to any number of guest operating systems 410a, 410b (generally 410). In some embodiments, a computing device 100 executes one or more types of hypervisors. In these embodiments, hypervisors may be used to emulate virtual hardware, partition physical hardware, virtualize physical hardware, and execute virtual machines that provide access to computing environments. Hypervisors may include those manufactured by VMWare, Inc., of Palo Alto, California; the XEN hypervisor, an open source product whose development is overseen by the open source Xen.org community; HyperV, VirtualServer or virtual PC hypervisors provided by Microsoft, or others. In some embodiments, a computing device 100 executing a hypervisor that creates a virtual machine platform on which guest operating systems may execute is referred to as a host server. In one of these embodiments, for example, the computing device 100 is a XEN SERVER provided by Citrix Systems, Inc., of Fort Lauderdale, FL.

In some embodiments, a hypervisor 401 executes within an operating system executing on a computing device. In one of these embodiments, a computing device executing an operating system and a hypervisor 401 may be said to have a host operating system (the operating system executing on the computing device), and a guest operating system (an operating system executing within a computing resource partition provided by the hypervisor 401). In other embodiments, a hypervisor 401 interacts directly with hardware on a computing device, instead of executing on a host operating system. In one of these

embodiments, the hypervisor 401 may be said to be executing on “bare metal,” referring to the hardware comprising the computing device.

In some embodiments, a hypervisor 401 may create a virtual machine 406a-c (generally 406) in which an operating system 410 executes. In one of these embodiments, for example, the hypervisor 401 loads a virtual machine image to create a virtual machine 406. In another of these embodiments, the hypervisor 401 executes an operating system 410 within the virtual machine 406. In still another of these embodiments, the virtual machine 406 executes an operating system 410.

In some embodiments, the hypervisor 401 controls processor scheduling and memory partitioning for a virtual machine 406 executing on the computing device 100. In one of these embodiments, the hypervisor 401 controls the execution of at least one virtual machine 406. In another of these embodiments, the hypervisor 401 presents at least one virtual machine 406 with an abstraction of at least one hardware resource provided by the computing device 100. In other embodiments, the hypervisor 401 controls whether and how physical processor capabilities are presented to the virtual machine 406.

A control operating system 405 may execute at least one application for managing and configuring the guest operating systems. In one embodiment, the control operating system 405 may execute an administrative application, such as an application including a user interface providing administrators with access to functionality for managing the execution of a virtual machine, including functionality for executing a virtual machine, terminating an execution of a virtual machine, or identifying a type of physical resource for allocation to the virtual machine. In another embodiment, the hypervisor 401 executes the control operating system 405 within a virtual machine 406 created by the hypervisor 401. In still another embodiment, the control operating system 405 executes in a virtual machine 406 that is authorized to directly access physical resources on the computing device 100. In some embodiments, a control operating system 405a on a computing device 100a may exchange data with a control operating system 405b on a computing device 100b, via communications between a hypervisor 401a and a hypervisor 401b. In this way, one or more computing devices 100 may exchange data with one or more of the other computing devices 100 regarding processors and other physical resources available in a pool of resources. In one of these embodiments, this functionality allows a hypervisor to manage a pool of resources distributed across a plurality of physical computing devices. In another of these embodiments, multiple hypervisors manage one or more of the guest operating systems executed on one of the computing devices 100.

In one embodiment, the control operating system 405 executes in a virtual machine 406 that is authorized to interact with at least one guest operating system 410. In another embodiment, a guest operating system 410 communicates with the control operating system 405 via the hypervisor 401 in order to request access to a disk or a network. In still another
5 embodiment, the guest operating system 410 and the control operating system 405 may communicate via a communication channel established by the hypervisor 401, such as, for example, via a plurality of shared memory pages made available by the hypervisor 401.

In some embodiments, the control operating system 405 includes a network back-end driver for communicating directly with networking hardware provided by the computing
10 device 100. In one of these embodiments, the network back-end driver processes at least one virtual machine request from at least one guest operating system 110. In other embodiments, the control operating system 405 includes a block back-end driver for communicating with a storage element on the computing device 100. In one of these embodiments, the block back-end driver reads and writes data from the storage element based upon at least one request
15 received from a guest operating system 410.

In one embodiment, the control operating system 405 includes a tools stack 404. In another embodiment, a tools stack 404 provides functionality for interacting with the hypervisor 401, communicating with other control operating systems 405 (for example, on a second computing device 100b), or managing virtual machines 406b, 406c on the computing
20 device 100. In another embodiment, the tools stack 404 includes customized applications for providing improved management functionality to an administrator of a virtual machine farm. In some embodiments, at least one of the tools stack 404 and the control operating system 405 include a management API that provides an interface for remotely configuring and controlling virtual machines 406 running on a computing device 100. In other embodiments,
25 the control operating system 405 communicates with the hypervisor 401 through the tools stack 404.

In one embodiment, the hypervisor 401 executes a guest operating system 410 within a virtual machine 406 created by the hypervisor 401. In another embodiment, the guest operating system 410 provides a user of the computing device 100 with access to resources
30 within a computing environment. In still another embodiment, a resource includes a program, an application, a document, a file, a plurality of applications, a plurality of files, an executable program file, a desktop environment, a computing environment, or other resource made available to a user of the computing device 100. In yet another embodiment, the resource may be delivered to the computing device 100 via a plurality of access methods

including, but not limited to, conventional installation directly on the computing device 100, delivery to the computing device 100 via a method for application streaming, delivery to the computing device 100 of output data generated by an execution of the resource on a second computing device 100' and communicated to the computing device 100 via a presentation layer protocol, delivery to the computing device 100 of output data generated by an execution of the resource via a virtual machine executing on a second computing device 100', or execution from a removable storage device connected to the computing device 100, such as a USB device, or via a virtual machine executing on the computing device 100 and generating output data. In some embodiments, the computing device 100 transmits output data generated by the execution of the resource to another computing device 100'.

In one embodiment, the guest operating system 410, in conjunction with the virtual machine on which it executes, forms a fully-virtualized virtual machine which is not aware that it is a virtual machine; such a machine may be referred to as a "Domain U HVM (Hardware Virtual Machine) virtual machine". In another embodiment, a fully-virtualized machine includes software emulating a Basic Input/Output System (BIOS) in order to execute an operating system within the fully-virtualized machine. In still another embodiment, a fully-virtualized machine may include a driver that provides functionality by communicating with the hypervisor 401. In such an embodiment, the driver may be aware that it executes within a virtualized environment. In another embodiment, the guest operating system 410, in conjunction with the virtual machine on which it executes, forms a paravirtualized virtual machine, which is aware that it is a virtual machine; such a machine may be referred to as a "Domain U PV virtual machine". In another embodiment, a paravirtualized machine includes additional drivers that a fully-virtualized machine does not include. In still another embodiment, the paravirtualized machine includes the network back-end driver and the block back-end driver included in a control operating system 405, as described above.

Referring now to FIG. 4B, a block diagram depicts one embodiment of a plurality of networked computing devices in a system in which at least one physical host executes a virtual machine. In brief overview, the system includes a management component 404 and a hypervisor 401. The system includes a plurality of computing devices 100, a plurality of virtual machines 406, a plurality of hypervisors 401, a plurality of management components referred to variously as tools stacks 404 or management components 404, and a physical resource 421, 428. The plurality of physical machines 100 may each be provided as computing devices 100, described above in connection with FIGs. 1E-1H and 4A.

In greater detail, a physical disk 428 is provided by a computing device 100 and stores at least a portion of a virtual disk 442. In some embodiments, a virtual disk 442 is associated with a plurality of physical disks 428. In one of these embodiments, one or more computing devices 100 may exchange data with one or more of the other computing devices 100

5 regarding processors and other physical resources available in a pool of resources, allowing a hypervisor to manage a pool of resources distributed across a plurality of physical computing devices. In some embodiments, a computing device 100 on which a virtual machine 406 executes is referred to as a physical host 100 or as a host machine 100.

The hypervisor executes on a processor on the computing device 100. The hypervisor
10 allocates, to a virtual disk, an amount of access to the physical disk. In one embodiment, the hypervisor 401 allocates an amount of space on the physical disk. In another embodiment, the hypervisor 401 allocates a plurality of pages on the physical disk. In some embodiments, the hypervisor provisions the virtual disk 442 as part of a process of initializing and executing a virtual machine 450.

15 In one embodiment, the management component 404a is referred to as a pool management component 404a. In another embodiment, a management operating system 405a, which may be referred to as a control operating system 405a, includes the management component. In some embodiments, the management component is referred to as a tools stack. In one of these embodiments, the management component is the tools stack 404
20 described above in connection with FIG. 4A. In other embodiments, the management component 404 provides a user interface for receiving, from a user such as an administrator, an identification of a virtual machine 406 to provision and/or execute. In still other embodiments, the management component 404 provides a user interface for receiving, from a user such as an administrator, the request for migration of a virtual machine 406b from one
25 physical machine 100 to another. In further embodiments, the management component 404a identifies a computing device 100b on which to execute a requested virtual machine 406d and instructs the hypervisor 401b on the identified computing device 100b to execute the identified virtual machine; such a management component may be referred to as a pool management component.

30 Referring now to Figure 4C, embodiments of a virtual application delivery controller or virtual appliance 450 are depicted. In brief overview, any of the functionality and/or embodiments of the appliance 200 (e.g., an application delivery controller) described above in connection with FIGs. 2A and 2B may be deployed in any embodiment of the virtualized environment described above in connection with FIGs 4A and 4B. Instead of the

functionality of the application delivery controller being deployed in the form of an appliance 200, such functionality may be deployed in a virtualized environment 400 on any computing device 100, such as a client 102, server 106 or appliance 200.

Referring now to FIG. 4C, a diagram of an embodiment of a virtual appliance 450 operating on a hypervisor 401 of a server 106 is depicted. As with the appliance 200 of FIGS. 2A and 2B, the virtual appliance 450 may provide functionality for availability, performance, offload and security. For availability, the virtual appliance may perform load balancing between layers 4 and 7 of the network and may also perform intelligent service health monitoring. For performance increases via network traffic acceleration, the virtual appliance may perform caching and compression. To offload processing of any servers, the virtual appliance may perform connection multiplexing and pooling and/or SSL processing. For security, the virtual appliance may perform any of the application firewall functionality and SSL VPN function of appliance 200.

Any of the modules of the appliance 200 as described in connection with FIGS. 2A may be packaged, combined, designed or constructed in a form of the virtualized appliance delivery controller 450 deployable as one or more software modules or components executable in a virtualized environment 300 or non-virtualized environment on any server, such as an off the shelf server. For example, the virtual appliance may be provided in the form of an installation package to install on a computing device. With reference to FIG. 2A, any of the cache manager 232, policy engine 236, compression 238, encryption engine 234, packet engine 240, GUI 210, CLI 212, shell services 214 and health monitoring programs 216 may be designed and constructed as a software component or module to run on any operating system of a computing device and/or of a virtualized environment 300. Instead of using the encryption processor 260, processor 262, memory 264 and network stack 267 of the appliance 200, the virtualized appliance 400 may use any of these resources as provided by the virtualized environment 400 or as otherwise available on the server 106.

Still referring to FIG. 4C, and in brief overview, any one or more vServers 275A-275N may be in operation or executed in a virtualized environment 400 of any type of computing device 100, such as any server 106. Any of the modules or functionality of the appliance 200 described in connection with FIG. 2B may be designed and constructed to operate in either a virtualized or non-virtualized environment of a server. Any of the vServer 275, SSL VPN 280, Intranet UP 282, Switching 284, DNS 286, acceleration 288, App FW 280 and monitoring agent may be packaged, combined, designed or constructed in a form of

application delivery controller 450 deployable as one or more software modules or components executable on a device and/or virtualized environment 400.

In some embodiments, a server may execute multiple virtual machines 406a-406n in the virtualization environment with each virtual machine running the same or different
5 embodiments of the virtual application delivery controller 450. In some embodiments, the server may execute one or more virtual appliances 450 on one or more virtual machines on a core of a multi-core processing system. In some embodiments, the server may execute one or more virtual appliances 450 on one or more virtual machines on each processor of a multiple processor device.

10

E. Systems and Methods for Providing A Multi-Core Architecture

In accordance with Moore's Law, the number of transistors that may be placed on an integrated circuit may double approximately every two years. However, CPU speed increases may reach plateaus, for example CPU speed has been around 3.5 – 4 GHz range
15 since 2005. In some cases, CPU manufacturers may not rely on CPU speed increases to gain additional performance. Some CPU manufacturers may add additional cores to their processors to provide additional performance. Products, such as those of software and networking vendors, that rely on CPUs for performance gains may improve their performance by leveraging these multi-core CPUs. The software designed and constructed
20 for a single CPU may be redesigned and/or rewritten to take advantage of a multi-threaded, parallel architecture or otherwise a multi-core architecture.

A multi-core architecture of the appliance 200 or appliance 205 (which either embodiment may be generally referred to as appliance, appliance 200 or appliance 205), referred to as nCore or multi-core technology, allows the appliance in some embodiments to
25 break the single core performance barrier and to leverage the power of multi-core CPUs. In some embodiments, a single network or packet engine is run. The multiple cores of the nCore technology and architecture allow multiple packet engines to run concurrently and/or in parallel. With a packet engine running on each core, the appliance architecture leverages the processing capacity of additional cores. In some embodiments, this provides up to a 7X
30 increase in performance and scalability.

Illustrated in FIG. 5A are some embodiments of work, task, load or network traffic distribution across one or more processor cores according to a type of parallelism or parallel computing scheme, such as functional parallelism, data parallelism or flow-based data parallelism. In brief overview, FIG. 5A illustrates embodiments of a multi-core system such

as an appliance 200 or 205 with n-cores, a total of cores numbers 1 through N. In one embodiment, work, load or network traffic can be distributed among a first core 505A, a second core 505B, a third core 505C, a fourth core 505D, a fifth core 505E, a sixth core 505F, a seventh core 505G, and so on such that distribution is across all or two or more of the n cores 505N (hereinafter referred to collectively as cores 505.) There may be multiple VIPs 275 each running on a respective core of the plurality of cores. There may be multiple packet engines 240 each running on a respective core of the plurality of cores. Any of the approaches used may lead to different, varying or similar work load or performance level 515 across any of the cores. For a functional parallelism approach, each core may run a different function of the functionalities provided by the packet engine, a VIP 275 or appliance 200. In a data parallelism approach, data may be paralleled or distributed across the cores based on the Network Interface Card (NIC) or VIP 275 receiving the data. In another data parallelism approach, processing may be distributed across the cores by distributing data flows to each core.

In further detail to FIG. 5A, in some embodiments, load, work or network traffic can be distributed among cores 505 according to functional parallelism 500. Functional parallelism may be based on each core performing one or more respective functions. In some embodiments, a first core may perform a first function while a second core performs a second function. In functional parallelism approach, the functions to be performed by the multi-core system are divided and distributed to each core according to functionality. In some embodiments, functional parallelism may be referred to as task parallelism and may be achieved when each processor or core executes a different process or function on the same or different data. The core or processor may execute the same or different code. In some cases, different execution threads or code may communicate with one another as they work.

Communication may take place to pass data from one thread to the next as part of a workflow.

In some embodiments, distributing work across the cores 505 according to functional parallelism 500, can comprise distributing network traffic according to a particular function such as network input/output management (NW I/O) 510A, secure sockets layer (SSL) encryption and decryption 510B and transmission control protocol (TCP) functions 510C. This may lead to a work, performance or computing load 515 based on a volume or level of functionality being used. In some embodiments, distributing work across the cores 505 according to data parallelism 540, can comprise distributing an amount of work 515 based on distributing data associated with a particular hardware or software component. In some

embodiments, distributing work across the cores 505 according to flow-based data parallelism 520, can comprise distributing data based on a context or flow such that the amount of work 515A-N on each core may be similar, substantially equal or relatively evenly distributed.

5 In the case of the functional parallelism approach, each core may be configured to run one or more functionalities of the plurality of functionalities provided by the packet engine or VIP of the appliance. For example, core 1 may perform network I/O processing for the appliance 200' while core 2 performs TCP connection management for the appliance. Likewise, core 3 may perform SSL offloading while core 4 may perform layer 7 or
10 application layer processing and traffic management. Each of the cores may perform the same function or different functions. Each of the cores may perform more than one function. Any of the cores may run any of the functionality or portions thereof identified and/or described in conjunction with FIGs. 2A and 2B. In this the approach, the work across the cores may be divided by function in either a coarse-grained or fine-grained manner. In some
15 cases, as illustrated in FIG. 5A, division by function may lead to different cores running at different levels of performance or load 515.

 In the case of the functional parallelism approach, each core may be configured to run one or more functionalities of the plurality of functionalities provided by the packet engine of the appliance. For example, core 1 may perform network I/O processing for the appliance
20 200' while core 2 performs TCP connection management for the appliance. Likewise, core 3 may perform SSL offloading while core 4 may perform layer 7 or application layer processing and traffic management. Each of the cores may perform the same function or different functions. Each of the cores may perform more than one function. Any of the cores may run any of the functionality or portions thereof identified and/or described in conjunction
25 with FIGs. 2A and 2B. In this the approach, the work across the cores may be divided by function in either a coarse-grained or fine-grained manner. In some cases, as illustrated in FIG. 5A division by function may lead to different cores running at different levels of load or performance.

 The functionality or tasks may be distributed in any arrangement and scheme. For
30 example, FIG. 5B illustrates a first core, Core 1 505A, processing applications and processes associated with network I/O functionality 510A. Network traffic associated with network I/O, in some embodiments, can be associated with a particular port number. Thus, outgoing and incoming packets having a port destination associated with NW I/O 510A will be directed towards Core 1 505A which is dedicated to handling all network traffic associated

with the NW I/O port. Similarly, Core 2 505B is dedicated to handling functionality associated with SSL processing and Core 4 505D may be dedicated handling all TCP level processing and functionality.

While FIG. 5A illustrates functions such as network I/O, SSL and TCP, other functions can be assigned to cores. These other functions can include any one or more of the functions or operations described herein. For example, any of the functions described in conjunction with FIGs. 2A and 2B may be distributed across the cores on a functionality basis. In some cases, a first VIP 275A may run on a first core while a second VIP 275B with a different configuration may run on a second core. In some embodiments, each core 505 can handle a particular functionality such that each core 505 can handle the processing associated with that particular function. For example, Core 2 505B may handle SSL offloading while Core 4 505D may handle application layer processing and traffic management.

In other embodiments, work, load or network traffic may be distributed among cores 505 according to any type and form of data parallelism 540. In some embodiments, data parallelism may be achieved in a multi-core system by each core performing the same task or functionally on different pieces of distributed data. In some embodiments, a single execution thread or code controls operations on all pieces of data. In other embodiments, different threads or instructions control the operation, but may execute the same code. In some embodiments, data parallelism is achieved from the perspective of a packet engine, vServers (VIPs) 275A-C, network interface cards (NIC) 542D-E and/or any other networking hardware or software included on or associated with an appliance 200. For example, each core may run the same packet engine or VIP code or configuration but operate on different sets of distributed data. Each networking hardware or software construct can receive different, varying or substantially the same amount of data, and as a result may have varying, different or relatively the same amount of load 515.

In the case of a data parallelism approach, the work may be divided up and distributed based on VIPs, NICs and/or data flows of the VIPs or NICs. In one of these approaches, the work of the multi-core system may be divided or distributed among the VIPs by having each VIP work on a distributed set of data. For example, each core may be configured to run one or more VIPs. Network traffic may be distributed to the core for each VIP handling that traffic. In another of these approaches, the work of the appliance may be divided or distributed among the cores based on which NIC receives the network traffic. For example, network traffic of a first NIC may be distributed to a first core while network traffic of a

second NIC may be distributed to a second core. In some cases, a core may process data from multiple NICs.

While FIG 5A illustrates a single vServer associated with a single core 505, as is the case for VIP1 275A, VIP2 275B and VIP3 275C. In some embodiments, a single vServer can be associated with one or more cores 505. In contrast, one or more vServers can be associated with a single core 505. Associating a vServer with a core 505 may include that core 505 to process all functions associated with that particular vServer. In some embodiments, each core executes a VIP having the same code and configuration. In other embodiments, each core executes a VIP having the same code but different configuration. In some embodiments, each core executes a VIP having different code and the same or different configuration.

Like vServers, NICs can also be associated with particular cores 505. In many embodiments, NICs can be connected to one or more cores 505 such that when a NIC receives or transmits data packets, a particular core 505 handles the processing involved with receiving and transmitting the data packets. In one embodiment, a single NIC can be associated with a single core 505, as is the case with NIC1 542D and NIC2 542E. In other embodiments, one or more NICs can be associated with a single core 505. In other embodiments, a single NIC can be associated with one or more cores 505. In these embodiments, load could be distributed amongst the one or more cores 505 such that each core 505 processes a substantially similar amount of load. A core 505 associated with a NIC may process all functions and/or data associated with that particular NIC.

While distributing work across cores based on data of VIPs or NICs may have a level of independency, in some embodiments, this may lead to unbalanced use of cores as illustrated by the varying loads 515 of FIG. 5A.

In some embodiments, load, work or network traffic can be distributed among cores 505 based on any type and form of data flow. In another of these approaches, the work may be divided or distributed among cores based on data flows. For example, network traffic between a client and a server traversing the appliance may be distributed to and processed by one core of the plurality of cores. In some cases, the core initially establishing the session or connection may be the core for which network traffic for that session or connection is distributed. In some embodiments, the data flow is based on any unit or portion of network traffic, such as a transaction, a request/response communication or traffic originating from an application on a client. In this manner and in some embodiments, data flows between clients

and servers traversing the appliance 200' may be distributed in a more balanced manner than the other approaches.

In flow-based data parallelism 520, distribution of data is related to any type of flow of data, such as request/response pairings, transactions, sessions, connections or application communications. For example, network traffic between a client and a server traversing the appliance may be distributed to and processed by one core of the plurality of cores. In some cases, the core initially establishing the session or connection may be the core for which network traffic for that session or connection is distributed. The distribution of data flow may be such that each core 505 carries a substantially equal or relatively evenly distributed amount of load, data or network traffic..

In some embodiments, the data flow is based on any unit or portion of network traffic, such as a transaction, a request/response communication or traffic originating from an application on a client. In this manner and in some embodiments, data flows between clients and servers traversing the appliance 200' may be distributed in a more balanced manner than the other approached. In one embodiment, data flow can be distributed based on a transaction or a series of transactions. This transaction, in some embodiments, can be between a client and a server and can be characterized by an IP address or other packet identifier. For example, Core 1 505A can be dedicated to transactions between a particular client and a particular server, therefore the load 515A on Core 1 505A may be comprised of the network traffic associated with the transactions between the particular client and server. Allocating the network traffic to Core 1 505A can be accomplished by routing all data packets originating from either the particular client or server to Core 1 505A.

While work or load can be distributed to the cores based in part on transactions, in other embodiments load or work can be allocated on a per packet basis. In these embodiments, the appliance 200 can intercept data packets and allocate them to a core 505 having the least amount of load. For example, the appliance 200 could allocate a first incoming data packet to Core 1 505A because the load 515A on Core 1 is less than the load 515B-N on the rest of the cores 505B-N. Once the first data packet is allocated to Core 1 505A, the amount of load 515A on Core 1 505A is increased proportional to the amount of processing resources needed to process the first data packet. When the appliance 200 intercepts a second data packet, the appliance 200 will allocate the load to Core 4 505D because Core 4 505D has the second least amount of load. Allocating data packets to the core with the least amount of load can, in some embodiments, ensure that the load 515A-N distributed to each core 505 remains substantially equal.

In other embodiments, load can be allocated on a per unit basis where a section of network traffic is allocated to a particular core 505. The above-mentioned example illustrates load balancing on a per/packet basis. In other embodiments, load can be allocated based on a number of packets such that every 10, 100 or 1000 packets are allocated to the core 505

5 having the least amount of load. The number of packets allocated to a core 505 can be a number determined by an application, user or administrator and can be any number greater than zero. In still other embodiments, load can be allocated based on a time metric such that packets are distributed to a particular core 505 for a predetermined amount of time. In these
10 embodiments, packets can be distributed to a particular core 505 for five milliseconds or for any period of time determined by a user, program, system, administrator or otherwise. After the predetermined time period elapses, data packets are transmitted to a different core 505 for the predetermined period of time.

Flow-based data parallelism methods for distributing work, load or network traffic among the one or more cores 505 can comprise any combination of the above-mentioned
15 embodiments. These methods can be carried out by any part of the appliance 200, by an application or set of executable instructions executing on one of the cores 505, such as the packet engine, or by any application, program or agent executing on a computing device in communication with the appliance 200.

The functional and data parallelism computing schemes illustrated in FIG. 5A can be
20 combined in any manner to generate a hybrid parallelism or distributed processing scheme that encompasses function parallelism 500, data parallelism 540, flow-based data parallelism 520 or any portions thereof. In some cases, the multi-core system may use any type and form of load balancing schemes to distribute load among the one or more cores 505. The load balancing scheme may be used in any combination with any of the functional and data
25 parallelism schemes or combinations thereof.

Illustrated in FIG. 5B is an embodiment of a multi-core system 545, which may be any type and form of one or more systems, appliances, devices or components. This system 545, in some embodiments, can be included within an appliance 200 having one or more processing cores 505A-N. The system 545 can further include one or more packet engines
30 (PE) or packet processing engines (PPE) 548A-N communicating with a memory bus 556. The memory bus may be used to communicate with the one or more processing cores 505A-N. Also included within the system 545 can be one or more network interface cards (NIC) 552 and a flow distributor 550 which can further communicate with the one or more

processing cores 505A-N. The flow distributor 550 can comprise a Receive Side Scaler (RSS) or Receive Side Scaling (RSS) module 560.

Further referring to FIG. 5B, and in more detail, in one embodiment the packet engine(s) 548A-N can comprise any portion of the appliance 200 described herein, such as any portion of the appliance described in FIGs. 2A and 2B. The packet engine(s) 548A-N can, in some embodiments, comprise any of the following elements: the packet engine 240, a network stack 267; a cache manager 232; a policy engine 236; a compression engine 238; an encryption engine 234; a GUI 210; a CLI 212; shell services 214; monitoring programs 216; and any other software or hardware element able to receive data packets from one of either the memory bus 556 or the one of more cores 505A-N. In some embodiments, the packet engine(s) 548A-N can comprise one or more vServers 275A-N, or any portion thereof. In other embodiments, the packet engine(s) 548A-N can provide any combination of the following functionalities: SSL VPN 280; Intranet UP 282; switching 284; DNS 286; packet acceleration 288; App FW 280; monitoring such as the monitoring provided by a monitoring agent 197; functionalities associated with functioning as a TCP stack; load balancing; SSL offloading and processing; content switching; policy evaluation; caching; compression; encoding; decompression; decoding; application firewall functionalities; XML processing and acceleration; and SSL VPN connectivity.

The packet engine(s) 548A-N can, in some embodiments, be associated with a particular server, user, client or network. When a packet engine 548 becomes associated with a particular entity, that packet engine 548 can process data packets associated with that entity. For example, should a packet engine 548 be associated with a first user, that packet engine 548 will process and operate on packets generated by the first user, or packets having a destination address associated with the first user. Similarly, the packet engine 548 may choose not to be associated with a particular entity such that the packet engine 548 can process and otherwise operate on any data packets not generated by that entity or destined for that entity.

In some instances, the packet engine(s) 548A-N can be configured to carry out the any of the functional and/or data parallelism schemes illustrated in FIG. 5A. In these instances, the packet engine(s) 548A-N can distribute functions or data among the processing cores 505A-N so that the distribution is according to the parallelism or distribution scheme. In some embodiments, a single packet engine(s) 548A-N carries out a load balancing scheme, while in other embodiments one or more packet engine(s) 548A-N carry out a load balancing scheme. Each core 505A-N, in one embodiment, can be associated with a particular packet

engine 548 such that load balancing can be carried out by the packet engine. Load balancing may in this embodiment, require that each packet engine 548A-N associated with a core 505 communicate with the other packet engines associated with cores so that the packet engines 548A-N can collectively determine where to distribute load. One embodiment of this process
5 can include an arbiter that receives votes from each packet engine for load. The arbiter can distribute load to each packet engine 548A-N based in part on the age of the engine's vote and in some cases a priority value associated with the current amount of load on an engine's associated core 505.

Any of the packet engines running on the cores may run in user mode, kernel or any
10 combination thereof. In some embodiments, the packet engine operates as an application or program running in user or application space. In these embodiments, the packet engine may use any type and form of interface to access any functionality provided by the kernel. In some embodiments, the packet engine operates in kernel mode or as part of the kernel. In some embodiments, a first portion of the packet engine operates in user mode while a second
15 portion of the packet engine operates in kernel mode. In some embodiments, a first packet engine on a first core executes in kernel mode while a second packet engine on a second core executes in user mode. In some embodiments, the packet engine or any portions thereof operates on or in conjunction with the NIC or any drivers thereof.

In some embodiments the memory bus 556 can be any type and form of memory or
20 computer bus. While a single memory bus 556 is depicted in FIG. 5B, the system 545 can comprise any number of memory buses 556. In one embodiment, each packet engine 548 can be associated with one or more individual memory buses 556.

The NIC 552 can in some embodiments be any of the network interface cards or mechanisms described herein. The NIC 552 can have any number of ports. The NIC can be
25 designed and constructed to connect to any type and form of network 104. While a single NIC 552 is illustrated, the system 545 can comprise any number of NICs 552. In some embodiments, each core 505A-N can be associated with one or more single NICs 552. Thus, each core 505 can be associated with a single NIC 552 dedicated to a particular core 505. The cores 505A-N can comprise any of the processors described herein. Further, the cores
30 505A-N can be configured according to any of the core 505 configurations described herein. Still further, the cores 505A-N can have any of the core 505 functionalities described herein. While FIG. 5B illustrates seven cores 505A-G, any number of cores 505 can be included within the system 545. In particular, the system 545 can comprise "N" cores, where "N" is a whole number greater than zero.

A core may have or use memory that is allocated or assigned for use to that core. The memory may be considered private or local memory of that core and only accessible by that core. A core may have or use memory that is shared or assigned to multiple cores. The memory may be considered public or shared memory that is accessible by more than one
5 core. A core may use any combination of private and public memory. With separate address spaces for each core, some level of coordination is eliminated from the case of using the same address space. With a separate address space, a core can perform work on information and data in the core's own address space without worrying about conflicts with other cores. Each packet engine may have a separate memory pool for TCP and/or SSL connections.

10 Further referring to FIG. 5B, any of the functionality and/or embodiments of the cores 505 described above in connection with FIG. 5A can be deployed in any embodiment of the virtualized environment described above in connection with FIGs. 4A and 4B. Instead of the functionality of the cores 505 being deployed in the form of a physical processor 505, such functionality may be deployed in a virtualized environment 400 on any computing device
15 100, such as a client 102, server 106 or appliance 200. In other embodiments, instead of the functionality of the cores 505 being deployed in the form of an appliance or a single device, the functionality may be deployed across multiple devices in any arrangement. For example, one device may comprise two or more cores and another device may comprise two or more cores. For example, a multi-core system may include a cluster of computing devices, a server
20 farm or network of computing devices. In some embodiments, instead of the functionality of the cores 505 being deployed in the form of cores, the functionality may be deployed on a plurality of processors, such as a plurality of single core processors.

In one embodiment, the cores 505 may be any type and form of processor. In some embodiments, a core can function substantially similar to any processor or central processing
25 unit described herein. In some embodiment, the cores 505 may comprise any portion of any processor described herein. While FIG. 5A illustrates seven cores, there can exist any "N" number of cores within an appliance 200, where "N" is any whole number greater than one. In some embodiments, the cores 505 can be installed within a common appliance 200, while in other embodiments the cores 505 can be installed within one or more appliance(s) 200
30 communicatively connected to one another. The cores 505 can in some embodiments comprise graphics processing software, while in other embodiments the cores 505 provide general processing capabilities. The cores 505 can be installed physically near each other and/or can be communicatively connected to each other. The cores may be connected by any

type and form of bus or subsystem physically and/or communicatively coupled to the cores for transferring data between to, from and/or between the cores.

While each core 505 can comprise software for communicating with other cores, in some embodiments a core manager (not shown) can facilitate communication between each core 505. In some embodiments, the kernel may provide core management. The cores may interface or communicate with each other using a variety of interface mechanisms. In some embodiments, core to core messaging may be used to communicate between cores, such as a first core sending a message or data to a second core via a bus or subsystem connecting the cores. In some embodiments, cores may communicate via any type and form of shared memory interface. In one embodiment, there may be one or more memory locations shared among all the cores. In some embodiments, each core may have separate memory locations shared with each other core. For example, a first core may have a first shared memory with a second core and a second share memory with a third core. In some embodiments, cores may communicate via any type of programming or API, such as function calls via the kernel. In some embodiments, the operating system may recognize and support multiple core devices and provide interfaces and API for inter-core communications.

The flow distributor 550 can be any application, program, library, script, task, service, process or any type and form of executable instructions executing on any type and form of hardware. In some embodiments, the flow distributor 550 may any design and construction of circuitry to perform any of the operations and functions described herein. In some embodiments, the flow distributor distribute, forwards, routes, controls and/or manage the distribution of data packets among the cores 505 and/or packet engine or VIPs running on the cores.. The flow distributor 550, in some embodiments, can be referred to as an interface master. In one embodiment, the flow distributor 550 comprises a set of executable instructions executing on a core or processor of the appliance 200. In another embodiment, the flow distributor 550 comprises a set of executable instructions executing on a computing machine in communication with the appliance 200. In some embodiments, the flow distributor 550 comprises a set of executable instructions executing on a NIC, such as firmware. In still other embodiments, the flow distributor 550 comprises any combination of software and hardware to distribute data packets among cores or processors. In one embodiment, the flow distributor 550 executes on at least one of the cores 505A-N, while in other embodiments a separate flow distributor 550 assigned to each core 505A-N executes on an associated core 505A-N. The flow distributor may use any type and form of statistical or probabilistic algorithms or decision making to balance the flows across the cores. The

hardware of the appliance, such as a NIC, or the kernel may be designed and constructed to support sequential operations across the NICs and/or cores.

In embodiments where the system 545 comprises one or more flow distributors 550, each flow distributor 550 can be associated with a processor 505 or a packet engine 548. The flow distributors 550 can comprise an interface mechanism that allows each flow distributor 550 to communicate with the other flow distributors 550 executing within the system 545. In one instance, the one or more flow distributors 550 can determine how to balance load by communicating with each other. This process can operate substantially similarly to the process described above for submitting votes to an arbiter which then determines which flow distributor 550 should receive the load. In other embodiments, a first flow distributor 550' can identify the load on an associated core and determine whether to forward a first data packet to the associated core based on any of the following criteria: the load on the associated core is above a predetermined threshold; the load on the associated core is below a predetermined threshold; the load on the associated core is less than the load on the other cores; or any other metric that can be used to determine where to forward data packets based in part on the amount of load on a processor.

The flow distributor 550 can distribute network traffic among the cores 505 according to a distribution, computing or load balancing scheme such as those described herein. In one embodiment, the flow distributor can distribute network traffic according to any one of a functional parallelism distribution scheme 550, a data parallelism load distribution scheme 540, a flow-based data parallelism distribution scheme 520, or any combination of these distribution scheme or any load balancing scheme for distributing load among multiple processors. The flow distributor 550 can therefore act as a load distributor by taking in data packets and distributing them across the processors according to an operative load balancing or distribution scheme. In one embodiment, the flow distributor 550 can comprise one or more operations, functions or logic to determine how to distribute packers, work or load accordingly. In still other embodiments, the flow distributor 550 can comprise one or more sub operations, functions or logic that can identify a source address and a destination address associated with a data packet, and distribute packets accordingly.

In some embodiments, the flow distributor 550 can comprise a receive-side scaling (RSS) network driver, module 560 or any type and form of executable instructions which distribute data packets among the one or more cores 505. The RSS module 560 can comprise any combination of hardware and software, In some embodiments, the RSS module 560 works in conjunction with the flow distributor 550 to distribute data packets across the cores

505A-N or among multiple processors in a multi-processor network. The RSS module 560 can execute within the NIC 552 in some embodiments, and in other embodiments can execute on any one of the cores 505.

In some embodiments, the RSS module 560 uses the MICROSOFT receive-side-scaling (RSS) scheme. In one embodiment, RSS is a Microsoft Scalable Networking initiative technology that enables receive processing to be balanced across multiple processors in the system while maintaining in-order delivery of the data. The RSS may use any type and form of hashing scheme to determine a core or processor for processing a network packet.

The RSS module 560 can apply any type and form hash function such as the Toeplitz hash function. The hash function may be applied to the hash type or any the sequence of values. The hash function may be a secure hash of any security level or is otherwise cryptographically secure. The hash function may use a hash key. The size of the key is dependent upon the hash function. For the Toeplitz hash, the size may be 40 bytes for IPv6 and 16 bytes for IPv4.

The hash function may be designed and constructed based on any one or more criteria or design goals. In some embodiments, a hash function may be used that provides an even distribution of hash result for different hash inputs and different hash types, including TCP/IPv4, TCP/IPv6, IPv4, and IPv6 headers. In some embodiments, a hash function may be used that provides a hash result that is evenly distributed when a small number of buckets are present (for example, two or four). In some embodiments, hash function may be used that provides a hash result that is randomly distributed when a large number of buckets were present (for example, 64 buckets). In some embodiments, the hash function is determined based on a level of computational or resource usage. In some embodiments, the hash function is determined based on ease or difficulty of implementing the hash in hardware. In some embodiments, the hash function is determined based on the ease or difficulty of a malicious remote host to send packets that would all hash to the same bucket.

The RSS may generate hashes from any type and form of input, such as a sequence of values. This sequence of values can include any portion of the network packet, such as any header, field or payload of network packet, or portions thereof. In some embodiments, the input to the hash may be referred to as a hash type and include any tuples of information associated with a network packet or data flow, such as any of the following: a four tuple comprising at least two IP addresses and two ports; a four tuple comprising any four sets of

values; a six tuple; a two tuple; and/or any other sequence of numbers or values. The following are example of hash types that may be used by RSS:

- 4-tuple of source TCP Port, source IP version 4 (IPv4) address, destination TCP Port, and destination IPv4 address.
- 5 - 4-tuple of source TCP Port, source IP version 6 (IPv6) address, destination TCP Port, and destination IPv6 address.
- 2-tuple of source IPv4 address, and destination IPv4 address.
- 2-tuple of source IPv6 address, and destination IPv6 address.
- 10 - 2-tuple of source IPv6 address, and destination IPv6 address, including support for parsing IPv6 extension headers.

The hash result or any portion thereof may used to identify a core or entity, such as a packet engine or VIP, for distributing a network packet. In some embodiments, one or more hash bits or mask are applied to the hash result. The hash bit or mask may be any number of bits or bytes. A NIC may support any number of bits, such as seven bits. The network stack
15 may set the actual number of bits to be used during initialization. The number will be between 1 and 7, inclusive.

The hash result may be used to identify the core or entity via any type and form of table, such as a bucket table or indirection table. In some embodiments, the number of hash-result bits are used to index into the table. The range of the hash mask may effectively define
20 the size of the indirection table. Any portion of the hash result or the hash result itself may be used to index the indirection table. The values in the table may identify any of the cores or processor, such as by a core or processor identifier. In some embodiments, all of the cores of the multi-core system are identified in the table. In other embodiments, a port of the cores of the multi-core system are identified in the table. The indirection table may comprise any
25 number of buckets for example 2 to 128 buckets that may be indexed by a hash mask. Each bucket may comprise a range of index values that identify a core or processor. In some embodiments, the flow controller and/or RSS module may rebalance the network rebalance the network load by changing the indirection table.

In some embodiments, the multi-core system 575 does not include a RSS driver or
30 RSS module 560. In some of these embodiments, a software steering module (not shown) or a software embodiment of the RSS module within the system can operate in conjunction with

or as part of the flow distributor 550 to steer packets to cores 505 within the multi-core system 575.

The flow distributor 550, in some embodiments, executes within any module or program on the appliance 200, on any one of the cores 505 and on any one of the devices or components included within the multi-core system 575. In some embodiments, the flow distributor 550' can execute on the first core 505A, while in other embodiments the flow distributor 550'' can execute on the NIC 552. In still other embodiments, an instance of the flow distributor 550' can execute on each core 505 included in the multi-core system 575. In this embodiment, each instance of the flow distributor 550' can communicate with other instances of the flow distributor 550' to forward packets back and forth across the cores 505. There exist situations where a response to a request packet may not be processed by the same core, i.e. the first core processes the request while the second core processes the response. In these situations, the instances of the flow distributor 550' can intercept the packet and forward it to the desired or correct core 505, i.e. a flow distributor instance 550' can forward the response to the first core. Multiple instances of the flow distributor 550' can execute on any number of cores 505 and any combination of cores 505.

The flow distributor may operate responsive to any one or more rules or policies. The rules may identify a core or packet processing engine to receive a network packet, data or data flow. The rules may identify any type and form of tuple information related to a network packet, such as a 4-tuple of source and destination IP address and source and destination ports. Based on a received packet matching the tuple specified by the rule, the flow distributor may forward the packet to a core or packet engine. In some embodiments, the packet is forwarded to a core via shared memory and/or core to core messaging.

Although FIG. 5B illustrates the flow distributor 550 as executing within the multi-core system 575, in some embodiments the flow distributor 550 can execute on a computing device or appliance remotely located from the multi-core system 575. In such an embodiment, the flow distributor 550 can communicate with the multi-core system 575 to take in data packets and distribute the packets across the one or more cores 505. The flow distributor 550 can, in one embodiment, receive data packets destined for the appliance 200, apply a distribution scheme to the received data packets and distribute the data packets to the one or more cores 505 of the multi-core system 575. In one embodiment, the flow distributor 550 can be included in a router or other appliance such that the router can target particular cores 505 by altering meta data associated with each packet so that each packet is targeted

towards a sub-node of the multi-core system 575. In such an embodiment, CISCO's vn-tag mechanism can be used to alter or tag each packet with the appropriate meta data.

Illustrated in FIG. 5C is an embodiment of a multi-core system 575 comprising one or more processing cores 505A-N. In brief overview, one of the cores 505 can be designated as a control core 505A and can be used as a control plane 570 for the other cores 505. The other cores may be secondary cores which operate in a data plane while the control core provides the control plane. The cores 505A-N may share a global cache 580. While the control core provides a control plane, the other cores in the multi-core system form or provide a data plane. These cores perform data processing functionality on network traffic while the control provides initialization, configuration and control of the multi-core system.

Further referring to FIG. 5C, and in more detail, the cores 505A-N as well as the control core 505A can be any processor described herein. Furthermore, the cores 505A-N and the control core 505A can be any processor able to function within the system 575 described in FIG. 5C. Still further, the cores 505A-N and the control core 505A can be any core or group of cores described herein. The control core may be a different type of core or processor than the other cores. In some embodiments, the control may operate a different packet engine or have a packet engine configured differently than the packet engines of the other cores.

Any portion of the memory of each of the cores may be allocated to or used for a global cache that is shared by the cores. In brief overview, a predetermined percentage or predetermined amount of each of the memory of each core may be used for the global cache. For example, 50% of each memory of each code may be dedicated or allocated to the shared global cache. That is, in the illustrated embodiment, 2GB of each core excluding the control plane core or core 1 may be used to form a 28GB shared global cache. The configuration of the control plane such as via the configuration services may determine the amount of memory used for the shared global cache. In some embodiments, each core may provide a different amount of memory for use by the global cache. In other embodiments, any one core may not provide any memory or use the global cache. In some embodiments, any of the cores may also have a local cache in memory not allocated to the global shared memory. Each of the cores may store any portion of network traffic to the global shared cache. Each of the cores may check the cache for any content to use in a request or response. Any of the cores may obtain content from the global shared cache to use in a data flow, request or response.

The global cache 580 can be any type and form of memory or storage element, such as any memory or storage element described herein. In some embodiments, the cores 505

may have access to a predetermined amount of memory (i.e. 32 GB or any other memory amount commensurate with the system 575). The global cache 580 can be allocated from that predetermined amount of memory while the rest of the available memory can be allocated among the cores 505. In other embodiments, each core 505 can have a

5 predetermined amount of memory. The global cache 580 can comprise an amount of the memory allocated to each core 505. This memory amount can be measured in bytes, or can be measured as a percentage of the memory allocated to each core 505. Thus, the global cache 580 can comprise 1 GB of memory from the memory associated with each core 505, or can comprise 20 percent or one-half of the memory associated with each core 505. In some
10 embodiments, only a portion of the cores 505 provide memory to the global cache 580, while in other embodiments the global cache 580 can comprise memory not allocated to the cores 505.

Each core 505 can use the global cache 580 to store network traffic or cache data. In some embodiments, the packet engines of the core use the global cache to cache and use data
15 stored by the plurality of packet engines. For example, the cache manager of FIG. 2A and cache functionality of FIG. 2B may use the global cache to share data for acceleration. For example, each of the packet engines may store responses, such as HTML data, to the global cache. Any of the cache managers operating on a core may access the global cache to server caches responses to client requests.

20 In some embodiments, the cores 505 can use the global cache 580 to store a port allocation table which can be used to determine data flow based in part on ports. In other embodiments, the cores 505 can use the global cache 580 to store an address lookup table or any other table or list that can be used by the flow distributor to determine where to direct incoming and outgoing data packets. The cores 505 can, in some embodiments read from
25 and write to cache 580, while in other embodiments the cores 505 can only read from or write to cache 580. The cores may use the global cache to perform core to core communications.

The global cache 580 may be sectioned into individual memory sections where each section can be dedicated to a particular core 505. In one embodiment, the control core 505A can receive a greater amount of available cache, while the other cores 505 can receiving
30 varying amounts or access to the global cache 580.

In some embodiments, the system 575 can comprise a control core 505A. While FIG. 5C illustrates core 1 505A as the control core, the control core can be any core within the appliance 200 or multi-core system. Further, while only a single control core is depicted, the system 575 can comprise one or more control cores each having a level of control over the

system. In some embodiments, one or more control cores can each control a particular aspect of the system 575. For example, one core can control deciding which distribution scheme to use, while another core can determine the size of the global cache 580.

The control plane of the multi-core system may be the designation and configuration of a core as the dedicated management core or as a master core. This control plane core may provide control, management and coordination of operation and functionality the plurality of cores in the multi-core system. This control plane core may provide control, management and coordination of allocation and use of memory of the system among the plurality of cores in the multi-core system, including initialization and configuration of the same. In some embodiments, the control plane includes the flow distributor for controlling the assignment of data flows to cores and the distribution of network packets to cores based on data flows. In some embodiments, the control plane core runs a packet engine and in other embodiments, the control plane core is dedicated to management and control of the other cores of the system.

The control core 505A can exercise a level of control over the other cores 505 such as determining how much memory should be allocated to each core 505 or determining which core 505 should be assigned to handle a particular function or hardware/software entity. The control core 505A, in some embodiments, can exercise control over those cores 505 within the control plan 570. Thus, there can exist processors outside of the control plane 570 which are not controlled by the control core 505A. Determining the boundaries of the control plane 570 can include maintaining, by the control core 505A or agent executing within the system 575, a list of those cores 505 controlled by the control core 505A. The control core 505A can control any of the following: initialization of a core; determining when a core is unavailable; re-distributing load to other cores 505 when one core fails; determining which distribution scheme to implement; determining which core should receive network traffic; determining how much cache should be allocated to each core; determining whether to assign a particular function or element to a particular core; determining whether to permit cores to communicate with one another; determining the size of the global cache 580; and any other determination of a function, configuration or operation of the cores within the system 575.

E. Systems and Methods for Dynamic Adaptation of Virtual Appliances

An appliance 200 may provide a virtualized networking platform. The virtualized networking platform may be designed and constructed to run a plurality of virtual appliances

in a virtualized environment 400, such a virtual application delivery controller 460 and/or a virtual networking optimization engine 450. The virtual application delivery controller and virtual network optimization engine may be network accelerators. The virtual application delivery controller may accelerate LAN traffic while the virtual network optimization engine accelerates WAN traffic. The virtualized networking platform may enable and support multiple instances of these different types of virtual appliances to run on a single device or appliance. The virtualized networking platform may support multi-tenancy that includes full resource isolation, per-instance fault tolerance, version control, data separation and policy management. The appliance may comprise a multi-core architecture in which instances of a virtual appliance execute on cores of the multi-core appliance.

In some deployments, the virtualized networking platform of the appliance may execute one or more virtualized application delivery controllers 460 that manage and/or load balance a plurality of virtualized network optimization engines 450 providing WAN optimization. For example, in a multi-core appliance 200, one core may execute one or more virtualized application delivery controllers 460 and each instance of the plurality of virtualized network optimization engines 450 (referred to sometimes as a branch repeater (BR) or WAN optimizer) may operate on other cores.

Systems and methods of the present solution provide an adaptable solution that dynamically and automatically deploys, configures and adjusts WAN optimizers such as virtualized WAN optimization engines based on changing network workload. The present solution may start in a learning mode. In such a mode, the system may examine traffic passively or by doing sub-optimal acceleration and learn information about the load and traffic. From examining traffic, the present solution may determine how many peer WAN optimizers are being serviced, how much traffic is coming from each peer WAN optimizers, and the type of traffic traversing the WAN optimizers. From this learning, the present solution may provide a better or improved baseline for the configuration of an appliance. In some embodiments, based on resources (e.g., CPU, Memory, Disk), the system from this knowledge may determine how many WAN optimization instances should be used and of what size, and how the load should be distributed across the instances of the WAN optimizer.

In some embodiments, a WAN optimizer may query, detect, or recognize the amount of resources, such as memory, storage and CPU, available on the appliance. Based on the results, the WAN optimizer may auto-configure itself. A WAN optimizer may select or determine the amount of resources allocated to itself based on the resources available on the appliance. For example, the WAN optimizer may allocate a predetermined amount or

percentage of each of the resources to itself based on the amount of each of the resources available on the appliance and/or based on the number of WAN optimizers executing or to execute on the appliance.

5 In some embodiments, a manager on the appliance may detect query, detect, or recognize the amount of resources and may configure a number of instances of WAN optimizers to execute on the appliance. In some embodiments, a manager on the appliance may detect query, detect, or recognize the amount of resources and may configure the size of each instance of a WAN optimizer. The manager may allocate a predetermined amount or percentage of each of the resources among the WAN optimizers based on the amount of each
10 of the resources available on the appliance and/or based on the number of WAN optimizers executing or to execute on the appliance.

In some embodiments, the system monitors and stores data on performance and operation of the WAN optimizers, such as load and bandwidth usage. In some embodiments, the system monitors and stores data on resource usage by each WAN optimizer. In some
15 embodiments, the system monitors and stores data on changed in resource usage by each WAN optimizer. Based on historical data, the system may auto-configure the number of instances of WAN optimizers to execute on the appliance and auto-configure each instance of a WAN optimizer. For example, based on historical data on resource usage, the system may auto-configure the allocation of each of the different types of resources across each WAN
20 optimizer. In another example, based on historical data on performance, the system may auto-configure to establish a predetermined number of instances of a WAN optimizer to meet an expected load or bandwidth usage by the system.

Referring now to FIG. 6A, an embodiment of an appliance providing a virtualized networking platform is depicted. An appliance 200 may be used to provide a virtualized
25 networking platform for executing a plurality of virtual appliances, such as the network optimization or WAN optimization engines 250/450 described herein. In brief overview, the appliance 200 may include a plurality of resources, such as networking resources, computing resources and storage resources. The appliance may provide a virtualized environment 400 that virtualizes the user of these resources. One or more hypervisors 401 may operate in the
30 virtualized environment. A plurality of virtual appliances (VAs or VMs) may execute in the virtualized environment. An application delivery controller VA 460 may execute on the appliance, such as on a core in multi-core embodiments of the appliance. The application delivery controller VA 460 may manage or load balance a plurality of WAN optimization VAs 450A-450N. The virtualized environment may include one or more management

services 404A-N for managing the virtualized environment and components thereof. As further discussed below, a packet schedule domain may handle network traffic between virtual appliances and the network resources of the appliance.

For virtual machines to communicate with other virtual machines or external hosts, the virtual machines may use network interfaces or network interface cards (NICs). In some embodiments, a virtualized environment may provide two types of Networks interfaces: para-virtualized NICs referred to as PV NICs and emulated NICs (referred to as VNICs or virtual NICS). When a virtual machine (VM) transmits a packet on a VNIC, Domain0 receives the packet and forwards the packet on a real physical NIC (Domain0 may be a special VM in the virtualized environment). Similarly when the virtual machine receives a packet, the VM forwards the packet to the appropriate VNIC and the VM to which the VNIC belongs receives the packet.

However, if a virtual machine requires very high networking performance, the virtualized networking architecture may become a bottleneck as all the RX (receive)/TX (transmit) packets go through Domain0. In some embodiment, Domain0 can run on only one CPU and this limits the networking performance. SR-IOV technology may be used in the virtualized environment to address the above identified IO bottleneck problem. The SR-IOV specification is a standard for a type of PCI pass-through which natively shares a single device to multiple guests. SR-IOV reduces hypervisor involvement by specifying virtualization compatible memory spaces, interrupts and DMA streams. SR-IOV improves device performance for virtualized guests. SR-IOV allows a PCIe device to appear to be multiple separate physical PCIe devices. With SR-IOV technology, the virtualized environment can create virtual functions(devices) from a physical device and these virtual devices can be assigned to virtual machines. Once a virtual machine has a virtual device assigned, the virtual machine can receive and transmit packets from the virtual device without the intervention of hypervisor. The virtual machine may use a driver for the virtual device. In an example embodiment, Intel 82599 NICs are used in the device and are SR-IOV capable.

Referring now to FIG. 6B, a diagram of embodiments of several load balancing methods for load balancing WAN optimizers is depicted. These load balancing methods may be performed by an application delivery controller. In brief overview, the load balancing methods include 1) a least connection, 2) a static configuration, 3) least accumulated load using receive bandwidth with agent id persistence 4) least connection with agent id persistence and 5) source internet protocol (IP) hashing. Each of the load balancing methods are considered for factors, such as load balance effectiveness, compression history

performance, bandwidth management, simplicity, fault tolerance and persistence unit. There may be no one-size-fits-all solution for all WAN optimization (e.g., branch repeater) deployments. The trade-off may be the ease of deployment versus optimal WAN optimization.

5 In further details, the least connection load balancing method may be easy to configure but may be sub-optimal. The least connection method may provide strong load balancing effectiveness and may be the better or best choice for such effectiveness of the load balancing methods. However, the least connection method may provide poor compression history performance from the WAN optimizer. The least connection method may offer little
10 or uncertain bandwidth management capabilities. However, the least connection method is simple to configure and offers one of the simplest configuration of all the least connection methods. Least connection load balancing provides strong fault tolerant capabilities. With least connections, persistence can be obtained for each client (e.g., the persistence unit is the client/clients' connection).

15 The static configuration load balancing method may be the most challenging to configure and manage but could provide one of the best or optimal WAN optimization. The static connection method may provide good load balancing effectiveness that is better than other choices of the load balancing methods. A static configuration may provide the best compression performance from the WAN optimizer. The static configuration method may
20 neutral bandwidth management capabilities. Fault tolerant capabilities for this type of load balancing is poor with a single spare option available for replacement. With static configuration, persistence can be managed on a branch level, such as WAN optimizer for a group of clients at a branch office (e.g., the persistence unit is the branch (e.g., WAN optimization peers).

25 The least accumulated load based on bandwidth method may provide strong bandwidth management and strong fault tolerance with neutral simplicity of configuration. while providing one of the best or optimal WAN optimization. The least accumulated load based on bandwidth method may provide good load balancing effectiveness that is better than other choices of the load balancing methods. The least accumulated load based on bandwidth
30 method may use agent id persistence. In some embodiments, the agent id is a name or identifier referring to a specific WAN optimizer.. Each WAN optimizer may have an agent id for a peer WAN optimizer to communicate with for performing WAN optimization. Each WAN optimizers may refer to a peer WAN optimizer by the agent id. With agent id persistence, the persistence unit is branch based.

The least connection with agent id persistence method may provide strong bandwidth management with good/better fault tolerance and simplicity of configuration. The least connection with agent id persistence may provide poor load balancing effectiveness with good to better compression performance. In some embodiments, this load balancing method may not enable equal balancing of load if the peers (identified by the agent ids) are of different sizes. This embodiment of the least connection method may use agent id persistence. The agent id may identify a peer of a WAN optimizer. With agent id persistence, the persistence unit is branch based.

The source IP hashing based method may be one of the best in simplicity of configuration and fault tolerance support with good/better bandwidth management capabilities and compression history performance. The source IP hashing method may load up one WAN optimizer at a time and thus offers uncertainty in effectiveness as a load balancing method for WAN optimizers. The source IP hashing, in some embodiments, offers branch level persistence.

Referring now to FIG. 6C, an embodiment of a system that performs adaptive deployment of WAN optimizers is depicted. In brief overview, a Branch Repeater (BR) Manager 625 may determine a number of instances 634 of WAN optimizers and a size 632 of each of the WAN optimizers deployed on the appliance 200. A monitor 640 may monitor the performance and operations of the appliance 200 and store this data 625 into a data storage. The BR manager may be responsive to one or more rules 630, which may be configured or learned by the BR Manager. From the data 615 and responsive to the rule 630, the BR Manager may determine the number of instances and size of WAN optimizers and a load balancing (LB) method 636.

In further details, the BR Manager 625 may comprise any type and form of executable instructions executable on a device, such as appliance 200. The BR Manager 625 may comprise a virtual machine. The BR Manager 625 may be designed and constructed as part of a virtual appliance. The BR Manager 625 may comprise an application executing on a core of a multi-core device. The BR Manager 625 may be designed and constructed as part of a vServer 275. The BR Manager 625 may be designed and constructed as part of an application delivery controller, such as a virtualized application delivery controller 460. The BR Manager 625 may be designed and constructed as part of the management services 404 of the virtualized environment. The BR Manager 625 may be designed and constructed as part of a virtualized networking optimization engine 450. The BR Manager may execute on a

separate device in communication with the appliance. For example, the BR Manager may execute on a client, a server, or another appliance.

The monitor 640 may comprise any type and form of executable instructions to monitor performance and operation of any aspect of the appliance. The monitor may monitor network traffic via the appliance. The monitor may monitor the load of any application or module executing on the appliance. The monitor may monitor the load of any virtualized appliance executing the appliance 200. The monitor may monitor the load of any virtualized application delivery controller 460. The monitor may monitor the load of any virtualized networking optimization engine 450 (e.g., WAN optimizer). The monitor may monitor the bandwidth received, transferred via and/or transmitted by the appliance and/or any module executing on the appliance, including any virtualized application delivery controller and/or virtualized network optimization engine.

The monitor may monitor the type of network traffic traversing the appliance and/or any module executing on the appliance, including any virtualized application delivery controller and/or virtualized network optimization engine. The monitor may monitor the number of transport layer connections established or used by the appliance and/or any module executing on the appliance, including any virtualized application delivery controller and/or virtualized network optimization engine. The monitor may monitor the number of transactions, requests and/or responses handled or processed over a period of time by the appliance and/or any module executing on the appliance, including any virtualized application delivery controller and/or virtualized network optimization engine. The monitor may monitor the rate and/or change in rate of requests, responses or transactions handled or processed by the appliance and/or any module executing on the appliance, including any virtualized application delivery controller and/or virtualized network optimization engine.

The monitor may monitor the size or resource footprint of each instance of an application or module executing on the appliance, such as any virtualized application delivery controller 460 and/or virtualized network optimization engine 450. The monitor may determine the amount of resources used by each instance of an application executing on the appliance, such as the amount of memory, computing, network and storage resources. The monitor may determine the amount of resources used by each instance of an application over a period of time and the change of rate of such use. The monitor may monitor the size of each peer of WAN optimizers, such as a first WAN optimizer working in conjunction with a second WAN optimizer. The monitor may identify or determine any difference in size or the resource footprint between peer WAN optimizers. The monitor may identify or determine

any difference in size or resource footprint between peer WAN optimizers exceeds a predetermined threshold.

The monitor may monitor the size of the compression histories of the WAN optimizers. The monitor may monitor the fragmentation of the compression histories stored on the appliance. The monitor may monitor the compression ratios or efficiencies obtained by each of the WAN optimizers. The monitor may monitor CPU or core utilization by each of the WAN optimizers. The monitor may determine whether or not the size and fragmentation of compression histories are within corresponding predetermined thresholds. The monitor may determine whether or not compression ratios are within a predetermined threshold. The monitor may determine whether or not CPU/processor/core utilization is within a predetermined threshold.

The monitor may store to a data store any of the performance and operational characteristics 615, (including but not limited to the above information monitored by the monitor) identified or determined via monitoring. The monitor may store such data 615 when identified upon monitoring. The monitor may store such data 615 on a predetermined frequency. In some embodiments, the WAN optimizer or entity being monitored stores or logs data to the data store, such as resource usage, performance metrics and/or operational metrics. The BR manager may query the virtualized environment to identify performance and operational data and store such data into a data store. The BR manager may query or inspect each virtualized application delivery controller and/or virtualized network optimization engine to identify performance and operational data and store such data into a data store.

The BR manager 625 may comprise logic, operations or functions to use the performance and operational data 615 to determine how to deploy instances of a virtualized network optimization engine on an appliance. The BR may comprise an intelligence engine that analyze the data 615 to learn how instances of the virtualized network optimization engine (e.g. WAN optimizer) and the appliance 200 are performing. The BR Manager may analyze the data 615 to determine a number of instances 634 of WAN optimizers to execute on the appliance. The BR Manager may analyze the data 615 to determine a size 632 of each of the WAN optimizers to execute on the appliance. The BR Manager may analyze the data 615 to determine a load balancing method 636 for load balancing virtual WAN optimizers executing on the appliance. The BR Manager may determine any combination of instances, size and load balancing method to use based on the data. The BR Manager may determine the load to distribute across each of the plurality of WAN optimizers.

The BR Manager may determine instances and size of virtual WAN optimizers upon startup, boot up or initialization of the appliance. The BR Manager may determine instances and size of virtual WAN optimizers upon startup, boot up or initialization of one or more WAN optimizers. The BR Manager may determine instances and size of virtual WAN optimizers dynamically and real-time during operation of the appliance. The BR Manager may determine instances and size of virtual WAN optimizers on demand to meet load and bandwidth used by the appliance. The BR manager may determine instances and size of virtual WAN optimizers at predetermined frequencies during operation of the appliance. The BR Manager may determine instances and size of virtual WAN optimizers upon a predetermined event, such as exceeding a performance or operational threshold. The BR Manager may switch between different configurations of sizes of WAN optimizers. The BR Manager may switch between different configurations of instances of WAN optimizers. .. The BR Manager may switch between different configurations of instances and sizes of WAN optimizers.

The BR Manager 625 may operate responsive to one or more rules 630. The rules may be configured by administrators. The BR Manager may create rules automatically responsive to the data 615. The BR Manager may comprise a learning mode in which the BR Manager monitors and learns the performance and operation of the appliance. As a result of learning mode, the BR Manager may automatically configure one or more rules responsive to examination and/or analyze of the data 615 collected via learning mode. The BR Manager may continuously operate in a learning mode to adjust the configuration of rules or create new rules according to what is learned.

A rule may identify or specify a number of instances of a virtual WAN optimizer, a size of a virtualized WAN optimizer and/or a load balancing (LB) method. A rule may identify or specify a set of conditions upon which to use a number of instances of a virtual WAN optimizer, a size of a virtualized WAN optimizer and/or a LB method. The conditions may be any type and form of performance or operation condition, including but not limited to number of connections, load, bandwidth usage, speed of network 104, number of clients and number of servers. A rule may identify or specify criteria upon which to use a number of instances of a virtual WAN optimizer, a size of a virtualized WAN optimizer and/or a LB method. The criteria may be the amount of resources on the appliance. The criteria may be the amount of resources available to a virtualized WAN optimizer.

In some embodiments, based on resources (e.g., CPU, memory, disk), the BR Manager responsive to the data and/or rules may determine how many WAN optimization

instances should be used and of what size, and how the load should be distributed across the instances of the WAN optimizer. The BR Manager responsive to the data and/or rules may automatically configure the number of WAN optimizers, the size of each WAN optimizer and/or the LB method to use for load balancing the WAN optimizers.

5 Some example rules are as follows. In some embodiments, a rule specifies if a small number of peer WAN optimizers exist, a smaller number of large WAN optimizer instances should be provisioned on an appliance because compression histories will be less fragmented and compression ratios higher (better). In some embodiments, a rule specifies when peer WAN optimizers are of significantly different sizes, they should be distributed unevenly
10 across the WAN optimizer instances (for example, in some embodiments, by using WAN optimizers instances of different sizes).

 In some embodiments, the WAN optimizer includes the BR Manager or a portion thereof. In some embodiments, the WAN optimizer may interface to or communicate with the BR Manager to obtain, receive or be provided a size configuration or resource allocation
15 from the BR Manager. In some embodiments, the WAN optimizer may be designed and constructed to identify, detect or query the resources available via the appliance or the virtualized environment. In some embodiments, the WAN optimizer may be designed and constructed to query the data store to obtain historical information for configuring itself. Based on the information obtained by the WAN optimizer, the WAN optimizer may
20 automatically configure the size of itself via the amount of resources used by the WAN optimizer. For example, the WAN optimizer may allocate a certain or predetermined amount of memory and storage.

 Referring now to FIG. 6D, an embodiment of a method for automatically and dynamically adapting the number of instances of WAN optimizers, configuration of WAN
25 optimizers and load balancing method used for load balancing the WAN optimizers. In brief overview, at step 650, the method includes establishing a plurality of WAN optimizers on an appliance. At step 655, the BR manager monitors performance of the appliance and the plurality of WAN optimizers. The BR Manager may store performance and operational data
615 from monitoring. At step 660, the BR manager may adjust the number of instances of
30 WAN optimizers executing on the appliance and the configuration, such as size, of each of the WAN optimizers.

 In further details, at step 650, the appliance may establish a plurality of WAN optimizers executing on the appliance. The appliance may establish the plurality of WAN optimizers according to one or more rules. The appliance may establish a plurality of

virtualized WAN optimizers 450A-450N across a virtualized networking platform provided by the appliance. In some embodiments, a BR Manager queries the resource capacity and/or availability of the appliance and/or virtualized environment 400. Based on the resource capacity and/or availability, the BR Manager may establish a number of instances of WAN optimizers to execute on the appliance. Based on the resource capacity and/or availability, the BR Manager may establish a resource footprint or allocation to be used for all instances of the WAN optimizer. Based on the resource capacity and/or availability, the BR Manager may establish a size of each of the WAN optimizers to execute on the appliance, such as the amount of memory, compute and storage to be used by each WAN optimizer. Based on the resource capacity and/or availability, the BR Manager may establish an amount of bandwidth to be allocated or used by each of the WAN optimizers. Based on the resource capacity and/or availability, the BR Manager may establish a size of compression history to be allocated or used by each of the WAN optimizers. Based on the resource capacity and/or availability, the BR Manager may establish a load to be allocated or used by each of the WAN optimizers. Based on the resource capacity and/or availability and/or the number and size of WAN optimizers, the BR Manager may establish a load balancing method to load balancing the plurality of WAN optimizers.

In some embodiments, each of the WAN optimizers may auto-configure the size of itself. Each WAN optimizer may query the available and/or capacity of resources on the appliance or provided by the virtualized environment. Based on the resource capacity and/or availability, the WAN optimizer may establish a size for itself, such as the amount of memory, compute and storage for the WAN optimizer to allocate or use. Based on the resource capacity and/or availability, the WAN optimizer may establish an amount of bandwidth for itself to allocate or use. Based on the resource capacity and/or availability, the WAN optimizer may establish an amount of compression history for itself to allocate or use. Based on the resource capacity and/or availability, the WAN optimizer may establish load for itself to manage.

The BR Manager may set an initial or predetermined number and size of WAN optimizers to execute on the appliance. Upon establishing, the BR Manager may execute in a learning mode to examine to network traffic, load, bandwidth and other operational and performance characteristics. For example, as described below in connection with steps 655 and 660, the BR Manager may monitor performance of the appliance and WAN optimizers and dynamically and automatically adjust the number of instances of WAN optimizers and the size of each WAN optimizer. In some embodiments, the BR manager may use a baseline

configuration established via monitoring and adjustment steps 655 and 660. Such a baseline configuration may identify the number of instances of WAN optimizers, the size of each instance of a WAN optimizer and Load balancing method to use to load balance the WAN optimizers. The load balancing method may include any of the load balancing methods described in connection with FIG. 6B.

At step 655, a monitor monitors the operations and/or performance of the appliance, such as any entities executing on the appliance, including but not limited to WAN optimizers. The monitor may monitor the performance and operation of each of the WAN optimizers. The monitor may monitor the performance and operation of each virtualized application delivery controller. The monitor may monitor the performance and operation of the virtualized environment. The monitor may monitor the performance and operation of each core of a multi-core appliance. The monitor may monitor the network traffic and bandwidth used by the appliance and any virtual appliance executing thereon, such as the WAN optimizers. The monitor may monitor the type and volume of network traffic traversing the appliance and any virtualized appliance executing thereon, such as the WAN optimizers. The monitor may monitor the load of each WAN optimizer, including but not limited to resource utilization, number of connections, compression history allocation, and bandwidth usage. The monitor may monitor the compression ratio or effectiveness of each WAN optimizer. The monitor may monitor the amount of traffic received, transmitted or processed by each WAN optimizer.

The monitor may store any monitored information to a data storage. The monitor may store the data 615 on demand, real-time or as monitored. The monitor may store the data 615 on a predetermined frequency or for predetermined time periods. The monitor may query the virtualized environment, management services 404 or operating system of the appliance for performance and operational information and store the information to a data store 615.

At step 660, responsive to the data 615 and/or one or more rules, the BR Manager may automatically and/or dynamically adjust the number of instances of WAN optimizers, the size of each WAN optimizer and/or the load balancing method for load balancing the WAN optimizers. The BR Manager may analyze the data and automatically create or specify one or more rules. The BR Manager may analyze the data and automatically identify or specify a number of instances 634. The BR Manager may analyze the data and automatically identify or specify a size of instances 636. The BR Manager may analyze the data and automatically identify or specify a load balancing method 636. Responsive to the data 615

and/or one or more rules, the BR Manager may create or establish a baseline configuration for the WAN optimizers on the appliance.

Responsive to the data 615 and/or one or more rules, the BR Manager may increase a number of instances of WAN optimizers executing on the appliance. Responsive to the data 615 and/or one or more rules, the BR Manager may decrease a number of instances of WAN optimizers executing on the appliance. Responsive to the data 615 and/or one or more rules, the BR Manager may change the size of a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may increase the size of a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may decrease the size of a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may increase the amount of a resource allocated to or used by a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may decrease the amount of a resource allocated to or used by a WAN optimizer.

Responsive to the data 615 and/or one or more rules, the BR Manager may decrease the amount of bandwidth used by a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may increase the amount of bandwidth used by a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may decrease the load used by a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may increase the load used by a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may increase the number of connections or clients managed by a WAN optimizer. Responsive to the data 615 and/or one or more rules, the BR Manager may decrease the number of connections or clients managed by a WAN optimizer.

Responsive to the data 615 and/or one or more rules, the BR Manager may change the load balancing method used to load balance the WAN optimizers. Responsive to the data 615 and/or one or more rules, the BR Manager may enumerate an ordered list of load balancing methods to switch between or use based on conditions, criteria or load meeting predetermined thresholds.

In some embodiments, the BR Manager establishes a baseline configuration or versioned configuration to use at step 650. For example, the BR Manager may establish a baseline configuration to use when the appliance reboots, is reset or initialized. As such, step 650 may use an output of step 660 and the process may continually adapt the configuration via steps 655 and 660.

In some embodiments, the BR manager dynamically and automatically adapts the number of instances of WAN optimizers, size of each WAN optimizer and/or the load balancing method continuously over time. For example, the BR manager may change the number of instances of WAN optimizers, size of each WAN optimizer and/or the load balancing method upon one instance of step 655 and step 660. During a next execution of steps 655 and step 660, the BR manager may further change the number of instances of WAN optimizers, size of each WAN optimizer and/or the load balancing method. The BR Manager may continually monitor the performance of the appliance and WAN optimizers and based on conditions, rules or thresholds, automatically and adjust the deployment of the WAN optimizers.

Although embodiments of the systems described herein an embodiments of the method of FIG. 6D are generally described in connection with automatically and dynamically adapting the number of instances of WAN optimizers, the systems and methods described herein may be used to automatically and dynamically adapt the number of instances of any application, such as any network optimizers. Although embodiments of the systems described herein an embodiments of the method of FIG. 6D are generally described in connection with automatically and dynamically adapting the number of instances of WAN optimizers, the systems and methods described herein may be used to automatically and dynamically adapt the number of instances of any type and form of virtual appliance, including without limitation a virtualized application delivery controllers.

What is claimed:

1. A method for managing a plurality of instances of a Wide Area Network (WAN) optimizer executing on an intermediary device, the method comprising:
 - 5 (a) establishing, on a device intermediary to a plurality of clients and a plurality of servers, a plurality of instances of a Wide Area Network (WAN) optimizer to accelerate WAN communications between the plurality of clients and the plurality of servers;
 - (b) monitoring, by the device, network traffic traversing the device for each of
10 the plurality of instances of the WAN optimizer; and
 - (c) selecting, by a manager executing on the device responsive to the monitoring, a change of a load balancing scheme to load balance the plurality of instances of the WAN optimizer.
- 15 2. The method of claim 1, wherein step (a) further comprises automatically establishing, by the device, a configuration of a size of each of the plurality of instances of the WAN optimizer based on data stored from monitoring of previous execution of the plurality of instances of the WAN optimizer.
- 20 3. The method of claim 1, wherein step (a) further comprises executing, by the device, each of the plurality of instances of the WAN optimizer as a virtual machine in a virtualized environment.
4. The method of claim 1, wherein step (b) further comprises monitoring, by the device, compression history allocation, compression fragmentation and compression ratios of each of the plurality of instances of the WAN optimizer.
- 25 5. The method of claim 1, wherein step (b) further comprises monitoring, by the device, or more of the following of each of the plurality of instances of the WAN optimizer: resource utilization, number of connections, number of claims and bandwidth usage.
6. The method of claim 1, wherein step (c) further comprises determining, by the device, that a metric computed from monitoring network traffic has exceeded a threshold and
30 responsive to the determination, automatically selecting by the device a second load balancing scheme to load balance the plurality of instances of the WAN optimizer.
7. The method of claim 1, wherein step (c) further comprises automatically switching, by the device, from the load balancing scheme to the selected load balancing scheme while executing the plurality of instances of the WAN optimizer.

8. The method of claim 1, further comprising automatically changing, by the device responsive to the monitoring, the number of instances of the WAN optimizer executing on the device.
9. The method of claim 1, further comprising automatically adjusting, by the device responsive to the monitoring, a size of resource usage used by one or more of the plurality of instances of the WAN optimizer.
10. The method of claim 1, further comprises applying, by the device one or more rules to data collected from monitoring, to determine to change one or more of the following: a number of instances of the WAN optimizer, a size of one or more WAN optimizers and the load balancing scheme.
11. A system for managing a plurality of instances of a Wide Area Network (WAN) optimizer executing on an intermediary device, the system comprising:
- a device intermediary to a plurality of clients and a plurality of servers;
 - a plurality of instances of a Wide Area Network (WAN) optimizer executing on the device to accelerate WAN communications between the plurality of clients and the plurality of servers;
 - a monitor that monitors network traffic traversing the device for each of the plurality of instances of the WAN optimizer; and
 - a manager executing on the device that, responsive to the monitor, selects a change of a load balancing scheme to load balance the plurality of instances of the WAN optimizer.
12. The system of claim 11, wherein the manager automatically establishes a configuration of a size of each of the plurality of instances of the WAN optimizer based on data stored from monitoring of previous execution of the plurality of instances of the WAN optimizer.
13. The system of claim 11, wherein each of the plurality of instances of the WAN optimizer execute as a virtual machine in a virtualized environment.
14. The system of claim 11, wherein the monitor monitors compression history allocation, compression fragmentation and compression ratios of each of the plurality of instances of the WAN optimizer.
15. The system of claim 11, wherein the monitor monitors one or more of the following of each of the plurality of instances of the WAN optimizer: resource utilization, number of connections, number of claims and bandwidth usage.

16. The system of claim 11, wherein the manager determines that a metric computed from monitoring network traffic has exceeded a threshold and responsive to the determination, automatically selects a second load balancing scheme to load balance the plurality of instances of the WAN optimizer.
- 5 17. The system of claim 11, wherein the manager automatically switches from a current load balancing scheme to the selected load balancing scheme while executing the plurality of instances of the WAN optimizer.
18. The system of claim 11, wherein the manager, responsive to the monitor, automatically changes the number of instances of the WAN optimizer executing on the device.
- 10 19. The system of claim 11, wherein the manager automatically adjusts, responsive to the monitor, a size of resource usage used by one or more of the plurality of instances of the WAN optimizer.
20. The system of claim 11, wherein the manager applies one or more rules to data collected from monitoring, to determine to change one or more of the following: a number of instances of the WAN optimizer, a size of one or more WAN optimizers and the load balancing scheme.
- 15




FIG. 1A

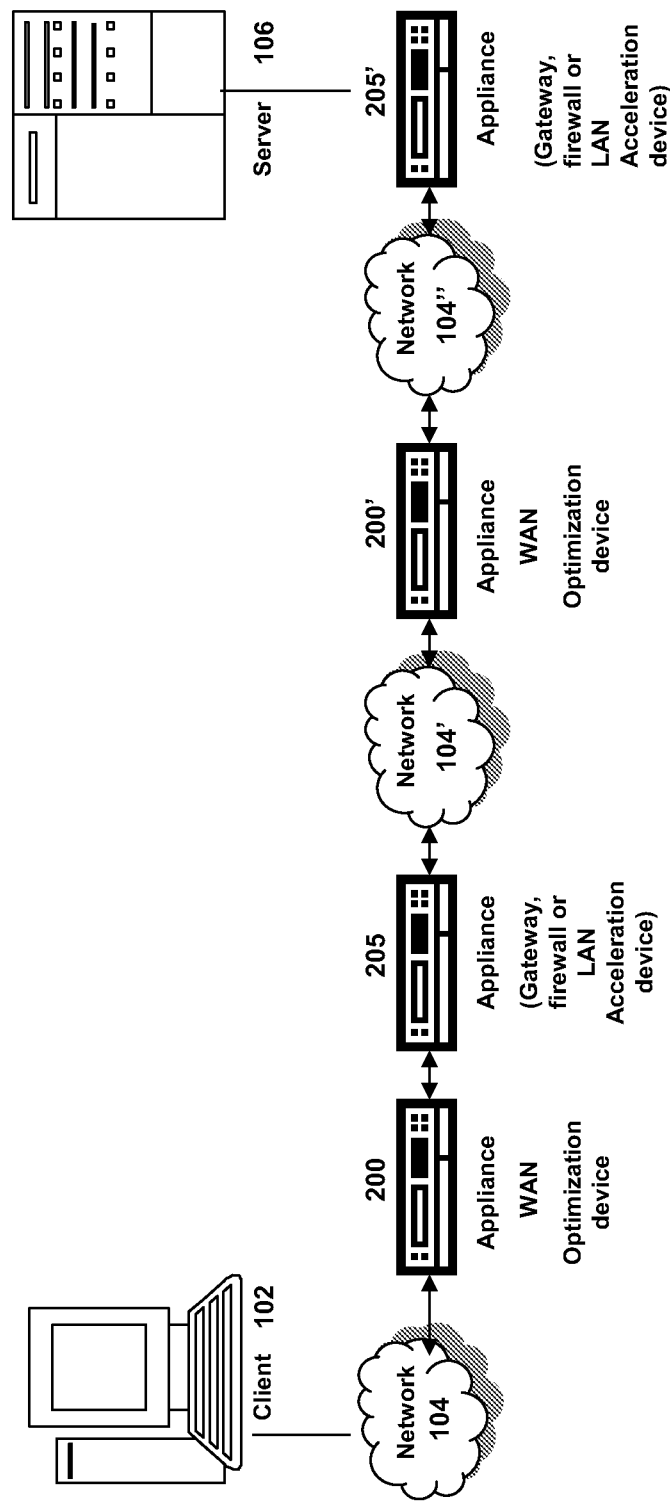


FIG. 1B

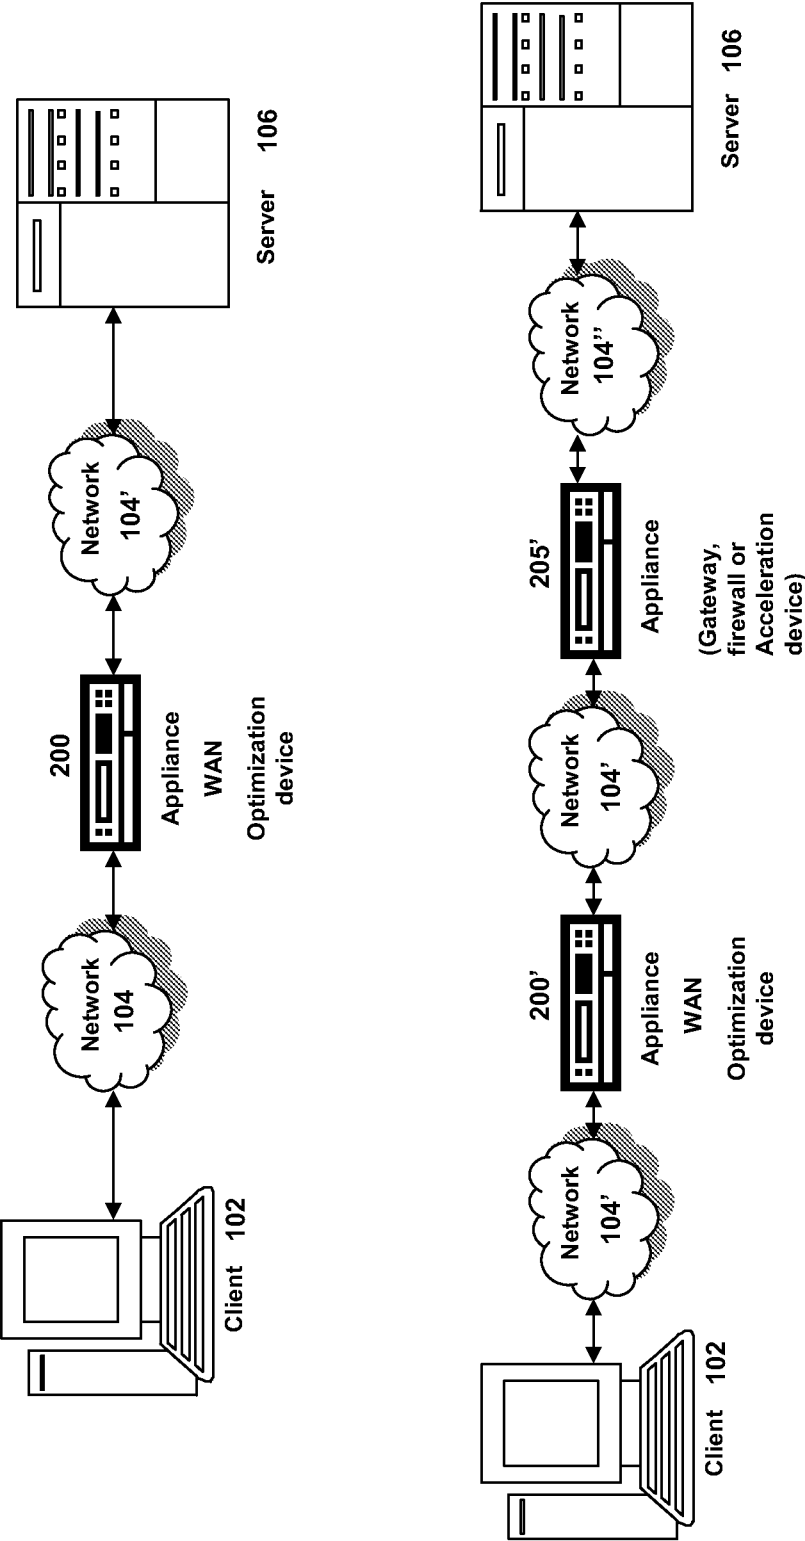




FIG. 1C

*Fig. 1D*

*Fig. 1E*




FIG. 2A




FIG. 2B

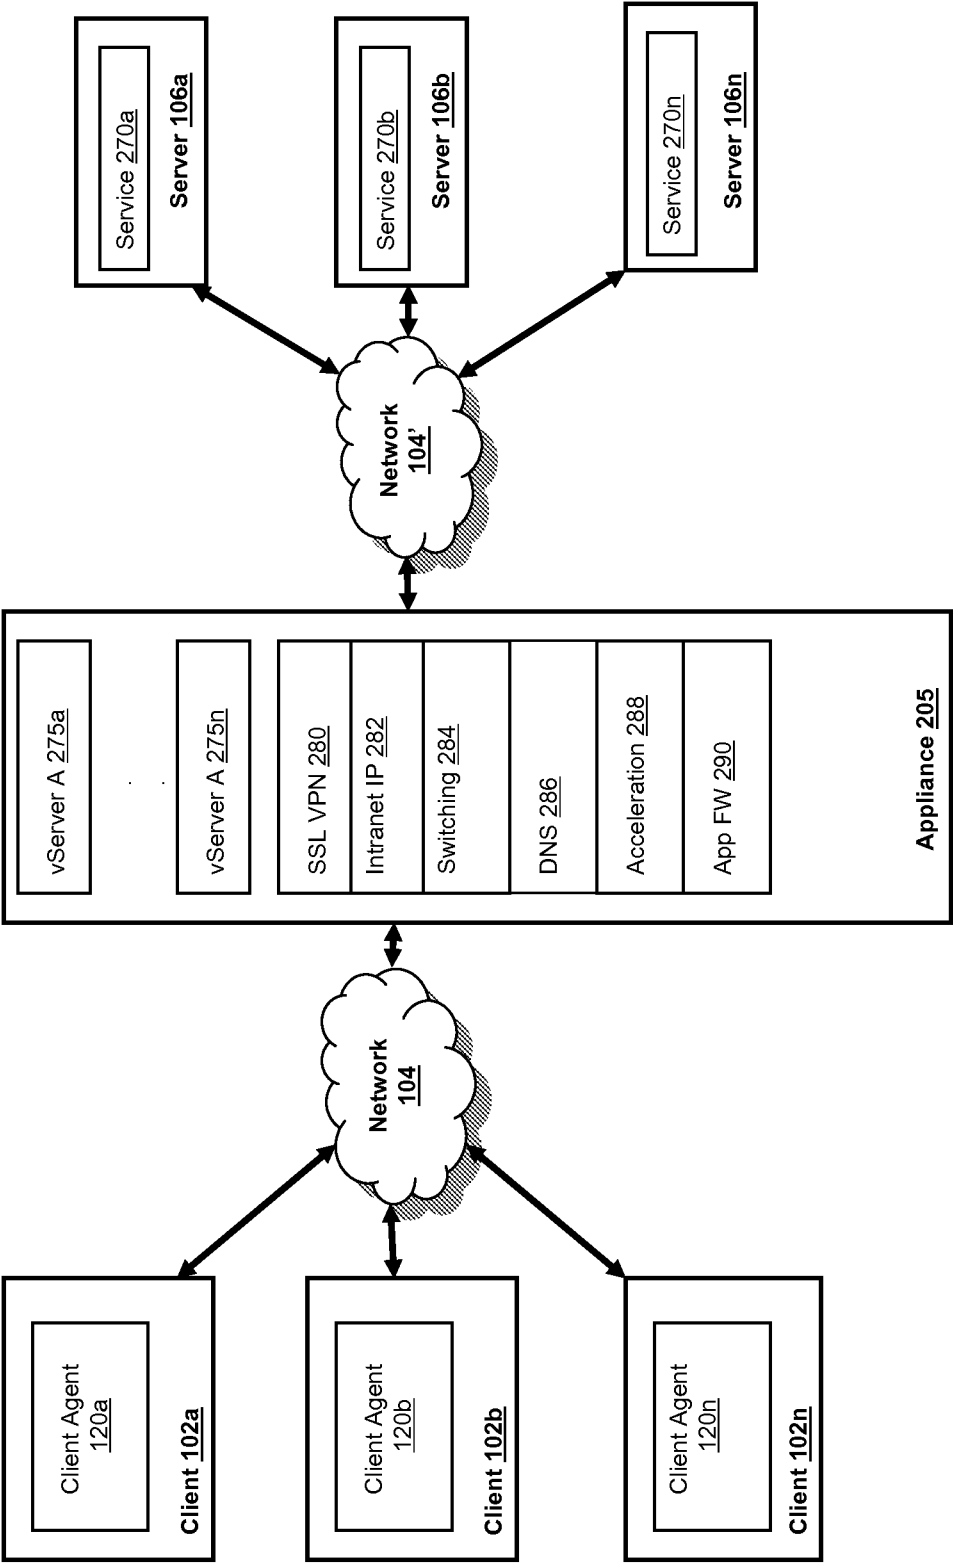
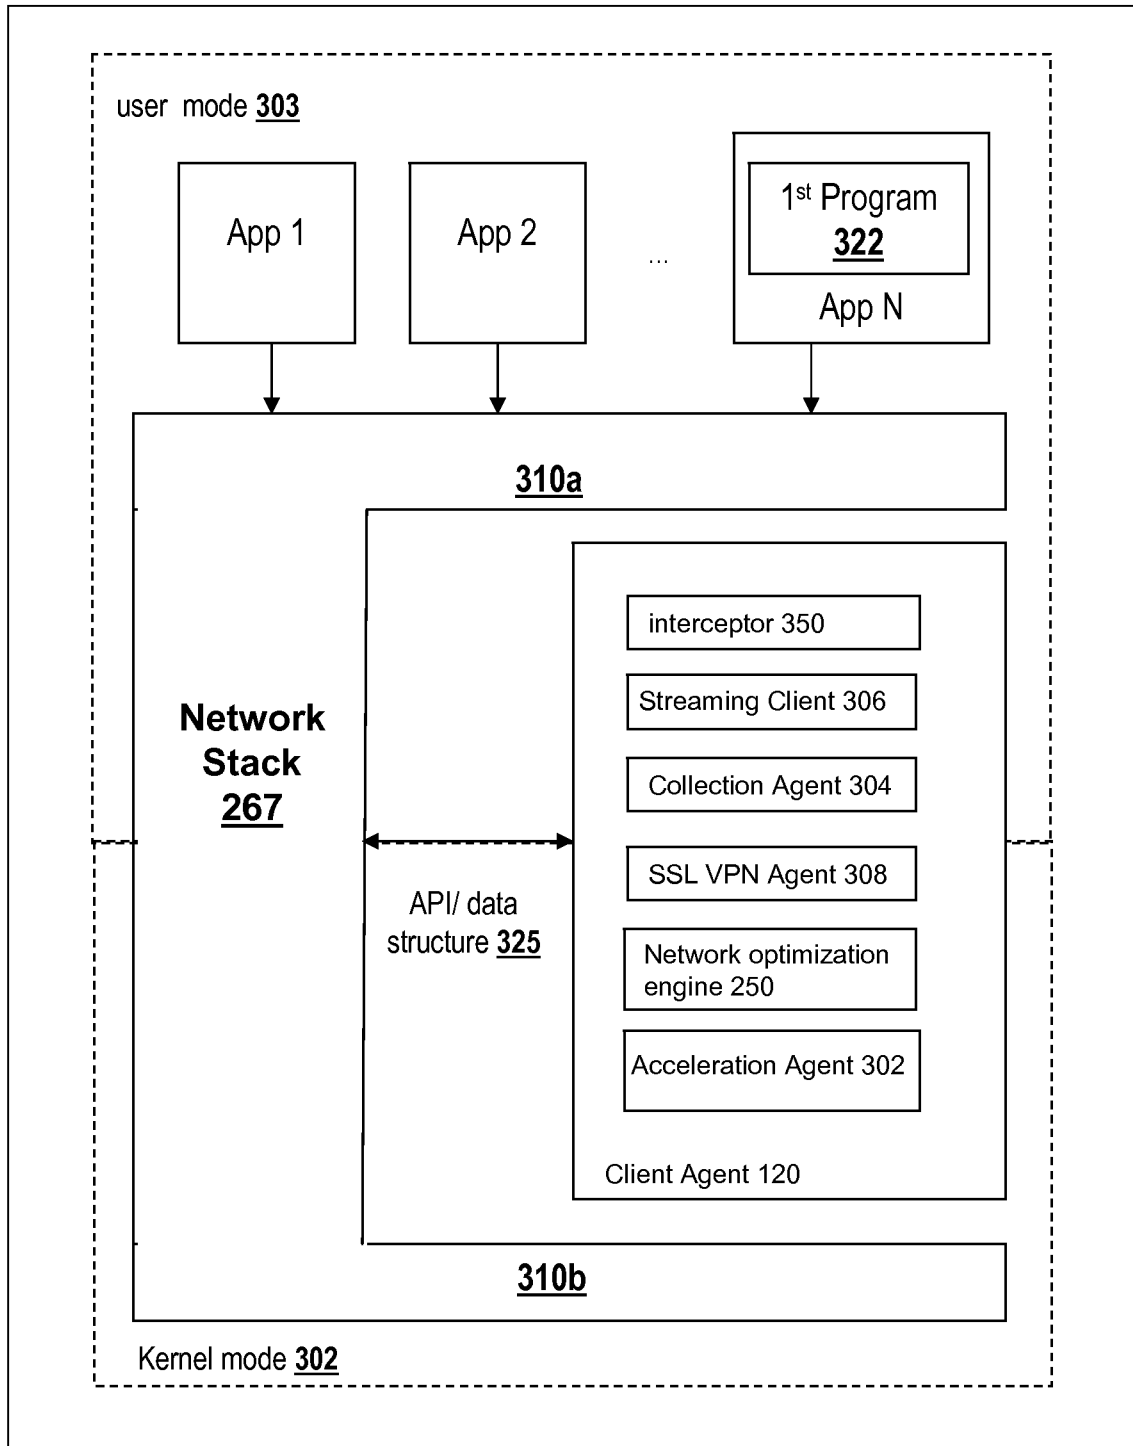
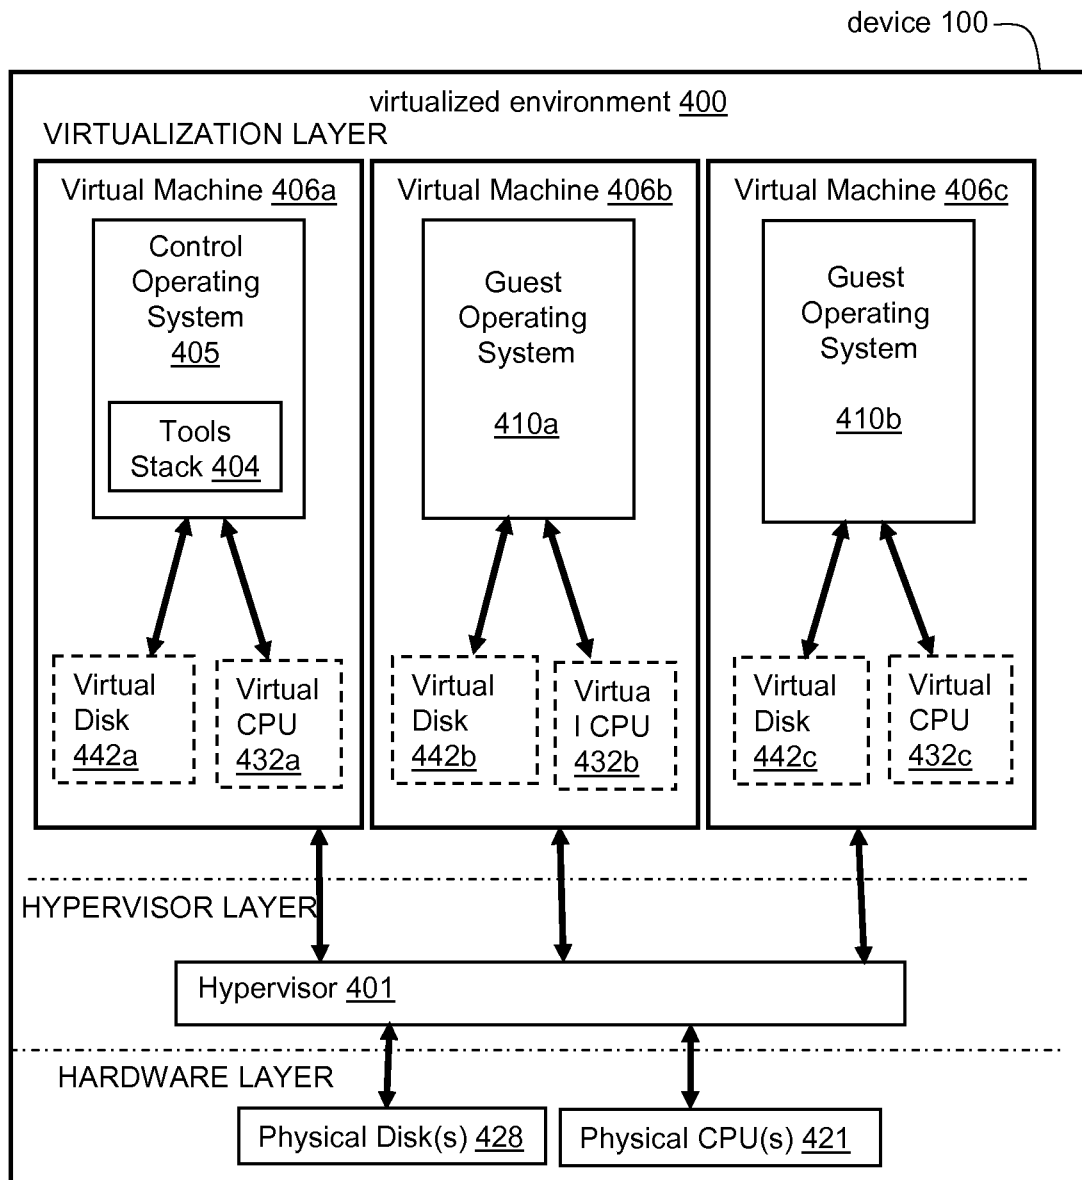


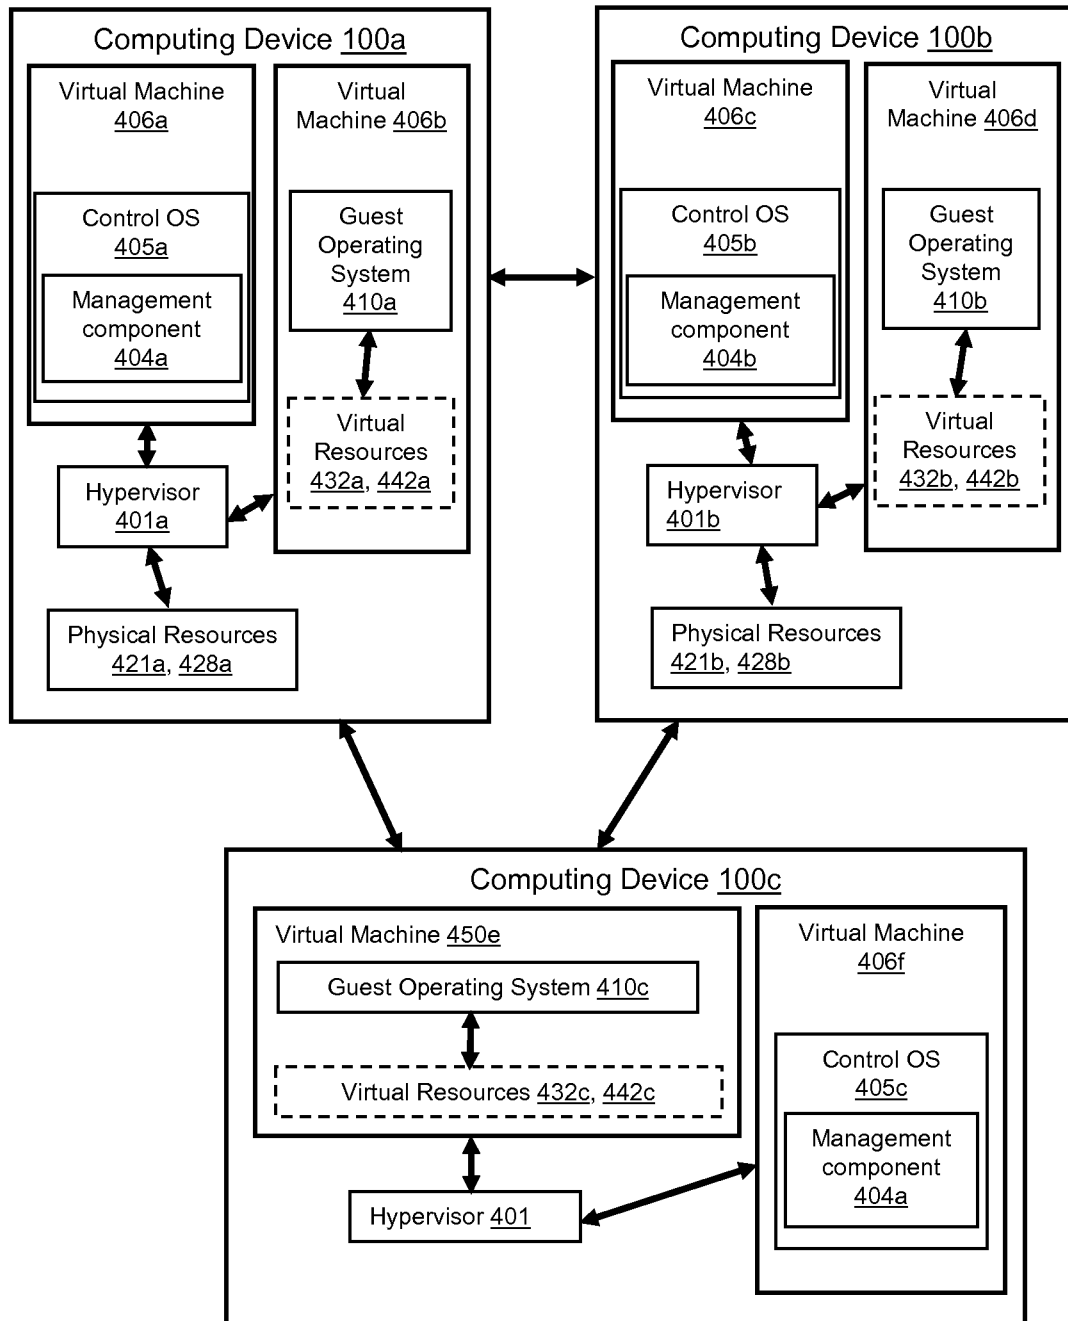
FIG. 2C

Client 102

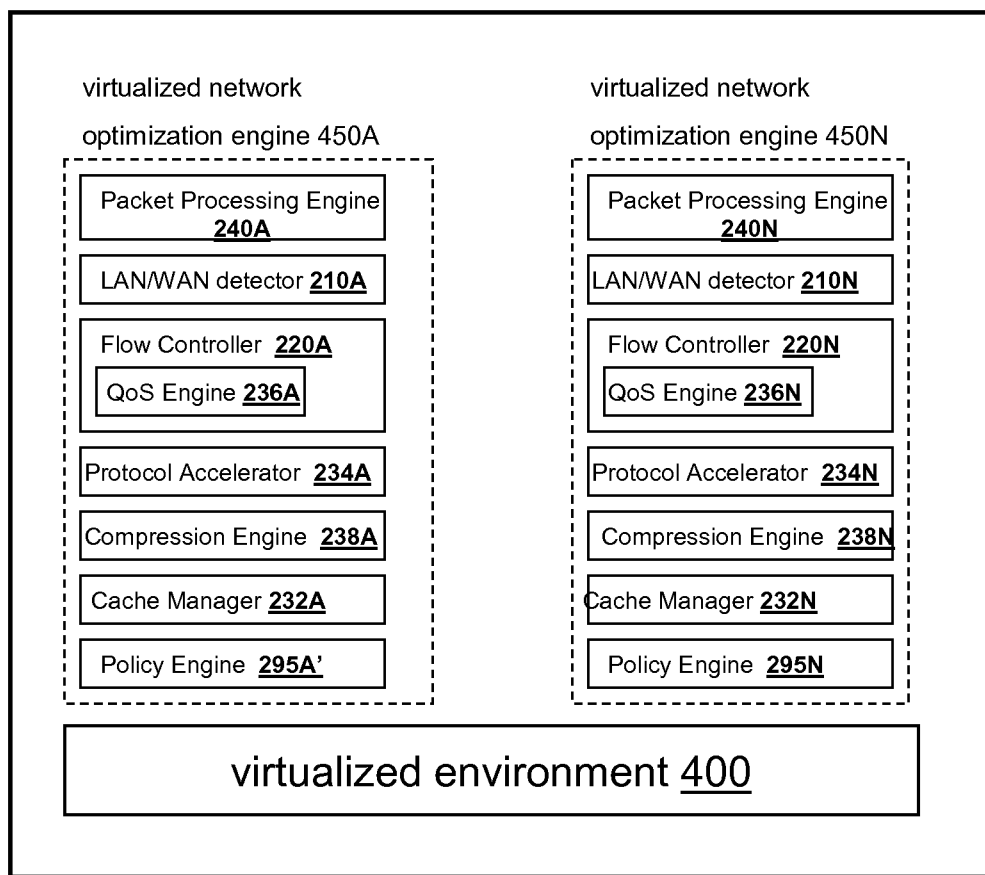
100

Fig. 3

**FIG. 4A**

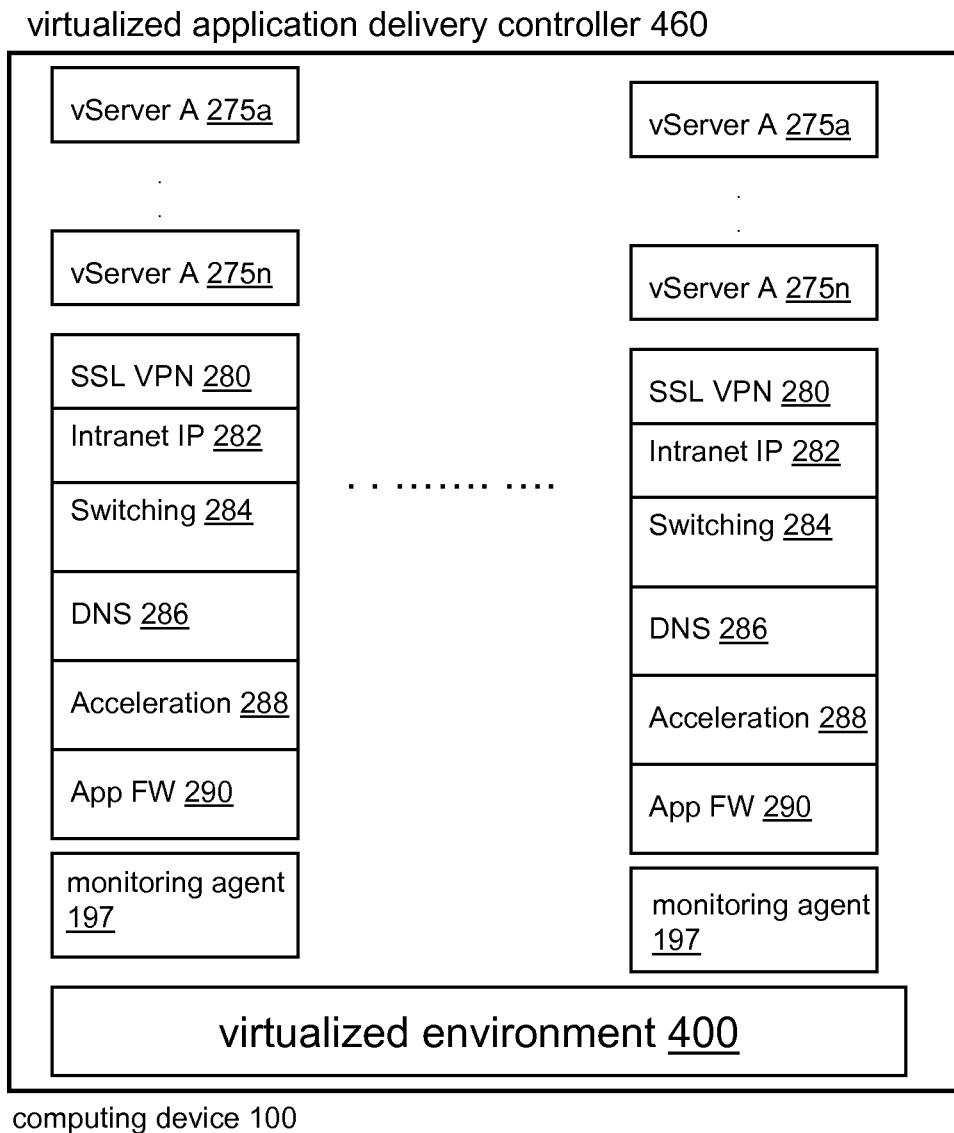
**FIG. 4B**

Virtualized appliance 450



device 100

FIG. 4C

**FIG. 4D**


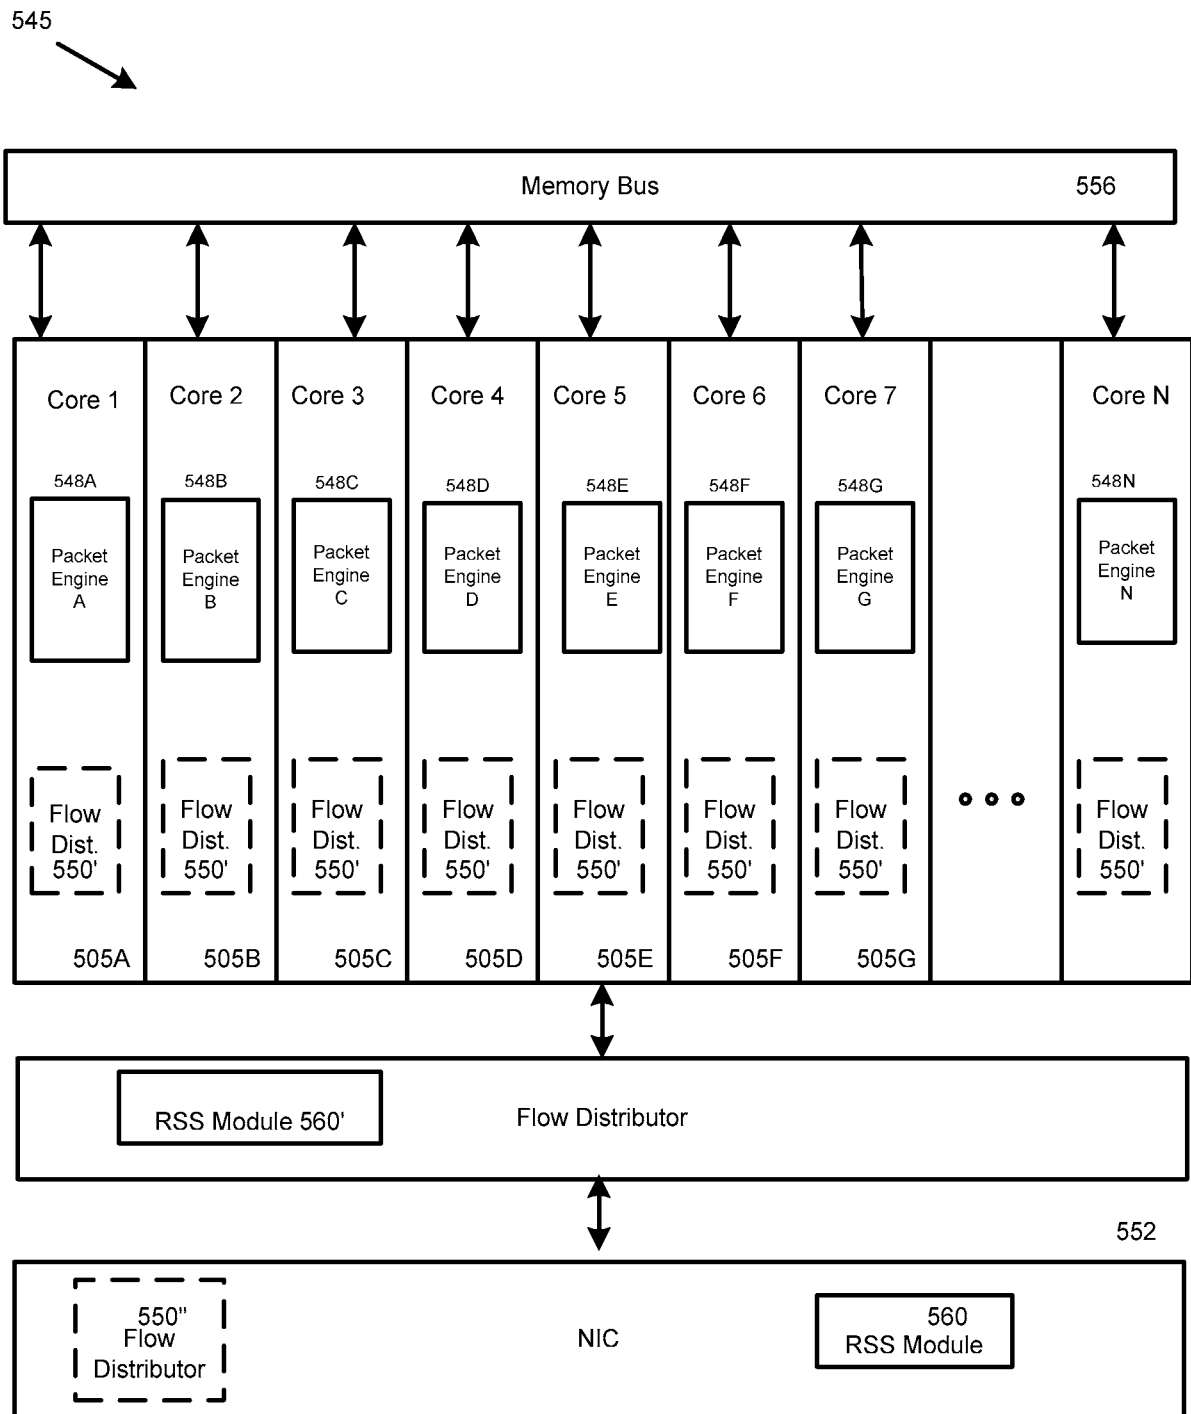


FIG. 5A

**FIG. 5B**

575
↘




FIG. 5C





FIG. 6A

	Load Balancing Method	Balance Effectiveness	Compression History	Bandwidth Management	Sim plici ty	Fault Tolerant	Persistent Unit
1	Least connection (sub-optimal, easiest to configure)	↑↑	↓↓	?	↑↑	↑↑	Client
2	Static configuration (optimal WAN optimization but difficult to manage)	↑	↑↑	→	↓↓	↓↓ Single spare option is available	Branch
3	Least accumulated load using receive bandwidth + AgentID persistence (better WAN optimization characteristics)	↑	↑↑	↑↑	→	↑↑	Branch
4	Least connection + AgentID persistence (requires branches to be same size,)	↓	↑	↑↑	↑	↑	Branch
5	Source IP hash to load up one BR at a time	?	↑	↑	↑↑	↑↑	Branch

Legend:

↑↑ Best
 ↑ Better
 → Neutral
 ↓ Poor
 ↓↓ Worst
 ? Don't know

FIG. 6B



Appliance 200

FIG. 6C

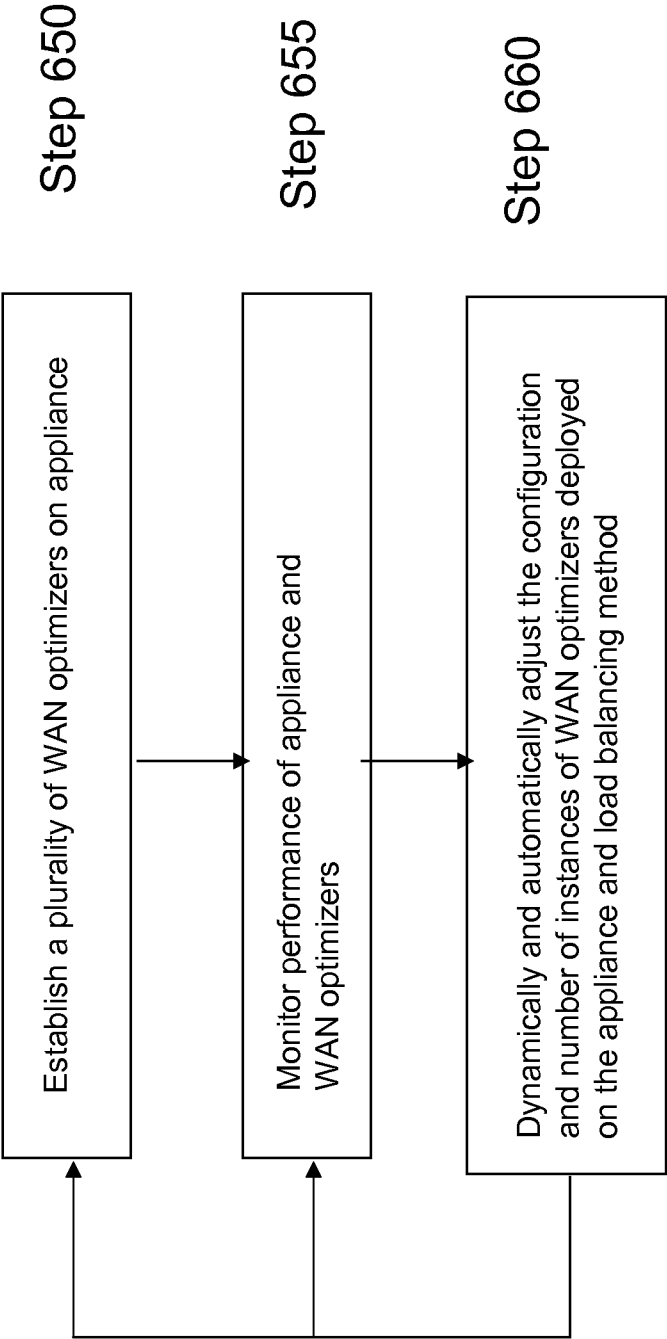


FIG. 6D

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2012/059586

A. CLASSIFICATION OF SUBJECT MATTER INV. H04L29/08 G06F9/455 G06F9/50 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 G06F		
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) EPO-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.
X	US 2001/049741 A1 (SKENE BRYAN D [US] ET AL) 6 December 2001 (2001-12-06) paragraphs [0024] - [0032] paragraphs [0070] - [0104] figures 3B, 7-11 -----	1-20
X	US 6 728 748 B1 (MANGIPUDI KRISHNA [US] ET AL) 27 April 2004 (2004-04-27) column 4, line 35 - column 5, line 50 column 12, line 62 - column 13, line 44 figure 7 -----	1-20
A	US 2010/131638 A1 (KONDAMURU RAVI [US]) 27 May 2010 (2010-05-27) paragraphs [0003] - [0009] paragraphs [0136] - [0228] paragraphs [0256] - [0258] -----	1-20
<div style="display: flex; justify-content: space-between;"> <div> <input type="checkbox"/> Further documents are listed in the continuation of Box C. </div> <div> <input checked="" type="checkbox"/> See patent family annex. </div> </div>		
<div style="display: flex;"> <div style="flex: 1;"> <p>* Special categories of cited documents :</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> </div> <div style="flex: 1;"> <p>"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</p> <p>"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</p> <p>"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</p> <p>"&" document member of the same patent family</p> </div> </div>		
Date of the actual completion of the international search <div style="text-align: center; font-size: 1.2em;">12 December 2012</div>		Date of mailing of the international search report <div style="text-align: center; font-size: 1.2em;">20/12/2012</div>
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 <div style="text-align: center; font-size: 1.2em;">Veloso González, J</div>

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2012/059586

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2001049741 A1	06-12-2001	NONE	
US 6728748 B1	27-04-2004	US 6728748 B1	27-04-2004
		US 2004162901 A1	19-08-2004
US 2010131638 A1	27-05-2010	US 2010131638 A1	27-05-2010
		WO 2010068458 A2	17-06-2010

摘要

本解決方案所涉及到的系統和方法以一種動態且自動的方式對系統工作負荷的改變做出回應，從而提供一種更加優化的解決方案。一種啟動緩慢的系統，通過以一種被動的方式對交通進行檢查或者通過執行次最優的加速，可以隨著時間的推移而學習得知一個設備當前正在為多少個對等 WAN 優化程序提供服務，每個對等 WAN 優化程序會發出多少的交通量，以及當前所看到的交通是何種類型。通過這種學習方法獲得的知識可以用來為一台設備的配置提供一種更佳的或者改進的基準線。在一些實施例中，根據資源的不同(例如 CPU，存儲器，磁盤)，來自該知識的系統可能決定了需要使用多少個 WAN 來優化實例以及這些實例的大小是多少，同時可能決定了負載是如何分佈到 WAN 優化程序的實例上。