



US007814182B2

(12) **United States Patent**
Cadigan, Jr. et al.

(10) **Patent No.:** **US 7,814,182 B2**
(45) **Date of Patent:** **Oct. 12, 2010**

(54) **ETHERNET VIRTUALIZATION USING
AUTOMATIC SELF-CONFIGURATION OF
LOGIC**

(75) Inventors: **Michael J. Cadigan, Jr.**, Brewster, NY
(US); **Howard M. Haynie**, Wappingers
Falls, NY (US); **Jeffrey M. Turner**,
Poughkeepsie, NY (US)

(73) Assignee: **International Business Machines
Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 321 days.

(21) Appl. No.: **12/052,303**

(22) Filed: **Mar. 20, 2008**

(65) **Prior Publication Data**

US 2009/0240788 A1 Sep. 24, 2009

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** 709/220; 709/248

(58) **Field of Classification Search** 709/220,
709/238

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,068,189 B2 * 6/2006 Brescia 340/995.1

7,117,241 B2 *	10/2006	Bloch et al.	709/201
7,425,962 B2 *	9/2008	Alcorn et al.	345/556
7,486,610 B1 *	2/2009	Bhupalam et al.	370/216
2001/0048661 A1 *	12/2001	Clear et al.	370/218
2004/0002348 A1 *	1/2004	Fraccaroli	455/456.3
2004/0117613 A1 *	6/2004	Edmondson	713/150
2005/0027778 A1 *	2/2005	Dimitrelis et al.	709/200
2006/0153156 A1 *	7/2006	Wentink et al.	370/338
2006/0155833 A1 *	7/2006	Matsuda et al.	709/220
2006/0174031 A1 *	8/2006	Yamakoshi et al.	709/237
2006/0268744 A1 *	11/2006	Sakai et al.	370/254
2007/0022211 A1 *	1/2007	Shimizu et al.	709/238
2007/0067361 A1 *	3/2007	Bailey et al.	707/204
2007/0253437 A1 *	11/2007	Radhakrishnan et al.	370/401
2008/0123647 A1 *	5/2008	Chen et al.	370/390
2008/0133721 A1 *	6/2008	Danz et al.	709/222
2009/0216852 A1 *	8/2009	Filippi	709/208

* cited by examiner

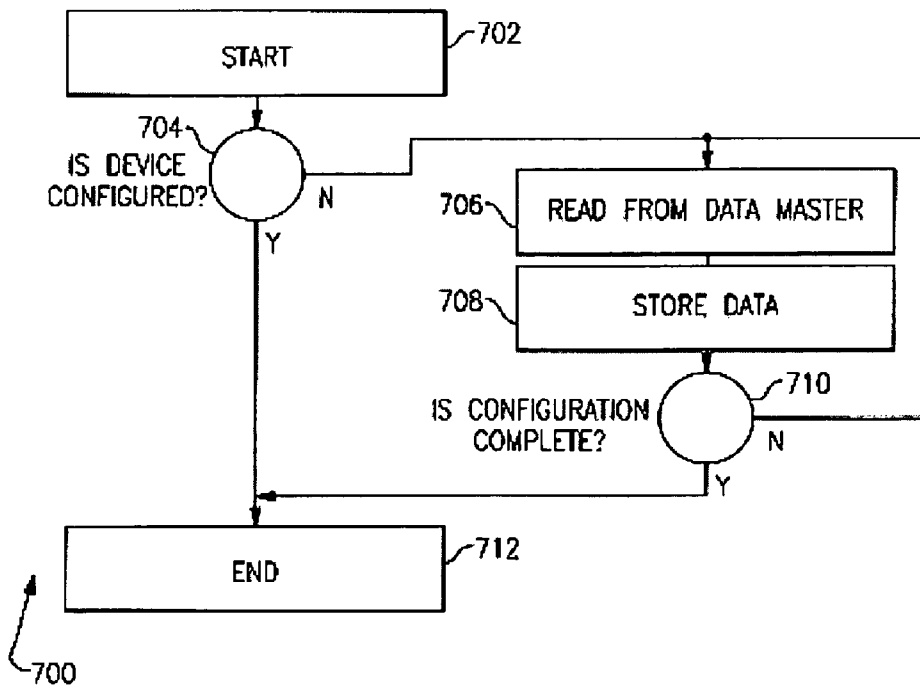
Primary Examiner—Bunjoo Jaroenchonwanit

(74) *Attorney, Agent, or Firm*—Steven Chiu

(57) **ABSTRACT**

A method of Ethernet virtualization using automatic self-configuration of logic of a data router. The method comprising maintaining control parameters at a master device, accessing, by a slave device, the control parameters at the master devices, and configuring the slave device based on the accessed control parameters.

1 Claim, 7 Drawing Sheets



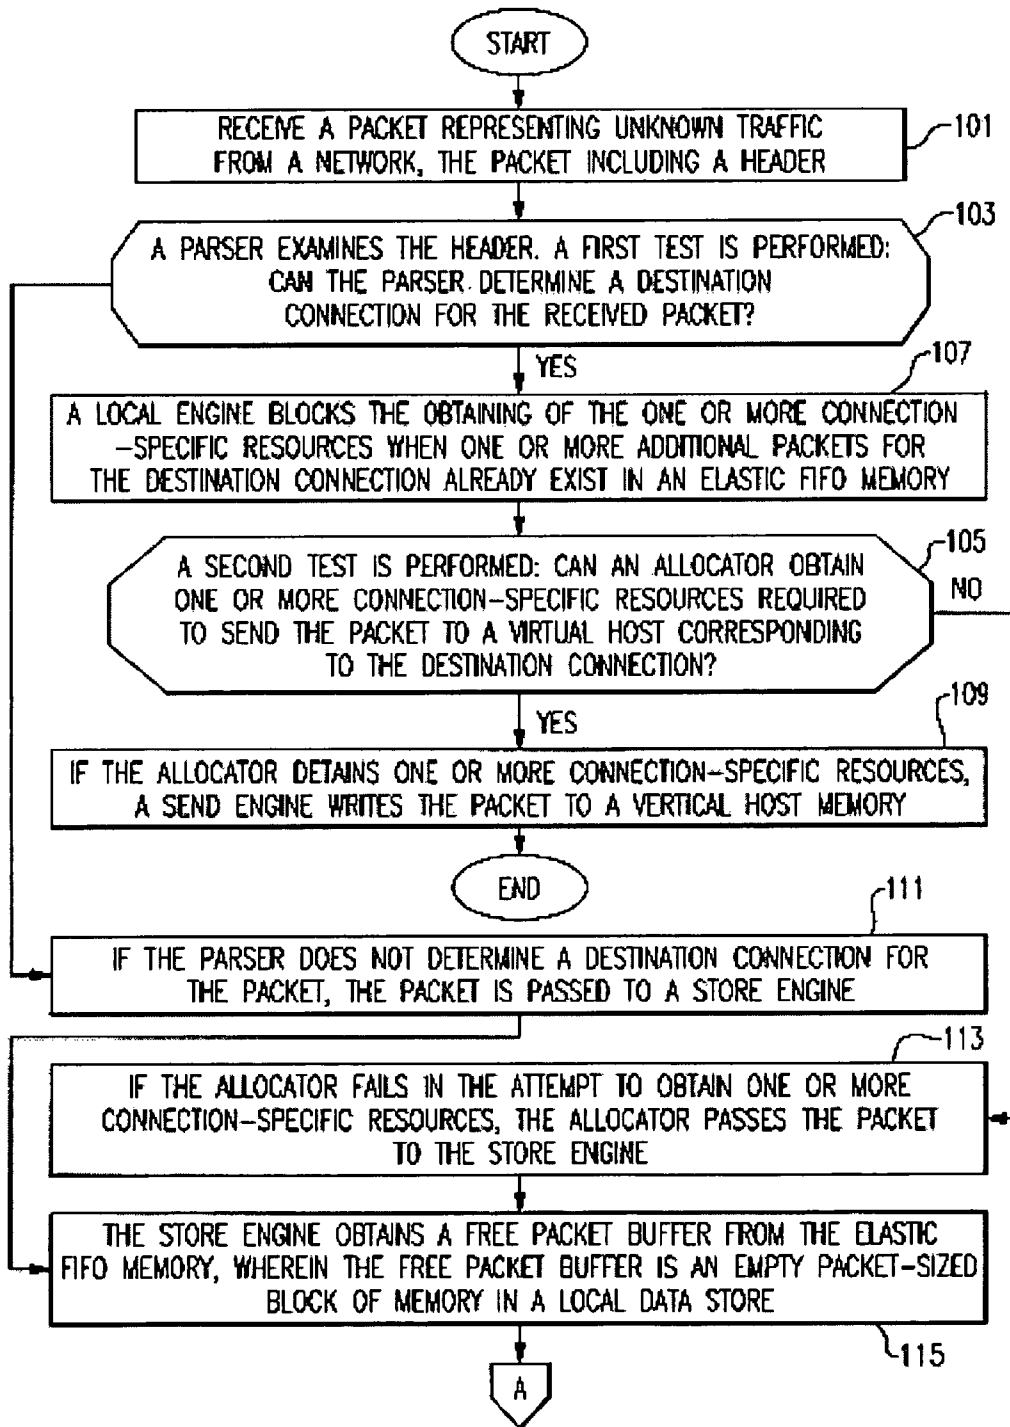


FIG. 1A

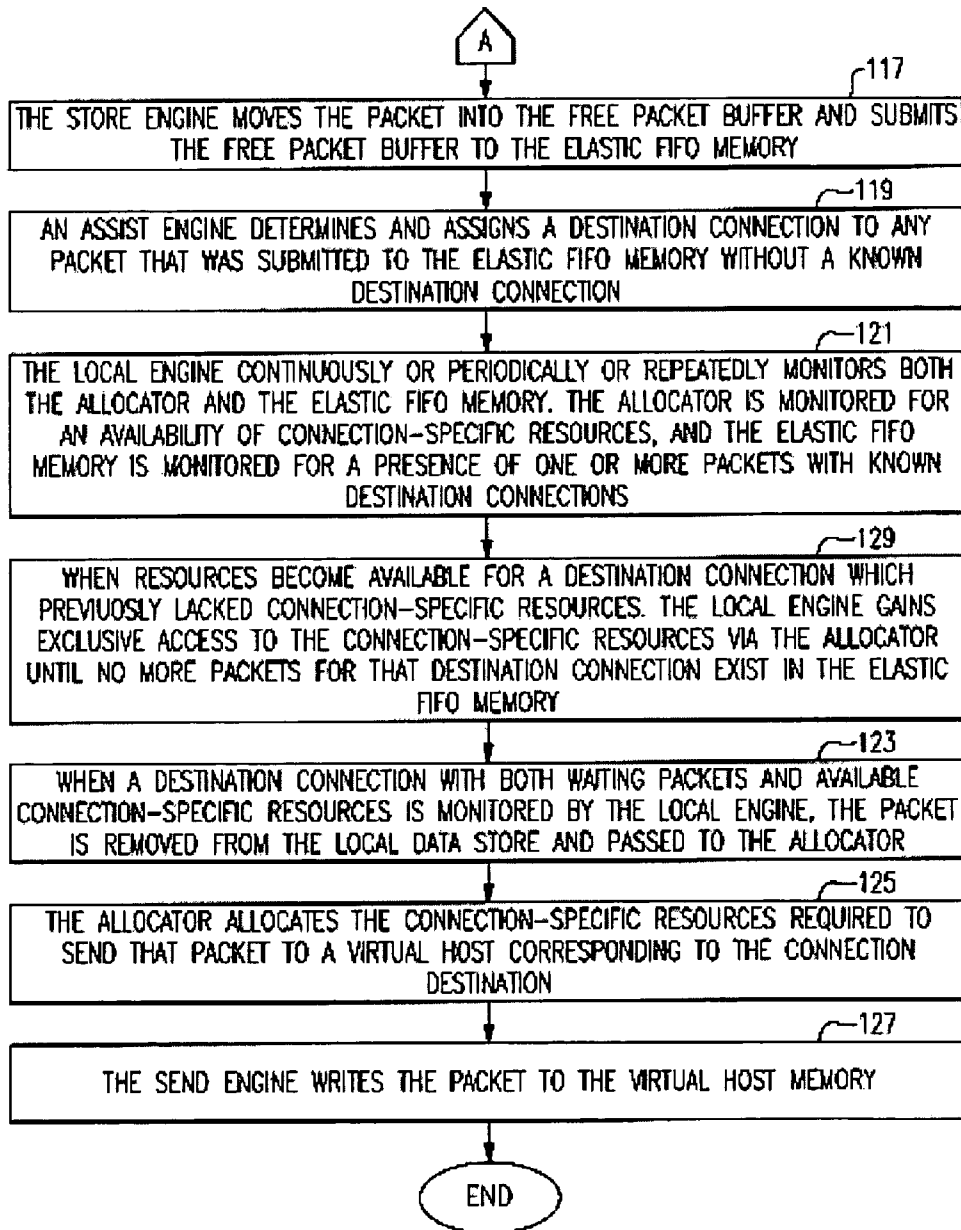


FIG.1B

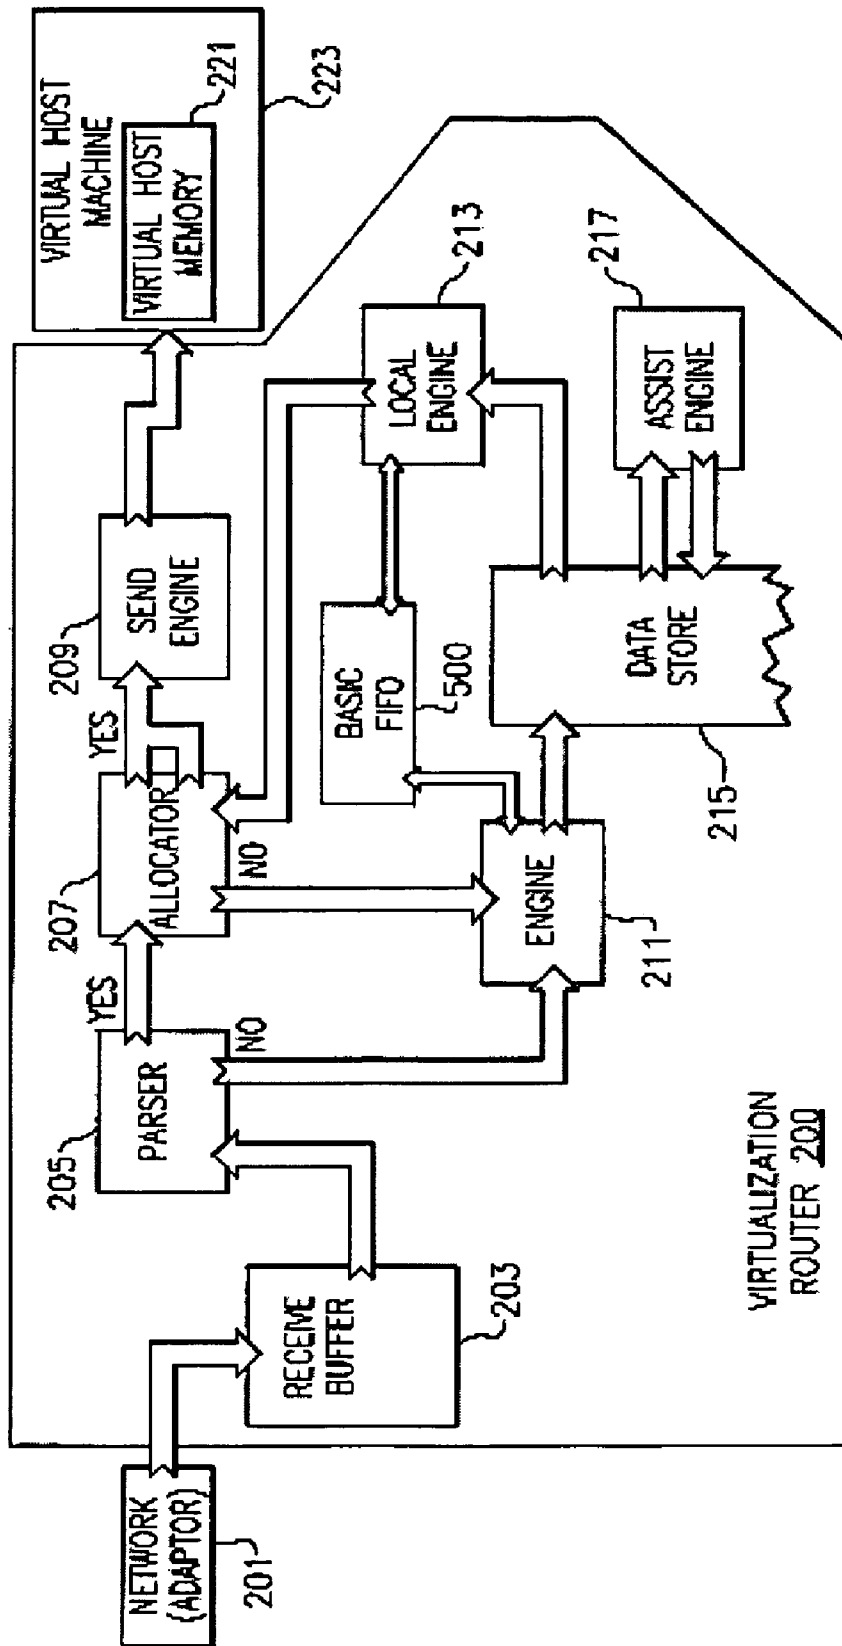


FIG. 2

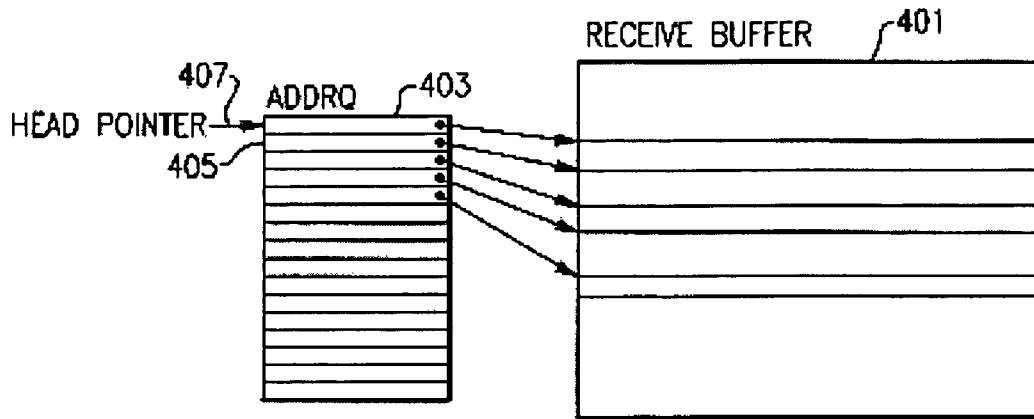


FIG.3

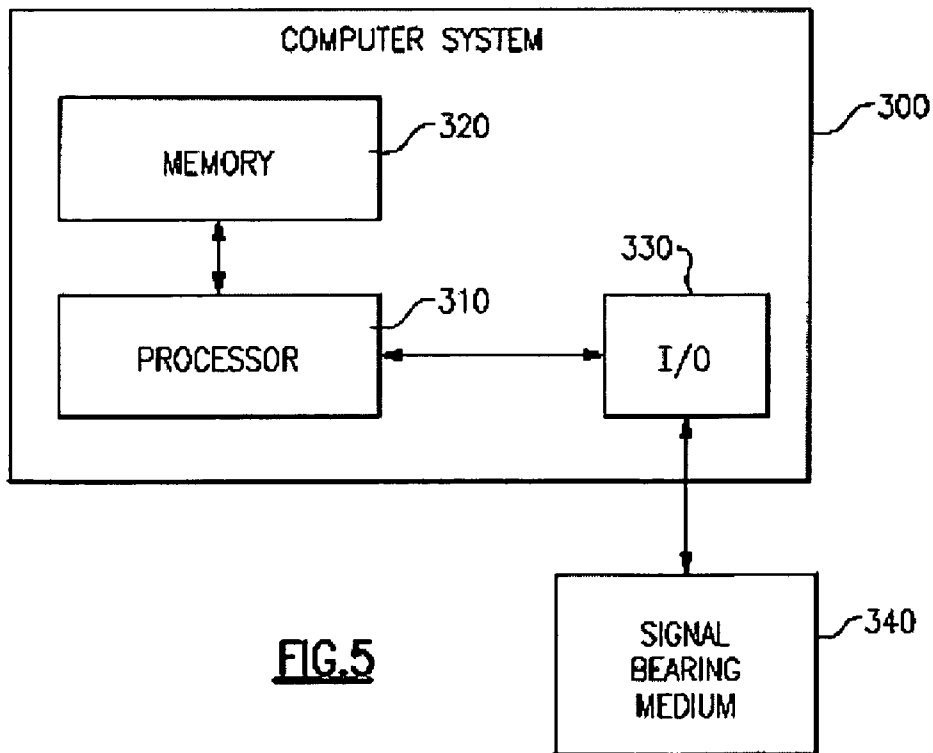


FIG.5

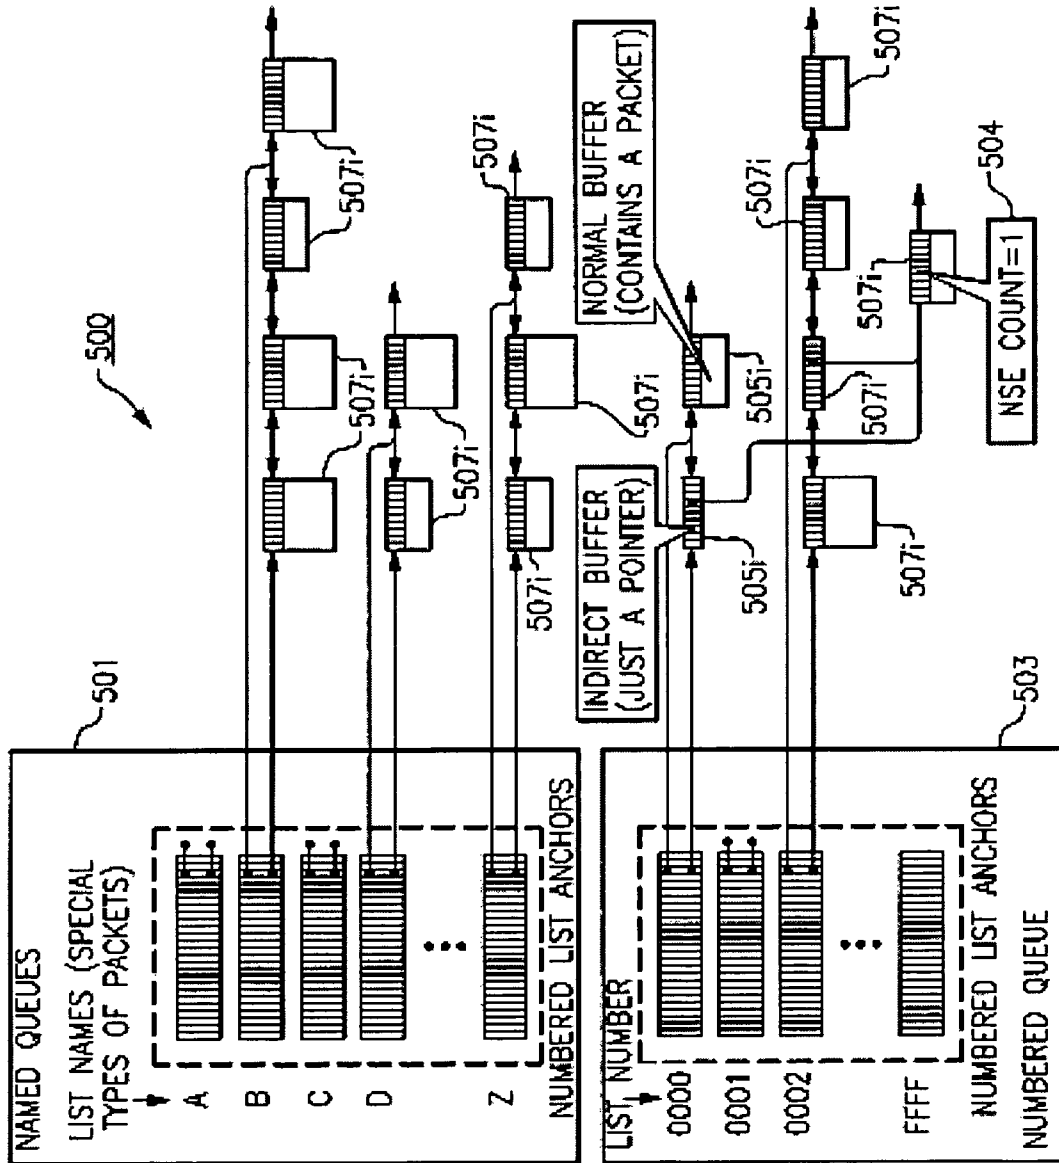


FIG. 4

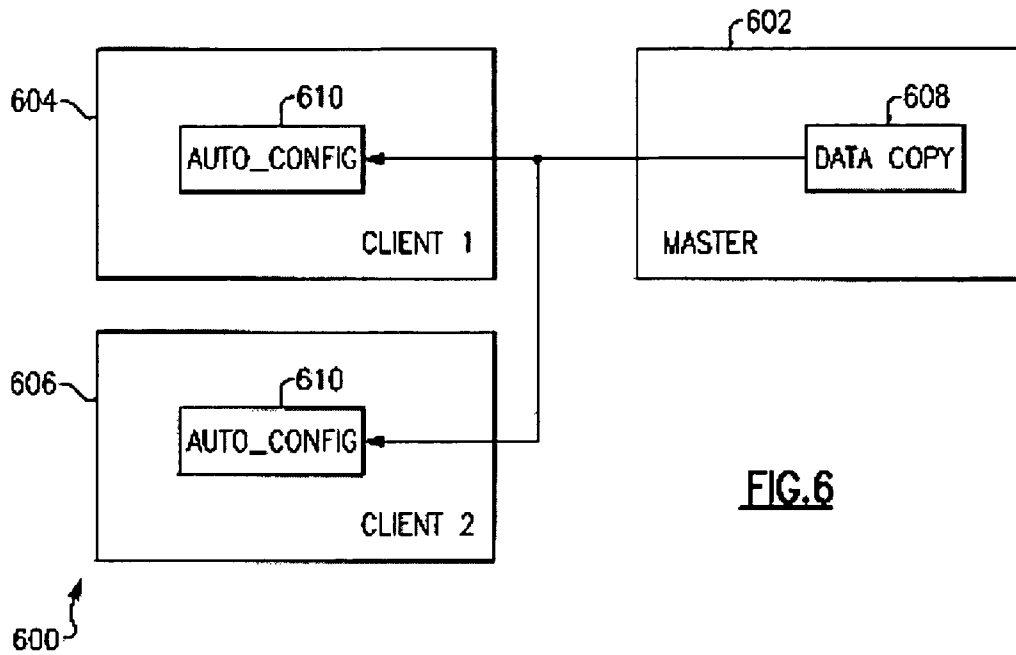


FIG.6

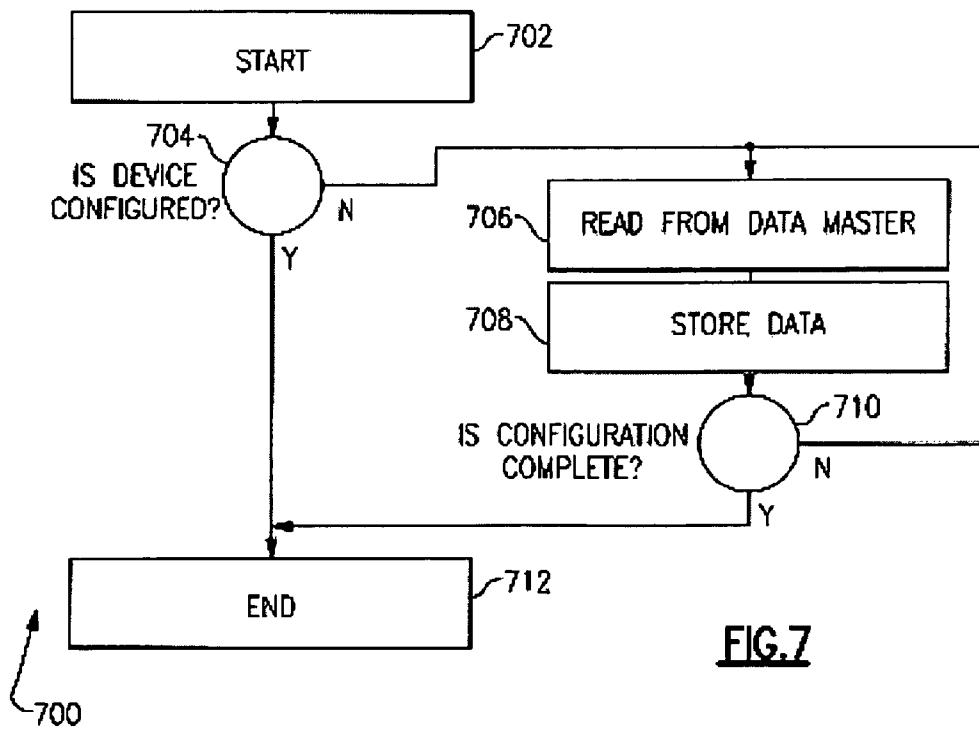


FIG.7

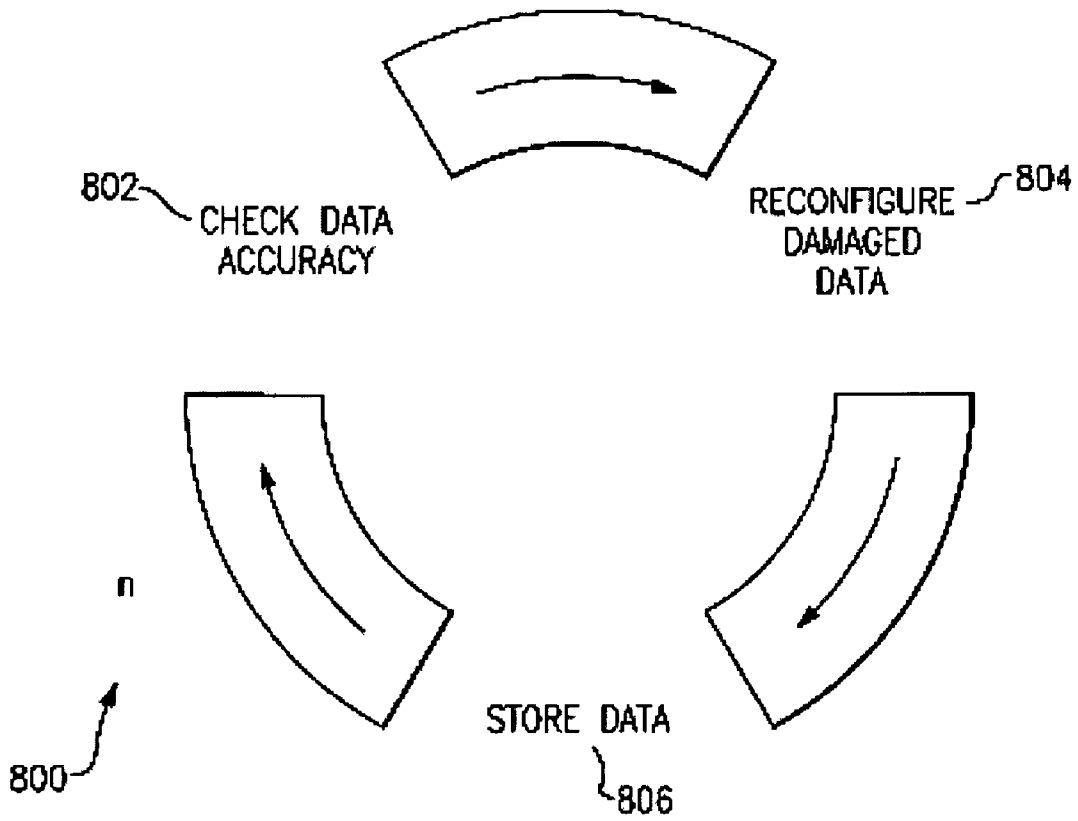


FIG.8

1

ETHERNET VIRTUALIZATION USING AUTOMATIC SELF-CONFIGURATION OF LOGIC

BACKGROUND OF THE INVENTION

This invention relates generally to network virtualization and, more particularly, to methods, systems, computer program products for implementing Ethernet virtualization routers using an automatic self-configuration of logic.

Current network virtualization techniques in multiple logic device systems require synchronization of common control parameters between them.

BRIEF SUMMARY OF THE INVENTION

In one embodiment, the present application provides an enhanced technique for Ethernet virtualization that may help overcome the foregoing deficiencies by self-configuration of logic. If one logic device is designated as the master, defined as the device containing the desired configuration values, one or more slaves can configure themselves based on those values. Logic can be implemented that handles auto-configuration as one or more slaves reading from this master. The slave devices can continually (or any periodic basis) check for changes in the configuration settings and update themselves with any such changes. They can also provide reporting and checking of these values. In such a manner, systems with multiple logic devices can easily stay in-sync without external intervention.

Exemplary embodiments include a method for Ethernet virtualization using an automatic self-configuration of logic. The method comprises maintaining control parameters at a master device, accessing, by a slave device, the control parameters at the master devices, and configuring the slave device based on the accessed control parameters.

A system and a computer program product corresponding to the above-summarized method is also described and claimed herein. Other systems, methods, computer program products, or hardware products according to embodiments will be or become apparent to one with skill in the art upon review of the following drawings and detailed description. It is intended that all such additional systems, methods, computer program products, and/or hardware products be included within this description, be within the scope of the present invention, and be protected by the accompanying claims.

Additional features and advantages are realized through the techniques of the present invention. Other embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed invention. For a better understanding of the invention with advantages and features, refer to the description and to the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Referring now to the drawings wherein like elements are numbered alike in the several FIGURES:

FIGS. 1A and 1B together comprise a flowchart setting forth an illustrative operational sequence for Ethernet virtualization to facilitate flow of unknown traffic to a virtual host;

FIG. 2 is a block diagram setting forth an illustrative system for Ethernet virtualization to facilitate flow of unknown traffic to a virtual host;

FIG. 3 is a data structure diagram setting forth an illustrative receive buffer for use with the system of FIG. 2;

2

FIG. 4 is a data structure diagram setting forth an illustrative elastic FIFO for use with the system of FIG. 2;

FIG. 5 is a block diagram setting forth an illustrative computer program product or hardware product for Ethernet virtualization to facilitate flow of unknown traffic to a virtual host;

FIG. 6 is a block diagram setting forth an illustrative system for implementing auto self-configuration of logic;

FIG. 7 is a flowchart illustrating the operational sequence for auto self-configuration of logic; and

FIG. 8 is a flowchart illustrating the operational sequence for auto self-configuration of logic including data verification and correction.

The detailed description explains the preferred embodiments of the invention, together with advantages and features, by way of example with reference to the drawings.

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

An exemplary embodiment of the present invention permits a hardware Ethernet virtualization router which serves a large number of communications channels to maintain high bandwidth by ensuring that packets can continue moving. The router avoids having to place back pressure on the network or having to drop packets, which are undesirable alternatives to moving packets. Thus, the technical effects and benefits include allowing an Ethernet virtualization router servicing any number of communications channels to continually move traffic efficiently regardless of packet types or shortages in channel-specific resources. High bandwidth can be maintained in spite of the unpredictable and bursty traffic patterns inherent to Ethernet networks. Packet order is maintained on a host-by-host basis and does not stall the flow of packets to one virtual machine because of a shortage of resources required for another virtual machine. Packets requiring special attention, such as multicast packets, packets for which connections cannot be determined by a parser, or address resolution protocol (ARP) packets, may also be handled using the methods disclosed herein.

FIGS. 1A and 1B together comprise a flowchart setting forth an illustrative operational sequence for Ethernet virtualization to facilitate flow of unknown traffic to a virtual host, and FIG. 2 is a block diagram setting forth an illustrative system for Ethernet virtualization using an elastic FIFO memory to facilitate flow of unknown traffic to a virtual host. Although FIGS. 1A and 1B show a linearly sequential series of operations, this is solely for purposes of clarity and illustration, as some of the steps could be performed in parallel or in a sequence other than what is shown in FIGS. 1A and 1B. The operational sequence of FIGS. 1A and 1B commences at block 101 (FIG. 1A) where a packet is received that represents unknown traffic destined for a virtual host. The packet is received from a network. The received packet includes a header. Illustratively, the received packet may be written into a relatively small optional receive buffer 203 (FIG. 2) as the packet arrives from a network adapter 201. The receive buffer 203 is not strictly required.

At block 103 (FIG. 1A), a parser 205 (FIG. 2) examines each packet's headers (for example, layer 2, layer 3, etc.) and a first test is performed to ascertain whether or not a destination connection can be determined for the received packet. The affirmative branch from block 103 leads to block 107, and the negative branch from block 103 leads to block 111 (described in greater detail hereinafter). If the parser 205 (FIG. 2) can determine a destination connection for the packet, following block 107 at block 105 (FIG. 1A), a second

test is performed to ascertain whether or not one or more connection-specific resources required to send the packet to a virtual host memory corresponding to the destination connection are available. The one or more connection-specific resources are available if they can be obtained for allocation by a resource allocation mechanism such as an allocator 207 (FIG. 2). For example, the allocator 207 attempts to obtain one or more connection-specific resources required to send that packet to a virtual machine 223 that corresponds to the destination connection. Note that the terms “virtual machine” and “virtual host” are used interchangeably herein. The affirmative branch from block 105 leads to block 109 (to be described in greater detail hereinafter), whereas the negative branch from block 105 leads to block 113 (to be described in greater detail hereinafter).

Although the illustrative configuration of FIG. 2 shows one virtual machine 223, this is for purposes of clarity, as the system of FIG. 2 may contain a plurality of virtual machines. For example, a practical system may include hundreds or thousands of virtual machines. With hundreds or thousands of virtual machines running on a single physical host, there may be scattered shortages of resources for some small subset of the virtual machines at any given time. Within the physical machine, a hypervisor may allocate differing resources to different virtual machines. For example, some virtual machines may have higher priority and therefore get more time-slices than others. Or, some virtual machines may have more physical storage allocated for use as receive buffers. A real-world system hosting thousands of virtual machines may not operate perfectly at all times with regard to resource management. Likewise, scattered resource shortages may be attributed to the inherently bursty nature of Ethernet traffic.

At block 107 (FIG. 1A), a local engine 213 (FIG. 2) blocks the allocator 207 from allocating resources to new packets when packets for the same connection already exist within an elastic first-in, first-out (FIFO) 500 memory because of a previous lack of resources. At block 109 (FIG. 1A), if the allocator 207 is successful in obtaining the resources, a send engine 209 (FIG. 2) writes the packet to a virtual host memory 221 associated with the virtual machine 223. If the parser 205 does not determine the packet's connection (FIG. 1A, block 111), it passes the packet to a store engine 211 (FIG. 2). At block 113 (FIG. 1A), if the allocator 207 (FIG. 2) fails in its attempt to obtain the required resources (which could be because the resources are not available or because the local engine currently has priority access to those resources), the allocator 207 passes the packet to the store engine 211.

Next, at block 115 (FIG. 1A), for each packet it is to service, the store engine 211 (FIG. 2) obtains a free packet buffer from the elastic FIFO 500. A free packet buffer is an empty packet-sized block of memory in a local data store 215. The store engine 211 moves the packet into that buffer (FIG. 1B, block 117) and submits the used buffer to the elastic FIFO 500 (FIG. 2). If a free packet buffer resource is not available, the packet is dropped or, optionally, the store engine 211 can wait for that shared resource to become available provided sufficient buffering, such as the receive buffer, is available. Since a packet's connection must be determined before it can be sent to the virtual host memory 221, at block 119 (FIG. 1B) an assist engine 217 (FIG. 2) determines and assigns connections to packets that were submitted to the elastic FIFO 500 without known connections (i.e. those packets which arrived from the parser 205).

The procedure of FIGS. 1A and 1B progresses to block 121 (FIG. 1B) where the local engine 213 (FIG. 2) continuously or periodically or repeatedly monitors both the allocator 207 for connection-specific resources and the elastic FIFO 500 for

the presence of packets with known destination connections. When resources become available for a connection which had previously lacked resources, block 129, the local engine 213 (FIG. 2) gains exclusive access to those resources, via the allocator 207, until no more packets for that connection exist in the elastic FIFO 500. The operations of block 129 (FIG. 1B) are critical to maintaining packet order within destination connections. At block 123 (FIG. 1B), when a connection with both waiting packets and available resources is seen, the packet is removed from the local data store 215 (FIG. 2) and passed to the allocator 207. The allocator 207 allocates the connection-specific resources required to send that packet to a virtual machine 223 corresponding to the connection destination (FIG. 1B, block 125). Since the local engine 213 (FIG. 2) already determined that the resources were available and claimed them for the packet, the allocator 207 is successful and the packet is written to virtual host memory 221 by the send engine 209 at block 127 (FIG. 1B).

The parser 205 (FIG. 2), allocator 207, send engine 209, store engine 211, local engine 213, local data store 215, elastic FIFO 500, and assist engine 217 together comprise a virtualization router 200. The router 200 is referred to as a virtualization router because it supports communication channels to a plurality of virtual machines which are called destination connections, such as virtual machine 223, providing the illusion to each virtual machine 223 that it possesses its own network interface card (such as the network adapter 201), when in fact only a single high-speed adapter (i.e., the network adapter 201) is present. The network adapter 201 is run in promiscuous mode to receive all packets off the network. The router 200 determines the correct connection for each packet and moves the packet to a memory space (i.e., virtual host memory 221) of a corresponding virtual machine 223.

Inbound Ethernet traffic flow is inherently bursty. Multiple shared and non-shared resources are required for each connection in order to move its packets. With inherently limited resources which must be carefully managed, one of the functions performed by the router 200 is to handle the bursty traffic in such a way as to minimize packet loss and retransmission for each connection in the context of virtual hosts, thereby maximizing network efficiency.

FIG. 3 is a data structure diagram setting forth an illustrative receive buffer 401 for implementing the optional receive buffer 203 of FIG. 2. Incoming bursts of packet data are initially written directly into the receive buffer 401 (FIG. 3), which is managed with a queue of records called an addrq 403. Each record 405 contains various information about a burst including a pointer into the receive buffer 401 where the data was written. All packets are removed from the receive buffer 401 in the same order that they were previously written (as the packets arrived from the network adapter 201, FIG. 2). If the packet at the head of the addrq 403 (FIG. 3) as indicated by a head pointer 407 is not a type of packet that requires special handling and if all required resources are available to move the packet, the packet is sent directly to its connection destination (virtual host memory 221, FIG. 2). Otherwise, if the packet is a special type or if one or more of the resources is not available, the packet is sent to the elastic FIFO 500.

As the head pointer 407 (FIG. 3) moves forward (i.e., in a downward direction wherein FIG. 3 is oriented such that the reference numerals are upright), packets for connections with ample resources continue to be sent directly to their connection destinations even while packets for other connections without resources are sent to the elastic FIFO 500 (FIG. 2). Similarly, by the same mechanism, packets that do not require special consideration can be sent directly to their destinations

while packets that require extra processing (e.g. multicast packets) can be sent to the elastic FIFO 500 to be serviced by an independently-running assist engine. In this way, the flow of all packets continues through the router 200 consistently without stalls or hesitations.

FIG. 4 is a data structure diagram setting forth an illustrative elastic FIFO 500 memory (FIGS. 2 and 4) for use with the system of FIG. 2 and the method of FIG. 1. The elastic FIFO 500 (FIG. 4) includes a large array of linked lists of packet buffers, as well as all the functionality required to maintain those lists and allow them to be put to a useful purpose. The elastic FIFO 500 serves to manage the flow of packets that can not be sent directly from the receive buffer 203 (FIG. 2) to one of the virtual machines such as the virtual machine 223. The elastic FIFO 500 (FIG. 4) is called "elastic" because it changes shape with changes in either or both of: (a) the number of active connections, or (b) the total number of packets contained by the FIFO. In its entirety, the memory footprint of the elastic FIFO 500 can be increased or decreased over time as free buffers are dynamically added or removed.

The elastic FIFO 500 is also called "elastic" because it is actually a collection of queues on which packet buffers and pointers to packet buffers are held. These pointers to packet buffers are shown in FIG. 4 as indirect buffers 505.i, where i is any positive integer greater than zero. The packet buffers are shown in FIG. 4 as normal buffers 507.j where j is any positive integer greater than zero. There are named queues 501 for free normal buffers 507.j of various sizes. These named queues 501 also include a named queue for free indirect buffers 505.i. Indirect buffers 505.i only occupy a small amount of memory, as they merely point to a normal buffer 507.j, but these indirect buffers 505.i do occupy some memory and are thus obtained as needed from the free list so that they may be enqueued to a numbered list. Since the elastic FIFO 500 is elastic, additional free indirect buffers 505.i can be added as necessary, and these buffers may also be taken away if the free list becomes excessively long. There is no direct relationship between the total number of indirect buffers 505.i and normal buffers 507.j. At any given time, some number of the normal buffers 507.j will be occupied by packets, though ideally most should be free most of the time.

There are two categories of queues within the elastic FIFO 500: named queues 501 and numbered queues 503. When a packet is sent to the elastic FIFO 500 because it requires special consideration, it is placed on a named queue of named queues 501, as opposed to being placed on a numbered connection queue of numbered queues 503. Each named queue of named queues 501 includes a list name A, B, C, D, Z that reflects a special packet type placed on that queue. Packets placed on named queues 501 must ultimately be moved to numbered queues 503. Packets on a numbered queue of numbered queues 503 can be sent to a corresponding destination connection's virtual host memory 221 (FIG. 2) as soon as the required connection-specific resources are available.

Named queues 501 (FIG. 4) are also used for implementing pools of various types of free buffers, which are buffers that are not currently associated with a packet. A free buffer is obtained for each packet that is to be locally stored, and that buffer is returned to the pool from which it originated once the packet has been removed from local storage and sent to virtual host memory 221 (FIG. 2). When the special packet type is multicast or broadcast, the packet must be moved from a "multicast" or "broadcast" named queue of named queues 501 (FIG. 4) to a plurality of numbered queues in numbered queues 503 so it may be sent to multiple virtual machines including virtual machine 223 (FIG. 2), wherein potentially

thousands of such virtual machines are present. The assist engine 217 efficiently performs this task through the use of indirect buffers 505.i (FIG. 4).

There are two types of packet buffers that may be put on a numbered queue of numbered queues 503: normal buffers 507.j and indirect buffers 505.i. Indirect buffers 505.i do not contain data themselves, but merely point to a normal buffer 507.j. The assist engine 217 (FIG. 2) removes a normal buffer 507.j (FIG. 4) from a named queue of named queues 501, obtains multiple free indirect buffers 505.i from the elastic FIFO 500, points those indirect buffers 505.i at the normal buffer 507.j, and enqueues those indirect buffers 505.i to the appropriate numbered queues 503.

Every normal buffer 507.j carries a use count 509. A normal buffer's use count 509 is usually 1 but can be higher when the normal buffer 507.j is the target of an indirect buffer 505.i. A normal buffer 507.j that is directly sitting on a queue has a use count of 1, while a normal buffer 507.j pointed to by one or more indirect buffers 505.i (which are sitting on one or more queues) has a use count equal to the number of indirect buffers 505.i pointing to it. A normal buffer 507.j that is the target of an indirect buffer 505.i can not itself directly exist on any queue. Each time a copy of the packet in the normal buffer 507.j is sent to virtual host memory 221 (FIG. 2), an indirect buffer 505.i (FIG. 4) pointing to it is removed from a numbered queue of numbered queues 503 and the normal buffer's use count 509 is decremented by 1 (provided it is still greater than 1). When the normal buffer's use count 509 reaches 1, it is returned to the pool of free normal buffers 507.j at the same time a final indirect buffer 505.i pointing to the normal buffer 507.j is dequeued and returned to the pool of free indirect buffers 505.i.

The local engine 213 (FIG. 2) performs the task of dequeuing buffers from the numbered queues 503 (FIG. 4), via the interface provided by the elastic FIFO 500, so the packets contained within or pointed to by those buffers can be sent to virtual host memory 221 (FIG. 2). As soon as the allocator 207 fails to obtain virtual host memory 221 for one packet for a given connection, it must send all subsequent packets for that connection to the local data store 215 in order to maintain packet order for that connection. One task performed by the local engine 213 is to empty queues containing packets that have accumulated, due to resource shortages, so packets may once again flow directly from the network adapter 201 to virtual host memory 221, i.e. without being stored first. The local engine 213 obtains exclusive access to connection-specific resources until the local engine determines that it has emptied a destination connection's queue and relinquishes that exclusivity.

FIG. 5 is a block diagram setting forth an illustrative computer program product or hardware product for Ethernet virtualization using an elastic FIFO memory to facilitate flow of unknown traffic to a virtual host. The system includes a computer 300 operatively coupled to a signal bearing medium 340 via an input/output interface (I/O) 330. The signal bearing medium 340 may include a representation of instructions for Ethernet virtualization using an elastic FIFO memory to facilitate flow of unknown traffic to a virtual host, and may be implemented as, e.g., information permanently stored on non-writeable storage media (e.g., read-only memory devices within a computer, such as CD-ROM disks readable by a CD-ROM drive), alterable information stored on a writeable storage media (e.g., floppy disks within a diskette drive or hard disk drive), information conveyed to a computer by a communications medium, such as through a computer or telephone network, including wireless or broadband communications networks, such as the Internet, etc.

The computer 300 includes a processor 310 that processes information for Ethernet virtualization using an elastic FIFO memory to facilitate flow of unknown traffic to a virtual host, wherein the information is represented, e.g., on the signal bearing medium 340 and communicated to the computer 300 via the I/O 330, wherein the processor 310 saves information as appropriate into a memory 320. Illustratively, the processor 310 corresponds to the processing mechanism 106 of FIG. 1. Returning now to FIG. 6, this information may also be saved into the memory 320, e.g., via communication with the I/O 330 and the signal bearing medium 340.

The processor 310 executes a program for implementing Ethernet virtualization to facilitate flow of unknown traffic to a virtual host. The processor 310 implements instructions for receiving a packet that represents unknown traffic destined for a virtual host on a network; performing a first test to ascertain whether or not a destination connection can be determined for the received packet wherein, if a destination connection can be determined, a second test is performed to ascertain whether or not one or more connection-specific resources required to send the packet to a virtual host memory corresponding to the destination connection are available; if a destination connection for the packet cannot be determined, passing the packet to a store engine; if the one or more connection-specific resources are not available, passing the packet to the store engine; the store engine obtaining a free packet buffer from an elastic first-in, first-out (FIFO) memory, wherein the free packet buffer is an empty packet-sized block of memory in a local data store; the store engine moving the packet into the free packet buffer and submitting the free packet buffer to the elastic FIFO memory; performing a monitoring procedure to detect both an availability of connection-specific resources and a presence of one or more waiting packets with a known destination connection; when a destination connection with: (i) one or more waiting packets, and (ii) available connection-specific resources; are both detected, removing the packet from the local data store; allocating the one or more connection-specific resources required to send the packet to the virtual host memory corresponding to the connection destination; and writing the packet to the virtual host memory. The foregoing steps may be implemented as a program or sequence of instructions within the memory 320, or on a signal bearing medium, such as the medium 340, and executed by the processor 310.

FIG. 6 shows a system 600 for automatically configuring a plurality of client devices from a master record located as a master device in accordance with an embodiment of the present invention. In one embodiment, the router 200 (FIG. 2) may be implemented as a plurality of hardware components. It is important to keep the configuration of the components synchronized for proper operation. The configuration of these components can be managed by automatic self-configuration of logic as disclosed in the present application. The system 600 of FIG. 6 can be used to automatically configure client devices from a master record. To illustrate, client devices 604, 606 have an interface to provide access to the master data copy from the master device 602.

In one embodiment, the interface comprise a read—write bus configuration allowing each client 604, 606 to read the master device 602. Data writes may also be performed for status updates or other uses. Client 604 or 606 raises a request line with an accompanying address and data size count, which targets a resource within the master copy 602. Since multiple devices may simultaneously request access to the same master copy 602, the targeted device must do internal arbitration and respond to each request separately. To complete the request, the master copy 602 responds with data and an

acknowledgement signal, and the client 604 or 606 drops the request line. A write is performed in a similar manner, except the client also provides data at the time the request is initiated, and the master copy responds only with an acknowledgement and no data. The client devices 604, 606 include auto-configuration logic 601 configured to allow the client devices to execute automatic self-configuration function. For example, such function can include configuration logic 700 of FIG. 7 and/or circular data checking logic 800 of FIG. 8, as explained below in detail. The client devices are typically implemented on an external programmable logic device, such as field programmable gate array (FPGA) and the master record is stored within another programmable logic device such as an application specific integrate circuit (ASIC).

In one embodiment, an Ethernet router that supports virtualization will scale to support multiple assistive devices simultaneously. That is, multiple parallel engines can be formed and tied together. In particular, engines such as 205 the packet parser might be duplicated or put into a programmable device such as 604, 606. All of these multiple engines need to be kept in synchronization with the master copy 602 of configuration data. This helps ensure consistency of operation within the virtualized router depicted in FIG. 2.

FIG. 7 is a flow chart 700 illustrating the operational sequence for automatic self-configuration of logic in accordance with an embodiment of the present application. The operational sequence can be implemented by the system 600 of FIG. 6. At block 702, a client device, such as device 604 (FIG. 6) begins execution of the automatic configuration function 610. At block 704, the client device checks if it is not already configured, or reconfiguration is desired. If it is not already configured, or reconfiguration, then, at block 706, the client device begins to read the master data 608 from the master device 602, via the interface as explained above. At block 708, when the master data arrives at the client device, it stores the master data in its memory or other storage. At block 710, the client device checks if there is additional data to be read. If there is additional data to be read, then the configuration of the client device is not complete and client device returns to block 706 to read additional master data and then to block 708 to store the read data. Once configuration is complete, the read—store loop exits and ends processing, at block 712.

FIG. 8 is a flow chart 800 illustrating the operational sequence for auto self-configuration of logic including data verification and correction in accordance with an embodiment of the present application. The flow chart 800 shows a circular data checking—reconfiguring loop which can be used with the system 600 of FIG. 6. At block 802, the client device checks the accuracy of the configuration data stored in the device. For example, the client device can check parity error or other failed check. If there is a data error, this event triggers the client device to proceed to block 804 to reconfigure itself and replace the damaged data. Then the client device proceeds to block 806 which include causing the client device to store the new data and check the data for accuracy. In this manner, the client device resumes normal operation.

It should be noted that the above techniques can be used with the computer system 300 and signal bearing medium 340 of FIG. 5.

As described above, the embodiments of the invention may be embodied in the form of computer-implemented processes and apparatuses for practicing those processes. Embodiments of the invention may also be embodied in the form of computer program code containing instructions embodied in tangible media, such as floppy diskettes, CD-ROMs, hard drives, or any other computer-readable storage medium, wherein,

when the computer program code is loaded into and executed by a computer, the computer becomes an apparatus for practicing the invention. The present invention can also be embodied in the form of computer program code, for example, whether stored in a storage medium, loaded into and/or executed by a computer, or transmitted over some transmission medium, such as over electrical wiring or cabling, through fiber optics, or via electromagnetic radiation, wherein, when the computer program code is loaded into and executed by a computer, the computer becomes an apparatus for practicing the invention. When implemented on a general-purpose microprocessor, the computer program code segments configure the microprocessor to create specific logic circuits.

While the invention has been described with reference to exemplary embodiments, it will be understood by those skilled in the art that various changes may be made and equivalents may be substituted for elements thereof without departing from the scope of the invention. In addition, many modifications may be made to adapt a particular situation or material to the teachings of the invention without departing from the essential scope thereof. Therefore, it is intended that the invention not be limited to the particular embodiment disclosed as the best mode contemplated for carrying out this invention, but that the invention will include all embodiments falling within the scope of the appended claims. Moreover, the use of the terms first, second, etc. do not denote any order or importance, but rather the terms first, second, etc. are used to distinguish one element from another.

We claim:

1. A method of Ethernet virtualization using automatic self-configuration of logic of a data router, the method comprising:

- 5 maintaining control parameters at a master device, wherein the master device comprises a first Ethernet virtualization router, the first Ethernet virtualization router facilitating flow of a first packet through a first communication channel to a first virtual machine, wherein the first Ethernet virtualization router comprising a plurality of the first engines comprising a first parser, a first allocator, a first send engine, a first store engine, a first local engine, a first local data store, a first memory and a first assist engine;
- 10 determining, by a slave device, that the slave device needs to be configured, wherein the slave device comprises a second Ethernet virtualization router, the second Ethernet virtualization router facilitating flow of a second packet through a second communication channel to a second virtual machine, wherein the second Ethernet virtualization router comprising a plurality of the first engines comprising a second parser, a second allocator, a second send engine, a second store engine, a second local engine, a second local data store, a second memory and a second assist engine;
- 15 based on the determining, accessing, by a the slave device, the control parameters at the master device; and
- 20 based on the determining, configuring the slave device based on the accessed control parameters, and synchronizing the plurality of the first engines with the plurality of the second engines.

* * * * *