

ROUTER

CROSS REFERENCE TO RELATED APPLICATION

[0001] The application claims the benefit of US provisional application No. 61/753,250 filed January 16, 2013, which is hereby incorporated by reference herein in its entirety.

FIELD

[0002] Some embodiments relate to communication network elements.

BACKGROUND

[0003] A communications network may include one or more network elements to facilitate the communication of data, for example, between computing devices.

SUMMARY

[0003a] According to a first aspect of the present invention, there is provided an apparatus including: a plurality of network interfaces including at least a first network interface, and a second network interface, wherein the first network interface is configured to a respective first network, and wherein the second network interface is configured to a respective second network, and a plurality of processors including at least a network processor for at least translating data, and a scheduling processor for at least transmitting the translated data by the network processor, the network processor configured to: obtain data received from the first network via the first network interface, perform network address translation on the data such that the data has an address on the second network, and communicate to the scheduling processor the data the network processor has translated from the first network to the second network, and obtain from the scheduling processor data that the scheduling processor has received from the second network via the second network interface, perform network address translation on the data such that the data has an address on the first network, and communicate the data to the first network via the first network interface, wherein the network processor obtains data received from the first network via the first network interface without use of the scheduling processor, and communicates data to the first network via the first network interface without use of the

scheduling processor; and the scheduling processor configured to: obtain from the network processor the data the network processor has translated from the first network to the second network and communicate the data to the second network via the second network interface; and obtain data received from the second network via the second network interface and communicate the data to the network processor for the network processor to perform network address translation on such that the data has an address on the first network, obtain data received from the second network work via the second network interface without use of the network processor, and communicates data to the second network via the second network interface without use of the network processor, and enable the scheduling processor to process interrupts from network interfaces that are configured to communicate with the scheduling processor, but disable the scheduling processor to process interrupts from network interfaces configured to communicate with the network processor.

[0003b] According to a second aspect of the present invention, there is provided a non-transitory computer-readable medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct: a network processor to at least translate data and a scheduling processor to at least transmit the translated data by the network processor, the network processor is configured to: obtain data received from a first network via a first network interface of a plurality of network interfaces, perform network address translation on the data such that the data has an address on a second network, and communicate to the scheduling processor the data the network processor has translated from the first network to the second network, and obtain from the scheduling processor data that the scheduling processor has received from the second network via a second network interface of the plurality of network interfaces, perform network address translation on the data such that the data has an address on the first network, and communicate the data to the first network via the first network interface, wherein the network processor obtains data received from the first network via the first network interface without use of the scheduling processor, and communicates data to the first network via the first network interface without use of the scheduling processor; the scheduling processor is configured to: obtain from the network processor the data the network processor has translated from the first network to the second network and communicate the data to the second network via the second network interface; and obtain data received from the second network via the second network interface and communicate the data to the network processor for the network processor to perform network address translation on such that the data has an address on the first network, obtain data received from the second network work via the

second network interface without use of the network processor, and communicates data to the second network via the second network interface without use of the network processor, enable the scheduling processor to process interrupts from network interfaces that are configured to communicate with the scheduling processor, but disable the scheduling processor to process interrupts from network interfaces configured to communicate with the network processor, wherein the first network interface is configured to be connected to a respective first network and the second network interface is configured to be connected to a respective second network.

BRIEF DESCRIPTION OF THE FIGURES

[0004] Figure 1A shows an example system according to some embodiments;

[0005] Figure 1B shows an example configuration of the system of Figure 1A;

[0006] Figure 2 shows another example system according to some embodiments; and

[0007] Figure 3 shows an example network element according to some embodiments.

DETAILED DESCRIPTION

[0008] Referring to Figure 1A, there is shown an example system 100. System 100 may include a plurality of entities, including entity 110 and a plurality of entities 130a-n (one of which is shown in detail in Figure 1A), one or more of which may be interconnected via network 103. Entity 110 may be, for example, a service provider that provides services, and each entity 130 may be, for example, a customer/user (e.g., such as a company, such as banks, investment funds, trading firm, etc.) of the services provided by service provider 110. For description purposes only, entity 110 will be referred to herein as a service provider and entity(s) 130 will be referred to as a user(s)/customer(s). Nonetheless, these terms are intended to be non-limiting and other example entities are possible. In addition, entity 110 being a service provider that provides services to entities 130 is also an example and other relationships between entities 110 and 130 are possible.

[0009] As one example, service provider 110 may provide one or more electronic marketplaces for the trading/buying-selling/matching of items (such as, for example, financial instruments, real-estate, bets/wagers, tangible goods, services, etc.) and as such, may provide

one or more electronic matching/trading engines. Similarly, customers 130 may seek to trade one or more items on the electronic market place(s) provided by service provider 130.

According to this example, one or more of customer 130 may electronically communicate data/messages, for example, to service provider 110 including, for example, orders to buy

5 and/or sell items at specified prices and/or quantities (e.g., bids, offers, hits, takes). Similarly, service provider 110 may electronically receive and execute such orders and communicate data/messages to customers 130 including, for example, prices and quantities of pending and executed orders. One skilled in the art will recognize that this is merely an example and that other and/or additional services may be provided by service provider 110 and that additional
10 and/or other messages/data may be transferred between services provider 110 and customers 130. For example, one or more of customer 130 may electronically communicate data/messages, for example, to service provider 110 including, for example, orders to back or lay a team or event, etc. at specified odds and/or stakes. Similarly, service provider 110 may electronically receive and execute/match such orders and communicate data/messages to
15 customers 130 including, for example, specified odds and/or stakes of pending and executed orders.

[0010] Service provider 110 may include one or more network elements 112, and one or more computing systems 114 (which may include/or be connected to one or more database
20 systems) that provide services, for example, to customers 130. One skilled in the art will recognize that service provider 110 may include additional and/or other computing systems, and/or network elements. Computing systems 114 may be referred to herein as servers for purposes of description. However, it should be understood that the use of the term server is non-limiting, and that other types of computing systems may be used. One or more of servers
25 114 may include one or more processors and one or more memories. One or more of servers 114 also may include one or more network hardware/software/firmware based interfaces/ports that enable the servers to connect to network elements 112 and thereby network 103. Such interfaces may be configured to support one or more different types of physical network connections, such as copper, fiber optic, and/or wireless, may be configured
30 to support one or more different types of protocols, such as Ethernet, and may be configured to operate at any speed, such as Gb rates. One skilled in the art will recognize that servers 114 may have additional and/or other configurations. Service provider 110 may also include one or more software and/or firmware and/or hardware based applications that may be stored

on one or more database systems and/or servers 114, and be configured to execute on one or more of servers 114. Each server may execute the same or different applications. As one example, the application(s) may be configured to provide one or more electronic matching/trading engines for the trading/matching of one or more items as described herein.

5

[0011] Network element(s) 112 may include, for example, one or more routers and/or switches including, for example, core and/or edge routers and/or switches. Each of network elements 112 may include one or more network hardware/software/firmware based interfaces/ports that enable the network elements to connect to one another, one or more of servers 114, and/or network 103. Such interfaces may be configured to support one or more different types of physical network connections, such as copper, fiber optic, and/or wireless, may be configured to support one or more different types of protocols, such as Ethernet, and may be configured to operate at any speed, such as Gb rates. One skilled in the art will recognize that network elements 112 may have additional and/or other configurations. One or more of network elements 112 may contain one or more physical connections (wired/wireless) to each other, servers 114, and/or network 103. Network elements 112 and one or more of servers 114 may be further configured such that one or more of servers 114 have private network addresses and thereby reside on a private network of service provider 110, and/or have public addresses and thereby reside on a public network. In this fashion, network elements 112 may be configured such that servers 114 may communicate with one another and/or with network 103 and thereby with one or more other computing systems, for example, connected to network 103, such as computing system(s) 138 of customers 130. One skilled in the art will also recognize that network elements 112 may include additional and/or other elements as described herein, and may provide additional and/or other type of functionality than described herein.

[0012] An example customer 130 (as shown by customer 130a) may include one or more network elements 132, 134, and 136, and one or more computing systems 138 (which may include/or be connected to one or more database systems). One skilled in the art will recognize that a customer 130 may include additional and/or other computing systems, and/or network elements. One skilled in the art will recognize that other customers 130 may include similar and/or other configurations as that of customer 130a shown in Figure 1A. Computing systems 138 may be referred to herein as servers for purposes of description. However, it

should be understood that the use of the term server is non-limiting, and that other types of computing systems may be used. One or more of servers 138 may include one or more processors and one or more memories. One or more of servers 138 also may include one or more network hardware/software/firmware based interfaces/ports that enable the servers to

5 connect to one or more of network elements 132-136 and thereby network 103. Such interfaces may be configured to support one or more different types of physical network connections, such as copper, fiber optic, and/or wireless, may be configured to support one or more different types of protocols, such as Ethernet, and may be configured to operate at any speed, such as Gb rates. One skilled in the art will recognize that servers 138 may have

10 additional and/or other configurations. Customer 130 may also include one or more software and/or firmware and/or hardware based applications that may be stored on one or more database systems and/or servers 138, and be configured to execute on one or more of servers 138. Each server may execute the same or different applications. As one example, the application(s) may be configured to use the services provided by server(s) 114 of service

15 provider 110 and in particular, may be configured to trade one or more items with one or more other customers 130, for example, through the use of electronic matching/trading engine(s) provided by server(s) 114 of service provider 110. According to this example, one or more of servers of 138 of customer 130 may electronically communicate via network 103 data/messages, for example, to servers 114 of service provider 110 including, for example,

20 orders to buy and/or sell items at specified prices and/or quantities (e.g., bids, offers, hits, takes). Similarly, servers 114 of service provider 110 may electronically receive and execute such orders and communicate data/messages to servers 138 of customer 130 including, for example, prices and quantities of pending and executed orders. One skilled in the art will recognize that this is merely an example and that other and/or additional services may be

25 provided by service provider 110 and that additional and/or other messages/data may be transferred between service provider 110 and customers 130, as described herein

[0013] Network element(s) 132-136 of customer 130 may include, for example, one or more routers and/or switches including, for example, core and/or edge routers and/or

30 switches. Each of network elements may include one or more network hardware/software/firmware based interfaces/ports that enable the network elements to connect to one another, to one or more of servers 138, and/or network 103. Such interfaces may be configured to support one or more different types of physical network connections,

such as copper, fiber optic, and/or wireless, may be configured to support one or more different types of protocols, such as Ethernet, and may be configured to operate at any speed, such as Gb rates. One skilled in the art will recognize that network elements 132-136 may have additional and/or other configurations. One or more of network elements 132-136 may contain one or more physical connections (wired/wireless) to each other, to one or more of servers 138, and/or network 103. Network elements 132-136 may be further configured such that one or more of servers 138 have private network addresses and thereby reside on a private network of a respective customer 130, and/or have public addresses and thereby reside on a public network. In this fashion, network elements 132-136 may be configured such that servers 114 may communicate with one another and/or with network 103 and thereby with one or more other computing systems, for example, connected to network 103, such as servers 114. One skilled in the art will also recognize that network elements 132-136 may include additional and/or other elements as described herein, and may provide additional and/or other type of functionality than that described herein.

15

[0014] Network 103 may include one or more network elements including, for example, one or more routers and/or switches. Such network elements may include one or more network hardware/software/firmware based interfaces/ports that may be configured to support one or more different types of physical network connections, such as copper, fiber optic, and/or wireless, may be configured to support one or more different types of protocols, such as Ethernet, and may be configured to operate at any speed, such as Gb rates. One or more elements of network 103 may contain one or more physical connections (wired/wireless) to each other and to each of entity 110 and entities 130. In this fashion, network 103 may be configured such that computing systems 114 of entity 110 and computing systems 138 of entity(s) 130 may communicate with at least one another. Network 103 may be configured as a public and/or private network(s). One skilled in the art will also recognize that network 103 may include additional and/or other network elements as described herein, and may be configured in additional and/or fashions than described herein.

30 **[0015]** Turning to Figure 1B where similar reference numbers refer to similar elements as described herein, there is shown a system 200 which may be an example configuration of system 100 of Figure 1A. According to this example, network elements 112 of service provider 110 may include a switch, such as core switch, that includes one or more

connections to each of servers 114. As an example, network element 112 may be an Arista 7124 application switch although other and/or additional network elements may be used. Network element 112 and one or more of servers 114 may be further configured such that one or more of servers 114 have network addresses on network 103, which address(es) may be referred to as “public” address(es) (although the address(es) may not actually be public). According to an example aspect of example system 200, network 103 may be a private network (possibly owned or leased) of service provider 110. According to this example, network element 112 and one or more of servers 114 would be part of network 103. In other words, the network interfaces of network element 112 that interface servers 114 and that also interface customers 130 may be on the same address space.

[0016] According to a further aspect of this example configuration of Figure 1B, network element 136 of an example customer 130 (as shown by customer 130a) may include a switch, such as core switch, network element 134 may be a switch, such as an edge switch, and network element 132 may be a router. Switch 136 may include one or more connections to each of servers 138 and one or more connections to switch 134. In turn, switch 134 may include one or more connections to router 136. Network elements 132-136 and one or more of servers 136 may be further configured such that one or more of servers 138 have private network addresses (i.e., addresses not on network 103) and thereby reside on a private network separate from network 103. Other customers 130n may have similar configurations.

[0017] According to a further aspect of this example configuration of Figure 1B, each customer 130 may have one (possibly more) addresses on network 103, which address(es) may be referred to as “public” address(es) (although the address(es) may not actually be public). Accordingly, each router 132 may also be configured as a network address and possibly port address translator (NAT/PAT) that maps between one or more private addresses of servers 138 on the customer 130 network and the public address assigned to the respective customer 130 on network 103, for example. Accordingly, when a server 138 is communicating a message/data (which may be encapsulated in a packet that includes addresses and/or ports for example) to a server 114, router 136 may translate a private address of server 138 in the packet to the public address assigned to customer 130 on network 103. Again, such translation may also include translating a port number being used by an application on server 138 to another port number. Similarly, when a server 110 is

communicating a message/data (which may be encapsulated in a packet that includes addresses and/or ports for example) to a server 138, router 136 may translate the public address assigned to customer 130 in the packet to the private address of server 138. Again, router 136 may also perform port translation as part of the address translation.

5

[0018] According to a further aspect of this example configuration of Figure 1B, network 103 may include point to point connections 113a-n (although again, non-point to point connections are possible) between switch 112 of service provider 110 and each of routers 132, for example, of customers 130. For example, each connection between switch 112 of service provider 110 and a router 132, for example, of a customer 130 may be a fiber connection such as a single mode fiber connection, running at 1Gb for example (although other types of connections and rates may be used). According to a still further aspect of this example configuration, one or more of: servers 114 and network element 112 of service provider 110, and network elements 132-136 and servers 138 of respective customers 130 may be co-located, such as in the same room. For example, one or more of servers 114 and network element 112 may reside in one rack. Similarly, one or more of network elements 132-136 and servers 138 of a first respective customer 130 may reside in another respective rack, etc. One skilled in the art will also recognize that system 200 may include additional and/or other elements than those described herein, and may include additional and/or other configurations than those described herein.

[0019] One example problem with example system 100 and 200 as shown in Figures 1A and 1B is that significant delays may be introduced when data is communicated between servers 114 and servers 138. As a specific example, router 132 of may introduce delays (e.g., greater than a 100us) as result of, for example, the network/port translation functionalities. Similarly, different routers 132 at different customers 130 may experience different delays.

[0020] Turning to Figure 2 where similar reference numbers refer to similar elements as described herein, there is shown an example system 300 that is similar to the example system 200 of Figure 1B. According to this example, system 200 includes network element 202. Network element 202 may be part of service provider 110 and as such, may be owned and/or operated by service provider 110. Network element may be co-located with servers 114 of service provider 110 and/or one or more network elements, including network element 112,

of service provider 110, and may reside in the same rack as these elements. One skilled in the art will recognize that network element 202 need not be owned and/or operated by service provider 110 and need not be co-located with the network elements and/or servers of service provider 110.

5

[0021] Network element 202 may include one or more network hardware/software/firmware based interfaces/ports 204a ... 204n that enable the network element to connect to servers 114, possibly via network element 112, for example. Network element 202 may also include one or more network hardware/software/firmware based
10 interfaces/ports 206a ... 206n that enable the network element to connect to server(s) 138, for example, of respective customers 130a-n. The network interfaces 204a-n and 206a-n of network element 202 may be configured to support one or more different types of physical network connections, such as copper, fiber optic, and/or wireless, may be configured to support one or more different types of protocols, such as Ethernet, and may be configured to
15 operate at any speed, such as Gb rates. Furthermore, different network interfaces 204a-n and 206a-n may have different configurations. One skilled in the art will recognize that network element 202 may have additional and/or other configurations.

[0022] As further shown in Figure 2, one or more of network interface(s) 204a-n of
20 network element 202 may be physically connected (wired/wirelessly), for example, by connection(s) 210a-n to network element 112, which may be a switch. According to another and/or additional example, one or more of network interface(s) 204a-n of network element 202 may each be physically connected by one or more connection(s) 210a-n directly to a respective server 114 of service provider 110. According to a further aspect of example
25 system 300, each customer 130a-n may be allocated one or more respective network interfaces 206a-n of network element 202. Accordingly, each respective network interface 206a-n of network element 202 may be physically connected (wired/wirelessly), for example, by a connection 212a-n to a respective server 138 of a respective customer 130a-n either directly and/or via one or more network elements (such as network elements 136, which may
30 be a switch for example) of a respective customer. For example, each connection 212a-n may be a fiber connection such as a single mode fiber connection, running at 1Gb for example (although other types of connections and rates may be used). One skilled in the art

will recognize that additional and/or other configurations of network element 202 and system 300 are possible.

[0023] According to a further aspect of example system 300, network interface(s) 204a-n of network element 202, connections 210a-n, network element 112, and one or more of servers 114 may reside on a network 214, which may be a private network of service provider 110 and have a network address space. Accordingly, network element 112 and one or more of servers 114 may be configured such that one or more of servers 114 have network addresses on network 214 within the network address space. According to a further aspect of this example configuration, network interface 206a, connection 212a, network element 136 of respective customers 130a, for example, and one or more of servers 138 of respective customers 130a may reside on a network 216a of customer 130a, which may be a private network of customer 130a and have a network address space. Accordingly, network element 136 and one or more of servers 138 of customer 130a may be configured such that the one or more of servers 138 have network addresses on network 216a within the respective network address space. Similarly, network interface 206n, connection 212n, network element 136 of customer 130n, for example, and one or more of servers 138 of customer 130n may reside on a network 216n of customer 130n, which may be a private network of customer 130n and have a respective network address space. Accordingly, network element 136 and one or more of servers 138 of customer 130n may be configured such that one or more of servers 138 of customer 130n have network addresses on network 216n within the respective network address space. Other customers 130 not shown in Figure 2 may have similar configurations.

[0024] According to a further aspect of this example system, network element 202 may be, for example, a bi-directional network address and possibly port address translator (NAT/PAT). More specifically, according to this example, each customer 130a-n may have one (possibly more) addresses on network 214 within the network address space of network 214. Accordingly, for each customer 130a-n, network element 202 may be configured as a NAT/PAT that maps between one or more addresses of servers 138 on the respective customer network 216a-n (i.e., within the network address space of the respective network) and the address assigned to the respective customer 130a-n on network 214, for example, within the network address space of network 214. Accordingly, when a server 138 is communicating a message/data (which may be encapsulated in a packet that includes

addresses and/or ports for example) to a server 114, network element 202 may translate an address on network 216a-n of server 138 in the packet to the address assigned to customer 130a-n on network 214. Again, such translation may also include translating a port number in the packet being used by an application on a server 138 to another port number. Similarly, when a server 110 is communicating a message/data (which may be encapsulated in a packet that includes addresses and/or ports for example) to a server 138, network element 202 may translate the address assigned to customer 130a-n on network 214 in the packet to the address of server 138 on network 216a-n. Again, network element 202 may also perform port translation as part of the address translation.

10

[0025] Turning to Figure 3 where similar reference numbers refer to similar elements as described herein, there is shown an example architecture of network element 202. Network element 202 may include a plurality of computing processors, including one or more of processors 301a ... 301n, and one or more of processors 302a, 302b, 302n. Processors 301a-n may be referred to herein as scheduling processors, and processors 302a-n may be referred to herein as network processors. The terms “scheduling” and “network” are intended to be non-limiting and are used herein for ease of discussion only. Each of processors 301a-n and 302a-n may or may not be similarly configured in terms of, for example, memory, processing speed, etc. One skilled in the art will recognize that network element 202 may include additional and or fewer processors. According to a further aspect of example network element 202, any one or more of scheduling processors 301a-n and network processors 302a-n may be interconnected to one another through a communications architecture, such as a bus architecture, which may include for example, a shared memory architecture. One skilled in the art will recognize that other and/or additional communication architectures are possible. The communication architecture may be such that any of the scheduling processor 301a-n may be configured to communicate with any one of the network processors 302a-n. One skilled in the art will recognize that other configurations are possible.

[0026] Network element 202 may also include, as discussed herein, one or more network interfaces 214a-n that may be configured, for example, to interface directly or indirectly with servers 114, and one or more network interfaces 206a-n that may be configured, for example, to interface with respective customers 130a-n. Network interfaces 204a-n and 206a-n may

have the same configuration and/or one or more different configurations. For example, interfaces 204a-n and 206a-n may be any combination of long range or short range, single mode or multi mode fiber interfaces running at 1Gb or 10Gb rates, for example. One or more of network interfaces 204a-n and 206a-n may be physically situated on its own network interface card and/or several of network interfaces 204a-n and 206a-n may be physically situated on a common network interface card or cards. One skilled in the art will recognize that other configurations are possible.

[0027] As discussed herein, each network interface 206a-n may be allocated to and interface with a respective customer 130a-n. One skilled in the art will recognize that a given customer may be allocated multiple network interfaces (for example, for load balancing purposes, backup, etc.). According to a further aspect of example network element 202, each network interface 206a-n may be assigned/allocated to a respective one of the network processors 302a-n such that all data communicated over the respective network interface is only handled by the respective network processor. One skilled in the art will recognize, however, that additional and/or other configurations are possible including assigning two or more network interfaces 206a-n to a given network processor 302a-n and/or assigning a two or more network processors 302a-n to a given network interface 206a-n. According to a further aspect of example network element 202, a given network interface 206a-n may be interconnected to its respective network processor 302a-n through a communications architecture, such as a bus architecture, such as a PCIe bus architecture. One skilled in the art will recognize that other and/or additional communication architectures are possible. The communication architecture may be such that any network interface 206a-n may be configured to communicate with any one or more of the network processors 302a-n and vice versa. According to a further aspect of network element 202, each network processor 302a-n may be shielded from all interrupts other than, for example, interrupts from its associated network interface 206a-n. One skilled in the art will recognize that other configurations are possible.

[0028] As further shown in Figure 3, example network element 202 may include one or more network applications 303a-n. Applications 303a-n may be software based applications, although other and/or additional configurations are possible including firmware and/or hardware based applications. Network element 202 may also include one or more memory

devices on which applications 303a-n may be stored and/or from which the applications may execute. Such memory devices may be electronically connected to one or more of processors 301a-n and 302a-n. According to an aspect of network element 202, each network processor may execute an application 303a-n. According to a further aspect of network element 202,
5 each application 303a-n may only execute on a respective processor. Hence, application 303a may only execute on processor 302a, application 303b may only execute on processor 302b, etc. This may be referred to as processor affinity. One skilled in the art will recognize, however, that additional and/or other configurations are possible including having multiple applications 303a-n execute on a single processor 302a-n and/or having one or more of
10 applications 303a-n execute on multiple processors.

[0029] As such, according to an example configuration of network element 202, a given network application 303a-n may be assigned/allocated to and execute on a respective network processor 302a-n, which may be assigned/allocated to a respective network interface 206a-n,
15 which may be assigned/allocated to a respective customer 130a-n. According to a further aspect of this example configuration, each network application 302a-n may be configured as a NAT and possibly a PAT and perform network/port translation for messages/data passing between network 214 and a respective network 216a-n of a respective customer 130a-n. Accordingly, network element 202 may be configured (such as by a network administrator) to
20 assign/allocate a given network interface 206a-n to a given network processor 302a-n and be further configured such that the network application 303a-n on each network processor performs NAT/PAT functions based on the customer 130a-n to which the assigned interface is connected. According to the physical hardware configuration/layout of network processors 302a-n and network interfaces 206a-n and/or the amount of data being generated by each
25 respective customer 130a-n and/or sent to each respective customer, certain assignments/allocations of network interfaces 206a-n to network processors 302a-n may be more beneficial than others in terms of total data throughput of network element 202 and be assigned accordingly, such as by an administrator. For example, it may be beneficial to allocate a given network interface 206a-n to a near network processor 302a-n. One skilled in the art
30 will recognize that network applications 302a-n need not be configured as NAT/PAT and that one network application may be configured to perform different functions than another network application, etc.

[0030] According to a further aspect of example network element 202, each network interface 204a-n may be assigned/allocated to a respective one of the scheduling processors 301a-n such that all data communicated over the respective network interface is only handled by the respective scheduling processor. One skilled in the art will recognize, however, that
5 additional and/or other configurations are possible. For example, all network interfaces 204a-n may be assigned/allocated to one of the scheduling processors 301a-n and/or each of the scheduling processors may be configured to communicate with any one or more of network interfaces 204a-n. According to a further aspect of example network element 202, a given network interface 204a-n may be interconnected to a respective scheduling processor(s)
10 301a-n through a communications architecture, such as a bus architecture, such as a PCIe bus architecture. One skilled in the art will recognize that other and/or additional communication architectures are possible. The communication architecture may be such that any network interface 204a-n may be configured to communicate with any one or more of the scheduling processors 301a-n and vice versa. According to a further aspect of network element 202, one
15 or more of scheduling processors 301a-n may be shielded from all interrupts other than, for example, interrupts from one or more of network interfaces 204a-n. One skilled in the art will recognize that other configurations are possible.

[0031] As further shown in Figure 3, network element 202 may also include one or more
20 scheduling applications and/or one or more administration applications, which are shown collectively in Figure 3 as applications 304a-n. Applications 304a-n may be software based applications, although other configurations are possible including firmware and/or hardware based applications. As discussed, network element 202 may also include one or more memory devices on which applications 304a-n may be stored and/or from which the
25 applications may execute. According to one example of network element 202, each scheduling processor 301a-n may execute one or more scheduling applications, one or more administration applications, or a combination of both. As another example, a scheduling application and/or an administration application may execute on multiple processors. As another example, a scheduling application may execute on one scheduling processor 301a-n
30 and be responsible for (as further discussed herein) a subset of network processors 302a-n, while another scheduling application may execute on another scheduling processor 301a-n and be responsible for a different subset of network processors 302a-n. As another example, a scheduling application may execute on one scheduling processor 301a-n and be responsible

for all network processors 302a-n. As a further example, a scheduling application may execute on multiple scheduling processors 301a-n and be responsible for all network processors 302a-n. One skilled in the art will recognize that other variations are possible. For ease of description, network element 202 will be described herein as having one
5 scheduling application and one administration application. One skilled in the art will recognize that network element 202 may include other types of applications than those described herein.

[0032] According to a further aspect of example network element 202, the scheduling
10 application 304 may be configured to pass messages/data between network interfaces 206a-n and network interfaces 204a-n as follows:

- a. For messages/data received on a network interfaces 206a-n, once the
respective network application 303a-n has performed NAT/PAT translation on
a respective packet, for example, scheduling application 304 may
15 retrieve/obtain the translated packet and forward/communicate the translated
packet onto/to one of network interfaces 204a-n.
- b. For messages/data received on a network interfaces 204a-n, scheduling
application 304 may obtain the messages/data, determine which of the
respective network applications 303a-n is performing NAT/PAT translation
20 for the respective packets, and forward/communicate the packets to that
network application, where the packet may be translated (NAT/PAT) and
forwarded/ communicated onto/to the respective network interface 206a-n.

[0033] As such, according to an example operation of network element 202, a given server
25 138 of a given customer 130a-n may transmit a message/data (which may be encapsulated in
a packet for example that includes addresses and/or ports) to a server 114 of service provider
110. That message/data/packet may be received at a given network interface 206a-n of
network element 202. A given network application 303a-n of the network processor 302a-n
that is assigned to that network interface 206a-n may then obtain/receive and/or be forwarded
30 the message/data/packet (or portion thereof) from the network interface 206 and perform
NAT/PAT translation on the respective packet. The network application 303a-n may then
forward and/or make available the translated message/data/packet to the scheduling
application 304 which in turn may then obtain and/or receive the translated

message/data/packet, and then forward and/or make available the translated message/data/packet to a respective network interface 204a-n, where the message/data/packet may then be communicated to the server 114. Similarly, a given server 114 of service provider 110 may transmit a message/data (which may be encapsulated in a packet for example) to a server 138 of the customer 130a-n. That message/data/packet may be received at a given network interface 204a-n of network element 202. The scheduling application 304 may then obtain/receive and/or be forwarded the message/data/packet (or a portion thereof) from the network interface 204a-n, determine which of the respective network applications 303a-n/network processors 302a-n is performing NAT/PAT translation for that message/data/packet, and forward and/or make available the message/data/packet to that network application 303a-n (as one example, and other implementations are possible, in order for the scheduling application to forward messages/data to the correct network processor 302a-n, the scheduling application may maintain a table – which may be configured by an administrator – that maps network addresses associated with customers 130 to respective network processors 302a-n). The network application 303a-n in turn may obtain and/or receive the message/data/packet, perform NAT/PAT translation on the respective packet and forward and/or make available the translated message/data/packet to a respective network interface 206a-n that is assigned to the network processor 302a-n on which the network application 303a-n is executing. That network interface 206 may then communicate the message/data/packet to server 138 of the customer 130a-n. One skilled in the art will recognize that this is only an example, and other and/or additional example operations of network element 202 are possible. One skilled in the art will also recognize that operations performed by scheduling application 304 as described herein and operations performed network applications 303a-n as described herein may be performed in whole or in part by the other.

[0034] Regarding the communications between scheduling application 304 and any given network application 303a-n, as one example the two applications may communicate via one or more (e.g., two) shared memory circular queues. For example, a network application 303 may place a translated message from a network interface 206 into a first memory queue, and update indexes/pointers, for example, to the queue to reflect that another message is in the queue. Similarly, scheduling application 304 may monitor the indexes/pointers to the queue, recognize there is a new message in the queue, read the message out, and update the

indexes/pointers, for example, to reflect that the message has been read out. Similarly, the scheduling application 304 may place a message from a network interface 204 into a second memory queue, and update indexes/pointers, for example, to that queue to reflect that another message is in the second queue. Similarly, network application 303 may monitor the
5 indexes/pointers to the second queue, recognize there is a new message in the queue, read the message out, and update the indexes/pointers, for example, to reflect that the message has been read out. A similar process may be used for other network applications 303 (which may have their own set of queues) to communicate with scheduling application 304. One skilled in the art will recognize that this is merely an example and other communication
10 techniques/processes may be used.

[0035] As discussed, network element 202 may also include an administration application 304 that may execute on a scheduling processor 301a-n, for example. Such an application may be used by an administrator to monitor the status of network element 202 and configure
15 the network element. For example, network element 202 may include, for example, one or more input/output devices such as a display interface, mouse, keyboard, touch screen, network interface (for remote access) etc. Through the use of such interfaces and administration application 304, for example, an administrator may monitor the status of network element 202 and configure the network element. For example, an administrator may
20 assign/allocate and/or re-assign/re-allocate a given network interface 206a-n to a given network processor 302a-n. The administrator may further configure the network application 303a-n of a given network processor 302 (e.g., as to the addresses and/or ports the network application 303 is translating between) according to the customer 130a-n that network application 303 is performing NAT/PAT translations on behalf of. One skilled in the art will
25 recognize that a network application may be configured in other fashions. One advantage, for example, of the hardware/software configuration of network element 202 is that an administrator, for example, may reconfigure a given customer 130a-n (in terms of, e.g., the NAT/PAT configuration, in terms of which network processor 302a-n is assigned to the customer, and/or in terms of which network interface 206a-n is assigned to the network
30 processor 302a-n and/or customer) without disrupting other customers. Administration application 304 may also include, for example, an IPMI subsystem so that an administrator can monitor the status of network element 202. Such a subsystem may also be separate from

administration application 304. One skilled in the art will recognize that these are only example administrative functions, and other/additional functions are possible.

[0036] According to a further aspect of network element 202, the network element may also include one or more memory devices, such as solid state drives, and may include one or more logging applications that capture all (or part) of the data moving through the network element and log the data to the drives. According to a further aspect, the logging application may filter the data and store only portions of the data and/or may run analysis on the data (such as latency calculations) and store such analysis. According to a still further aspect, the logging application may filter the data (such as looking for market data prices, completed transactions, etc.) and/or perform analysis on the data and forward the filtered data and/or analysis onto another network interface(s) (e.g. different from network interfaces 204a-n and 206a-n). Other systems and users (like customers 130), external to network element 202, may receive such data and/or analysis, including receiving the data and/or analysis for a price. One skilled in the art will recognize that these are only example logging functions, and other/additional logging functions are possible.

[0037] According to an example implementation of network element 202, an Intel Sandy Bridge processor that includes a plurality of cores (e.g., 8) may be used to provide one or more of processors 301a-n and 302a-n. For example, for a given Sandy Bridge processor, one or more cores may be allocated as scheduling processors 302a-n and one or more cores may be allocated as network processors 303a-n. Such cores may be configured and operate as described herein for processors 301a-n and 302a-n. Network element 202 may also include a multiple Sandy Bridge Processors. Here, for each Sandy Bridge processor, one or more cores may be allocated as network processors 302a-n and one or more cores may be allocated as scheduling processors 301a-n (that may, for example, only work with the network processors of that chip for purposes of moving messages between network interfaces 204a-n and 206a-n). In such a configuration, any of network interfaces 206a-n may be allocated to any of the network processors (cores) of the Sandy Bridges. Alternatively, only certain network interfaces 206a-n may be allocated to one Sandy Bridge processor while another set of network interfaces 206a-n may be allocated to another Sandy Bridge processor, etc. Similarly, any of network interfaces 204a-n may be allocated to any of the scheduling processors (cores) of the Sandy Bridge processors. Alternatively, only certain network

interfaces 204a-n may be allocated to one Sandy Bridge processor while another set of network interfaces 206a-n may be allocated to another Sandy Bridge processor, etc. One skilled in the art will recognize that these are merely examples, and that other configurations and other chip sets may be used.

5

[0038] According to a further example implementation of network element 202, network interfaces 204a-n and/or 206a-n may be provided through one or more network interface cards from Hotlava, for example, including any one or more of the Tambora 120G6, Tambora 64G6, Tambora 80G4, Tambora 64G4, and Bosavi 12G6. One skilled in the art will recognize that these are merely examples, and that other network interface cards, including those from other suppliers, may be used.

10

[0039] According to a further example implementation of network element 202, the Sandy Bridge processor(s), for example, and network interface cards from Hotlava, for example, may be inserted onto a single motherboard and such system may run a the Linux operating system. Again, one skilled in the art will recognize that these are merely examples, and that other configurations are possible.

15

[0040] One skilled in the art will recognize that network element 202 may also be configured (e.g., while also being configured as discussed herein in Figure 2 and 3) in a reverse direction. For example, a given server 114 on network 214 may have one (possibly more) addresses on another network, such as network 216a, although some other network interfaced to a network interface 206 is possible. Accordingly, for given server 114, network element 202 may be configured as a NAT/PAT (as discussed herein) that maps between one or more addresses of server 114 on customer network 214 and the address assigned to the server on network 216a, for example.

20

25

[0041] While this disclosure has been described in terms of certain embodiments and generally associated methods, alterations and permutations of the embodiments and methods will be apparent to those skilled in the art. Accordingly, the above description of example embodiments does not constrain this disclosure. Other changes, substitutions, and alterations are also possible without departing from the spirit and scope of this disclosure.

30

[0042] The following sections provide a guide to interpreting the present application.

I. Terms

[0043] The term “product” means a machine, manufacture and/or composition of matter,
5 unless expressly specified otherwise.

[0044] The term “process” means a process, algorithm, method or the like, unless expressly specified otherwise.

10 [0045] Each process (whether called a method, algorithm or otherwise) inherently includes one or more steps, and therefore all references to a “step” or “steps” of a process have an inherent antecedent basis in the mere description of a process, or in the mere recitation of the term ‘process’ or a like term. Accordingly, any reference in a claim to a ‘step’ or ‘steps’ of a process has sufficient antecedent basis.

15 [0046] The term “invention” and the like mean “the one or more inventions disclosed in this application”, unless expressly specified otherwise.

[0047] The terms “an embodiment”, “embodiment”, “embodiments”, “the embodiment”,
20 “the embodiments”, “one or more embodiments”, “some embodiments”, “certain embodiments”, “one embodiment”, “another embodiment” and the like mean “one or more (but not all) embodiments of the invention”, unless expressly specified otherwise.

[0048] The term “variation” of an invention means an embodiment of the invention, unless
25 expressly specified otherwise.

[0049] The term “indication” is used in an extremely broad sense. An “indication” of a thing should be understood to include anything that may be used to determine the thing.

30 [0050] An indication of a thing may include an electronic message that identifies the thing (e.g., an identification of a widget by a serial number affixed to the widget, an identification of a widget by one or more characteristics of the widget). An indication of a thing may include information that may be used to compute and/or look-up a thing (e.g., information identifying a machine of which a widget is a part that may be used to determine the widget).
35 An indication of a thing may specify things that are related to the thing (e.g., characteristics of the thing, a name of the thing, a name of a thing related to the thing). An indication of a

thing may not specify things that are related to the thing (e.g., a letter “a” may be an indication of a widget of a computer system that is configured to interpret the letter “a” to identify the widget). An indication of a thing may include a sign, a symptom, and/or a token of the thing. An indication, for example, may include a code, a reference, an example, a link, a signal, and/or an identifier. An indication of a thing may include information that represents, describes, and/or otherwise is associated with the thing.

[0051] A transformation of an indication of a thing may be an indication of the thing (e.g., an encrypted indication of a thing may be an indication of the thing). An indication of a thing may include the thing itself, a copy of the thing, and/or a portion of the thing. An indication of a thing may be meaningless to a thing that is not configured to understand the indication (e.g., a person may not understand that a letter “a” indicates a widget but it may nonetheless be an indication of the widget because the computer system may determine the widget from the letter “a”). It should be understood that the fact that an indication of a thing may be used to determine the thing does not mean that the thing or anything else is determined. An indication of a thing may include an indication of any number of the thing unless specified otherwise. An indication of a thing may include an indication of other things (e.g., an electronic message that indicates many things). (Indication can be used as a very broad term in claim language. For example: receiving an indication of a financial instrument.)

20

[0052] The term “represent” means (1) to serve to express, designate, stand for, or denote, as a word, symbol, or the like does; (2) to express or designate by some term, character, symbol, or the like; (3) to portray or depict or present the likeness of, as a picture does; or (4) to serve as a sign or symbol of.

25

[0053] A reference to “another embodiment” in describing an embodiment does not imply that the referenced embodiment is mutually exclusive with another embodiment (e.g., an embodiment described before the referenced embodiment), unless expressly specified otherwise. Similarly, the mere fact that two (or more) embodiments are referenced does not imply that those embodiments are mutually exclusive.

30

[0054] One embodiment of the invention may include or cover or embrace more than one other embodiment of the invention. For example, a first embodiment comprising elements a,

b, and c may cover a second embodiment that comprises elements a, b, c, and d as well as a third embodiment covering elements a, b, c, and e. Similarly, each of the first, second, and third embodiments may cover a fourth embodiment comprising elements a, b, c, d, and e.

5 **[0055]** The terms “including”, “comprising” and variations thereof mean “including but not necessarily limited to”, unless expressly specified otherwise. Thus, for example, the sentence “the machine includes a red widget and a blue widget” means the machine includes the red widget and the blue widget, but may possibly include one or more other items as well.

10 **[0056]** The term “consisting of” and variations thereof mean “including and also limited to”, unless expressly specified otherwise. Thus, for example, the sentence “the machine consists of a red widget and a blue widget” means the machine includes the red widget and the blue widget, but does not include anything else.

15 **[0057]** The term “compose” and variations thereof mean “to make up the constituent parts of, component of or member of”, unless expressly specified otherwise. Thus, for example, the sentence “the red widget and the blue widget compose a machine” means the machine includes the red widget and the blue widget.

20 **[0058]** The term “exclusively compose” and variations thereof mean “to make up exclusively the constituent parts of, to be the only components of, or to be the only members of”, unless expressly specified otherwise. Thus, for example, the sentence “the red widget and the blue widget exclusively compose a machine” means the machine consists of the red widget and the blue widget (i.e. and nothing else).

25 **[0059]** The terms “a”, “an” and “the” refer to “one or more”, unless expressly specified otherwise. Thus, for example, the phrase “a widget” means one or more widgets, unless expressly specified otherwise. Similarly, after reciting the phrase “a widget”, a subsequent recitation of the phrase “the widget” means “the one or more widgets”. Accordingly, it
30 should be understood that the word “the” may also refer to a specific term having antecedent basis. For example, if a paragraph mentions “a specific single feature” and then refers to “the feature,” then the phrase “the feature” should be understood to refer to the previously mentioned “a specific single feature.” (It should be understood that the term “a” in “a

specific single feature” refers to “one” specific single feature and not “one or more” specific single features.)

[0060] The term “plurality” means “two or more”, unless expressly specified otherwise.

5

[0061] The term “herein” means “in the present application, including anything which may be incorporated by reference”, unless expressly specified otherwise.

[0062] The phrase “at least one of”, when such phrase modifies a plurality of things (such as an enumerated list of things), means any combination of one or more of those things, unless expressly specified otherwise. For example, the phrase “at least one of a widget, a car and a wheel” means either (i) a widget, (ii) a car, (iii) a wheel, (iv) a widget and a car, (v) a widget and a wheel, (vi) a car and a wheel, or (vii) a widget, a car and a wheel. The phrase “at least one of”, when such phrase modifies a plurality of things does not mean “one of each of” the plurality of things. For example, the phrase “at least one of a widget, a car and a wheel” does not mean “one widget, one car and one wheel”.

[0063] Numerical terms such as “one “, “two “, etc. when used as cardinal numbers to indicate quantity of something (e.g., one widget, two widgets), mean the quantity indicated by that numerical term, but do not mean at least the quantity indicated by that numerical term. For example, the phrase “one widget” does not mean “at least one widget”, and therefore the phrase “one widget” does not cover, e.g., two widgets.

[0064] The phrase “based on” does not mean “based only on”, unless expressly specified otherwise. In other words, the phrase “based on” covers both “based only on” and “based at least on”. The phrase “based at least on” is equivalent to the phrase “based at least in part on”. For example, the phrase “element A is calculated based on element B and element C” covers embodiments where element A is calculated as the product of B times C (in other words, $A = B \times C$), embodiments where A is calculated as the sum of B plus C (in other words, $A = B + C$), embodiments where A is calculated as a product of B times C times D, embodiments where A is calculated as a sum of the square root of B plus C plus D times E, and so on.

[0065] The term “represent” and like terms are not exclusive, unless expressly specified otherwise. For example, the term “represents” does not mean “represents only”, unless

35

expressly specified otherwise. For example, the phrase “the data represents a credit card number” covers both “the data represents only a credit card number” and “the data represents a credit card number and the data also represents something else”.

5 **[0066]** The term “whereby” is used herein only to precede a clause or other set of words that express only the intended result, objective or consequence of something that is explicitly recited before the term “whereby”. Thus, when the term “whereby” is used in a claim, the clause or other words that the term “whereby” modifies do not establish specific further limitations of the claim or otherwise restrict the meaning or scope of the claim.

10 **[0067]** The terms “e.g.”, “such as” and like terms mean “for example”, and thus do not limit the term or phrase they explain. For example, in the sentence “the computer sends data (e.g., instructions, a data structure) over the Internet”, the term “e.g.” explains that “instructions” are an example of “data” that the computer may send over the Internet, and
15 also explains that “a data structure” is an example of “data” that the computer may send over the Internet. However, both “instructions” and “a data structure” are merely examples of “data”, and other things besides “instructions” and “a data structure” can be “data”.

20 **[0068]** The term “respective” and like terms mean “taken individually”. Thus if two or more things have “respective” characteristics, then each such thing has its own characteristic, and these characteristics can be different from each other but need not be. For example, the phrase “each of two machines has a respective function” means that the first of the two machines has a function and the second of the two machines has a function as well. The function of the first machine may or may not be the same as the function of the second
25 machine.

30 **[0069]** The term “i.e.” and like terms mean “that is”, and thus limits the term or phrase it explains. For example, in the sentence “the computer sends data (i.e., instructions) over the Internet”, the term “i.e.” explains that “instructions” are the “data” that the computer sends over the Internet.

[0070] A numerical range includes integers and non-integers in the range, unless expressly specified otherwise. For example, the range “1 to 10” includes the integers from 1 to 10 (e.g., 1, 2, 3, 4, ... 9, 10) and non-integers (e.g., 1.0031415926, 1.1, 1.2, ... 1.9).

[0071] Where two or more terms or phrases are synonymous (e.g., because of an explicit statement that the terms or phrases are synonymous), instances of one such term or phrase does not mean instances of another such term or phrase must have a different meaning. For example, where a statement renders the meaning of “including” to be synonymous with
5 “including but not limited to”, the mere usage of the phrase “including but not limited to” does not mean that the term “including” means something other than “including but not limited to”.

II. Determining

10 [0072] The term “determining” and grammatical variants thereof (e.g., to determine a price, determining a value, the determination of an object which meets a certain criterion) is used in an extremely broad sense. The term “determining” encompasses a wide variety of actions and therefore “determining” can include calculating, computing, processing, deriving, investigating, looking up (e.g., looking up in a table, a database or another data structure),
15 rendering into electronic format or digital representation, ascertaining and the like. Also, “determining” can include receiving (e.g., receiving information), accessing (e.g., accessing data in a memory) and the like. Also, “determining” can include resolving, selecting, choosing, establishing, and the like.

20 [0073] The term “determining” does not imply certainty or absolute precision, and therefore “determining” can include estimating, extrapolating, predicting, guessing, averaging and the like.

[0074] The term “determining” does not imply that mathematical processing must be
25 performed, and does not imply that numerical methods must be used, and does not imply that an algorithm is used.

[0075] The term “determining” does not imply that any particular device must be used. For example, a computer need not necessarily perform the determining.

30 [0076] The term “determining” may include “calculating”. The term “calculating” should be understood to include performing one or more calculations. Calculating may include computing, processing, and/or deriving. Calculating may be performed by a computing device. For example, calculating a thing may include applying an algorithm to data by a
35 computer processor and generating the thing as an output of the processor.

[0077] The term “determining” may include “referencing”. The term “referencing” should be understood to include making one or more reference, e.g., to a thing. Referencing may include querying, accessing, selecting, choosing, reading, and/or looking-up. The act of referencing may be performed by a computing device. For example, referencing a thing may include reading a memory location in which the thing is stored by a processor.

[0078] The term “determining” may include “receiving”. For example, receiving a thing may include taking in the thing. In some embodiments, receiving may include acts performed to take in a thing, such as operating a network interface through which the thing is taken in. In some embodiments, receiving may be performed without acts performed to take in the thing, such as in a direct memory write or a hard wired circuit. Receiving a thing may include receiving a thing from a remote source that may have calculated the thing.

III. Forms of Sentences

[0079] Where a limitation of a first claim would cover one of a feature as well as more than one of a feature (e.g., a limitation such as “at least one widget” covers one widget as well as more than one widget), and where in a second claim that depends on the first claim, the second claim uses a definite article “the” to refer to that limitation (e.g., “the widget”), this mere usage does not imply that the first claim covers only one of the feature, and this does not imply that the second claim covers only one of the feature (e.g., “the widget” can cover both one widget and more than one widget).

[0080] When an ordinal number (such as “first”, “second”, “third” and so on) is used as an adjective before a term, that ordinal number is used (unless expressly specified otherwise) merely to indicate a particular feature, such as to distinguish that particular feature from another feature that is described by the same term or by a similar term, but that ordinal number does not have any other meaning or limiting effect – it is merely a convenient name. For example, a “first widget” may be so named merely to distinguish it from, e.g., a “second widget”. Thus, the mere usage of the ordinal numbers “first” and “second” before the term “widget” does not indicate any other relationship between the two widgets, and likewise does not indicate any other characteristics of either or both widgets. For example, the mere usage of the ordinal numbers “first” and “second” before the term “widget” (1) does not indicate that either widget comes before or after any other in order or location; (2) does not indicate that either widget occurs or acts before or after any other in time; and (3) does not indicate

that either widget ranks above or below any other, as in importance or quality. The mere usage of ordinal numbers does not define a numerical limit to the features identified with the ordinal numbers. For example, the mere usage of the ordinal numbers “first” and “second” before the term “widget” does not indicate that there are exactly two widgets.

5

[0081] When a single device, article or other product is described herein, in another embodiment more than one device or article (whether or not they cooperate) may alternatively be used in place of the single device or article that is described. Accordingly, the functionality that is described as being possessed by a device may alternatively be
10 possessed by more than one device or article (whether or not they cooperate) in another embodiment.

[0082] Similarly, where more than one device, article or other product is described herein (whether or not they cooperate), in another embodiment a single device or article may
15 alternatively be used in place of the more than one device or article that is described. For example, a plurality of computer-based devices may be substituted with a single computer-based device. In some embodiments, such a plurality of computer-based devices may operate together to perform one step of a process such as is common in grid computing systems. In some embodiments, such a plurality of computer-based devices may operate provide added
20 functionality to one another so that the plurality may operate to perform one step of a process such as is common in cloud computing systems. (Conversely, a single computer-based device may be substituted with multiple computer-based devices operating in cooperation with one another. For example, a single computing device may be substituted with a server and a workstation in communication with one another over the internet) Accordingly, the
25 various functionality that is described as being possessed by more than one device or article may alternatively be possessed by a single device or article.

[0083] The functionality and/or the features of a single device that is described may, in another embodiment, be alternatively embodied by one or more other devices which are
30 described but are not explicitly described as having such functionality or features. Thus, other embodiments need not include the described device itself, but rather can include the one or more other devices which would, in those other embodiments, have such functionality or features.

IV. Disclosed Examples and Terminology Are Not Limiting

[0084] Neither the Title (set forth at the beginning of the first page of the present application) nor the Abstract (set forth at the end of the present application) is to be taken as limiting in any way the scope of the disclosed invention, is to be used in interpreting the meaning of any claim or is to be used in limiting the scope of any claim. An Abstract has been included in this application merely because an Abstract is required under 37 C.F.R. § 1.72(b).

[0085] The headings of sections provided in the present application are for convenience only, and are not to be taken as limiting the disclosure in any way.

[0086] Numerous embodiments are described in the present application, and are presented for illustrative purposes only. The described embodiments are not, and are not intended to be, limiting in any sense. The disclosed invention is widely applicable to numerous embodiments, as is readily apparent from the disclosure. One of ordinary skill in the art will recognize that the disclosed invention may be practiced with various modifications and alterations, such as structural, logical, software, and electrical modifications. Although particular features of the disclosed invention may be described with reference to one or more particular embodiments and/or drawings, it should be understood that such features are not limited to usage in the one or more particular embodiments or drawings with reference to which they are described, unless expressly specified otherwise.

[0087] Though an embodiment may be disclosed as including several features, other embodiments of the invention may include fewer than all such features. Thus, for example, a claim may be directed to less than the entire set of features in a disclosed embodiment, and such claim would not be interpreted as requiring features beyond those features that the claim expressly recites.

[0088] No embodiment of method steps or product elements described in the present application constitutes the invention claimed herein, or is essential to the invention claimed herein, or is coextensive with the invention claimed herein, except where it is either expressly stated to be so in this specification or (with respect to a claim and the invention defined by that claim) expressly recited in that claim.

[0089] Any preambles of the claims that recite anything other than a statutory class shall be interpreted to recite purposes, benefits and possible uses of the claimed invention, and such preambles shall not be construed to limit the claimed invention.

5 **[0090]** The present disclosure is not a literal description of all embodiments of the invention. Also, the present disclosure is not a listing of features of the invention which must be present in all embodiments.

[0091] All disclosed embodiments are not necessarily covered by the claims (even
10 including all pending, amended, issued and canceled claims). In addition, a disclosed embodiment may be (but need not necessarily be) covered by several claims. Accordingly, where a claim (regardless of whether pending, amended, issued or canceled) is directed to a particular embodiment, such is not evidence that the scope of other claims do not also cover that embodiment.

15 **[0092]** Devices that are described as in communication with each other need not be in continuous communication with each other, unless expressly specified otherwise. On the contrary, such devices need only transmit to each other as necessary or desirable, and may actually refrain from exchanging data most of the time. For example, a machine in
20 communication with another machine via the Internet may not transmit data to the other machine for long period of time (e.g. weeks at a time). In addition, devices that are in communication with each other may communicate directly or indirectly through one or more intermediaries. Devices are in communication with one another if they are capable of at least one-way communication with one another. For example, a first device is in communication
25 with a second device if the first device is capable of transmitting information to the second device. Similarly, the second device is in communication with the first device if the second device is capable of receiving information from the first device.

[0093] A description of an embodiment with several components or features does not imply
30 that all or even any of such components or features are required. On the contrary, a variety of optional components are described to illustrate the wide variety of possible embodiments of the present invention. Unless otherwise specified explicitly, no component or feature is essential or required.

[0094] Although process steps, algorithms or the like may be described or claimed in a particular sequential order, such processes may be configured to work in different orders. In other words, any sequence or order of steps that may be explicitly described or claimed does not necessarily indicate a requirement that the steps be performed in that order. The steps of
5 processes described herein may be performed in any order possible. Further, some steps may be performed simultaneously despite being described or implied as occurring non-simultaneously (e.g., because one step is described after the other step). Moreover, the illustration of a process by its depiction in a drawing does not imply that the illustrated process is exclusive of other variations and modifications thereto, does not imply that the
10 illustrated process or any of its steps are necessary to the invention, and does not imply that the illustrated process is preferred.

[0095] Although a process may be described as including a plurality of steps, that does not imply that all or any of the steps are preferred, essential or required. Various other
15 embodiments within the scope of the described invention include other processes that omit some or all of the described steps. Unless otherwise specified explicitly, no step is essential or required.

[0096] Although a process may be described singly or without reference to other products or methods, in an embodiment the process may interact with other products or methods. For example, such interaction may include linking one business model to another business model. Such interaction may be provided to enhance the flexibility or desirability of the process.

[0097] Although a product may be described as including a plurality of components, aspects, qualities, characteristics and/or features, that does not indicate that any or all of the plurality are preferred, essential or required. Various other embodiments within the scope of the described invention include other products that omit some or all of the described plurality.

[0098] An enumerated list of items (which may or may not be numbered) does not imply
30 that any or all of the items are mutually exclusive, unless expressly specified otherwise. Likewise, an enumerated list of items (which may or may not be numbered) does not imply that any or all of the items are comprehensive of any category, unless expressly specified otherwise. For example, the enumerated list “a computer, a laptop, and a PDA” does not imply that any or all of the three items of that list are mutually exclusive and does not imply
35 that any or all of the three items of that list are comprehensive of any category.

[0099] An enumerated list of items (which may or may not be numbered) does not imply that any or all of the items are equivalent to each other or readily substituted for each other.

[0100] All embodiments are illustrative, and do not imply that the invention or any
5 embodiments were made or performed, as the case may be.

V. Computing

[0102] It will be readily apparent to one of ordinary skill in the art that the various processes described herein may be implemented by, e.g., appropriately programmed general
10 purpose computers, special purpose computers and computing devices. Typically a processor (e.g., one or more microprocessors, one or more microcontrollers, one or more digital signal processors) will receive instructions (e.g., from a memory or like device), and execute those instructions, thereby performing one or more processes defined by those instructions. Instructions may be embodied in, e.g., one or more computer programs, one or more scripts.

15

[0103] The term “compute” shall mean to determine using a processor in accordance with a software algorithm.

[0104] A “processor” means one or more microprocessors, central processing units (CPUs),
20 computing devices, microcontrollers, digital signal processors, graphics processing units (GPUs) or like devices or any combination thereof, regardless of the architecture (e.g., chip-level multiprocessing or multi-core, RISC, CISC, Microprocessor without Interlocked Pipeline Stages, pipelining configuration, simultaneous multithreading, microprocessor with integrated graphics processing unit, GPGPU).

25

[0105] A “computing device” means one or more microprocessors, central processing units (CPUs), computing devices, microcontrollers, digital signal processors, graphics card, mobile gaming device, or like devices or any combination thereof, regardless of the architecture (e.g., chip-level multiprocessing or multi-core, RISC, CISC, Microprocessor without
30 Interlocked Pipeline Stages, pipelining configuration, simultaneous multithreading).

[0106] Thus a description of a process is likewise a description of an apparatus for performing the process. The apparatus that performs the process can include, e.g., a processor and those input devices and output devices that are appropriate to perform the
35 process. For example, a description of a process is a description of an apparatus comprising a

processor and memory that stores a program comprising instructions that, when executed by the processor, direct the processor to perform the method.

[0107] The apparatus that performs the process can include a plurality of computing devices that work together to perform the process. Some of the computing devices may work together to perform each step of a process, may work on separate steps of a process, may provide underlying services that other computing devices that may facilitate the performance of the process. Such computing devices may act under instruction of a centralized authority. In another embodiment, such computing devices may act without instruction of a centralized authority. Some examples of apparatus that may operate in some or all of these ways may include grid computer systems, cloud computer systems, peer-to-peer computer systems, computer systems configured to provide software as a service, and so on. For example, the apparatus may comprise a computer system that executes the bulk of its processing load on a remote server but outputs display information to and receives user input information from a local user computer, such as a computer system that executes VMware software.

[0108] Further, programs that implement such methods (as well as other types of data) may be stored and transmitted using a variety of media (e.g., computer readable media) in a number of manners. In some embodiments, hard-wired circuitry or custom hardware may be used in place of, or in combination with, some or all of the software instructions that can implement the processes of various embodiments. Thus, various combinations of hardware and software may be used instead of software only.

[0109] The term “computer-readable medium” refers to any non-transitory medium, a plurality of the same, or a combination of different media, that participate in providing data (e.g., instructions, data structures) which may be read by a computer, a processor or a like device. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media include, for example, optical or magnetic disks and other persistent memory. Volatile media include dynamic random access memory (DRAM), which typically constitutes the main memory. Transmission media include coaxial cables, copper wire and fiber optics, including the wires that comprise a system bus coupled to the processor. Transmission media may include or convey acoustic waves, light waves and electromagnetic emissions, such as those generated during radio frequency (RF) and infrared (IR) data communications. Common forms of

computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, any other magnetic medium, a CD-ROM, DVD, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, a RAM, a PROM, an EPROM, a FLASH-EEPROM, any other memory chip or cartridge, a carrier wave as
5 described hereinafter, or any other medium from which a computer can read.

[0110] The term “tangible computer-readable medium” refers to a “computer-readable medium” that comprises a hardware component, such as optical or magnetic disks.

10 **[0111]** Various forms of computer readable media may be involved in carrying data (e.g. sequences of instructions) to a processor. For example, data may be (i) delivered from RAM to a processor; (ii) carried over a wireless transmission medium; (iii) formatted and/or transmitted according to numerous formats, standards or protocols, such as Ethernet (or IEEE 802.3), wireless local area network communication defined by the IEEE 802.11 specifications
15 whether or not they are approved by the WiFi Alliance, SAP, ATP, Bluetooth™, and TCP/IP, TDMA, CDMA, and 3G; and/or (iv) encrypted to ensure privacy or prevent fraud in any of a variety of ways well known in the art.

[0112] The term “database” refers to any electronically-stored collection of data that is
20 stored in a retrievable format.

[0113] The term “data structure” refers to a database in a hardware machine such as a computer.

25 **[0114]** The term “network” means a series of points or nodes interconnected by communication paths. For example, a network can include a plurality of computers or communication devices interconnected by one or more wired and/or wireless communication paths. Networks can interconnect with other networks and contain subnetworks.

30 **[0115]** The term “predetermined” means determined beforehand, e.g., before a present time or a present action. For example, the phrase “displaying a predetermined value” means displaying a value that was determined before the act of displaying.

[0116] The term “condition” means (1) a premise upon which the fulfillment of an agreement depends, or (2) something essential to the appearance or occurrence of something else.

- 5 **[0117]** The term “transaction” means (1) an exchange or transfer of goods, services, or funds, or (2) a communicative action or activity involving two parties or things that reciprocally affect or influence each other.

[0118] Thus a description of a process is likewise a description of a computer-readable medium storing a program for performing the process. The computer-readable medium can store (in any appropriate format) those program elements which are appropriate to perform the method. For example, a description of a process is a description of a computer-readable storage medium that stores a program comprising instructions that, when executed by a processor, direct the processor to perform the method.

15

[0119] Just as the description of various steps in a process does not indicate that all the described steps are required, embodiments of an apparatus include a computer or computing device operable to perform some (but not necessarily all) of the described process.

- 20 **[0120]** Likewise, just as the description of various steps in a process does not indicate that all the described steps are required, embodiments of a computer-readable medium storing a program or data structure include a computer-readable medium storing a program that, when executed, can cause a processor to perform some (but not necessarily all) of the described process.

25

[0121] Where databases are described, it will be understood by one of ordinary skill in the art that (i) alternative database structures to those described may be readily employed, and (ii) other memory structures besides databases may be readily employed. Any illustrations or descriptions of any sample databases presented herein are illustrative arrangements for stored representations of information. Any number of other arrangements may be employed besides those suggested by, e.g., tables illustrated in drawings or elsewhere. Similarly, any illustrated entries of the databases represent exemplary information only; one of ordinary skill in the art will understand that the number and content of the entries can be different from those

described herein. Further, despite any depiction of the databases as tables, other formats (including relational databases, object-based models and/or distributed databases) could be used to store and manipulate the data types described herein. Likewise, object methods or behaviors of a database can be used to implement various processes, such as the described
5 herein. In addition, the databases may, in a known manner, be stored locally or remotely from a device which accesses data in such a database.

[0122] Various embodiments can be configured to work in a network environment including a computer that is in communication (e.g., via a communications network) with one
10 or more devices. The computer may communicate with the devices directly or indirectly, via any wired or wireless medium (e.g. the Internet, LAN, WAN or Ethernet, Token Ring, a telephone line, a cable line, a radio channel, an optical communications line, commercial on-line service providers, bulletin board systems, a satellite communications link, a combination of any of the above). Each of the devices may themselves comprise computers or other
15 computing devices, such as those based on the Intel[®], Pentium[®], or Centrino[™], Atom[™] or Core[™] processor, that are adapted to communicate with the computer. Any number and type of devices may be in communication with the computer.

[0123] In an embodiment, a server computer or centralized authority may not be necessary
20 or desirable. For example, the present invention may, in an embodiment, be practiced on one or more devices without a central authority. In such an embodiment, any functions described herein as performed by the server computer or data described as stored on the server computer may instead be performed by or stored on one or more such devices.

25 **[0124]** Where a process is described, in an embodiment the process may operate without any user intervention. In another embodiment, the process includes some human intervention (e.g., a step is performed by or with the assistance of a human).

[0125] As used herein, the term “encryption” refers to a process for obscuring or hiding
30 information so that the information is not readily understandable without special knowledge. The process of encryption may transform raw information, called plaintext, into encrypted information. The encrypted information may be called ciphertext, and the algorithm for transforming the plaintext into ciphertext may be referred to as a cipher. A cipher may also

be used for performing the reverse operation of converting the ciphertext back into plaintext. Examples of ciphers include substitution ciphers, transposition ciphers, and ciphers implemented using rotor machines.

- 5 **[0126]** In various encryption methods, ciphers may require a supplementary piece of information called a key. A key may consist, for example, of a string of bits. A key may be used in conjunction with a cipher to encrypt plaintext. A key may also be used in conjunction with a cipher to decrypt ciphertext. In a category of ciphers called symmetric key algorithms (e.g., private-key cryptography), the same key is used for both encryption and decryption.
- 10 The sanctity of the encrypted information may thus depend on the key being kept secret. Examples of symmetric key algorithms are DES and AES. In a category of ciphers called asymmetric key algorithms (e.g., public-key cryptography), different keys are used for encryption and decryption. With an asymmetric key algorithm, any member of the public may use a first key (e.g., a public key) to encrypt plaintext into ciphertext. However, only the
- 15 holder of a second key (e.g., the private key) will be able to decrypt the ciphertext back in to plaintext. An example of an asymmetric key algorithm is the RSA algorithm.

VI. Continuing Applications

- 20 **[0127]** The present disclosure provides, to one of ordinary skill in the art, an enabling description of several embodiments and/or inventions. Some of these embodiments and/or inventions may not be claimed in the present application, but may nevertheless be claimed in one or more continuing applications that claim the benefit of priority of the present application.
- 25 **[0128]** Applicants intend to file additional applications to pursue patents for subject matter that has been disclosed and enabled but not claimed in the present application.

VII. 35 U.S.C. § 112, paragraph 6

- 30 **[0129]** In a claim, a limitation of the claim which includes the phrase “means for” or the phrase “step for” means that 35 U.S.C. § 112, paragraph 6, applies to that limitation.

[0130] In a claim, a limitation of the claim which does not include the phrase “means for” or the phrase “step for” means that 35 U.S.C. § 112, paragraph 6 does not apply to that

limitation, regardless of whether that limitation recites a function without recitation of structure, material or acts for performing that function. For example, in a claim, the mere use of the phrase “step of” or the phrase “steps of” in referring to one or more steps of the claim or of another claim does not mean that 35 U.S.C. § 112, paragraph 6, applies to that step(s).

5

[0131] With respect to a means or a step for performing a specified function in accordance with 35 U.S.C. § 112, paragraph 6, the corresponding structure, material or acts described in the specification, and equivalents thereof, may perform additional functions as well as the specified function.

10

[0132] Computers, processors, computing devices and like products are structures that can perform a wide variety of functions. Such products can be operable to perform a specified function by executing one or more programs, such as a program stored in a memory device of that product or in a memory device which that product accesses. Unless expressly specified otherwise, such a program need not be based on any particular algorithm, such as any particular algorithm that might be disclosed in the present application. It is well known to one of ordinary skill in the art that a specified function may be implemented via different algorithms, and any of a number of different algorithms would be a mere design choice for carrying out the specified function.

20

[0133] Therefore, with respect to a means or a step for performing a specified function in accordance with 35 U.S.C. § 112, paragraph 6, structure corresponding to a specified function includes any product programmed to perform the specified function. Such structure includes programmed products which perform the function, regardless of whether such product is programmed with (i) a disclosed algorithm for performing the function, (ii) an algorithm that is similar to a disclosed algorithm, or (iii) a different algorithm for performing the function.

25

[0134] Where there is recited a means for performing a function that is a method, one structure for performing this method includes a computing device (e.g., a general purpose computer) that is programmed and/or configured with appropriate hardware to perform that function.

30

[0135] Also included is a computing device (e.g., a general purpose computer) that is programmed and/or configured with appropriate hardware to perform that function via other algorithms as would be understood by one of ordinary skill in the art.

5 VIII. Disclaimer

[0136] Numerous references to a particular embodiment do not indicate a disclaimer or disavowal of additional, different embodiments, and similarly references to the description of embodiments which all include a particular feature do not indicate a disclaimer or disavowal of embodiments which do not include that particular feature. A clear disclaimer or disavowal
10 in the present application will be prefaced by the phrase “does not include” or by the phrase “cannot perform”.

 IX. Incorporation By Reference

[0137] Any patent, patent application or other document referred to herein is incorporated
15 by reference into this patent application as part of the present disclosure, but only for purposes of written description and enablement in accordance with 35 U.S.C. § 112, paragraph 1, and should in no way be used to limit, define, or otherwise construe any term of the present application, unless without such incorporation by reference, no ordinary meaning would have been ascertainable by a person of ordinary skill in the art. Such person of
20 ordinary skill in the art need not have been in any way limited by any embodiments provided in the reference. Conversely, the definitions provided in this application should not be used to limit, define, or otherwise construe any term of any document incorporated herein by reference. The definitions set forth explicitly in this application are controlling notwithstanding the description of particular embodiments that may be incompatible with the
25 definition(s).

[0138] Any incorporation by reference does not, in and of itself, imply any endorsement of, ratification of or acquiescence in any statements, opinions, arguments or characterizations contained in any incorporated patent, patent application or other document, unless explicitly
30 specified otherwise in this patent application.

 X. Prosecution History

[0139] In interpreting the present application (which includes the claims), one of ordinary skill in the art refers to the prosecution history of the present application, but not to the prosecution history of any other patent or patent application, regardless of whether there are other patent applications that are considered related to the present application, and regardless
5 of whether there are other patent applications that share a claim of priority with the present application.

CLAIMS

1. An apparatus comprising:

a plurality of network interfaces including at least a first network interface, and a second network interface,

wherein the first network interface is configured to a respective first network, and

wherein the second network interface is configured to a respective second network, and

a plurality of processors including at least a network processor for at least translating data, and

a scheduling processor for at least transmitting the translated data by the network processor,

the network processor configured to:

obtain data received from the first network via the first network interface,

perform network address translation on the data such that the data has an address on the second network, and communicate to the scheduling processor the data the network processor has translated from the first network to the second network, and

obtain from the scheduling processor data that the scheduling processor has received from the second network via the second network interface, perform network address translation on the data such that the data has an address on the first network, and communicate the data to the first network via the first network interface,

wherein the network processor obtains data received from the first network via the first network interface without use of the scheduling processor, and communicates data to the first network via the first network interface without use of the scheduling processor; and

the scheduling processor configured to:

obtain from the network processor the data the network processor has translated from the first network to the second network and communicate the data to the second network via the second network interface; and

obtain data received from the second network via the second network interface and communicate the data to the network processor for the network processor to perform network address translation on such that the data has an address on the first network,

obtain data received from the second network work via the second network interface without use of the network processor, and communicates data to the second network via the second network interface without use of the network processor, and

enable the scheduling processor to process interrupts from network interfaces that are configured to communicate with the scheduling processor, but disable the scheduling processor to process interrupts from network interfaces configured to communicate with the network processor.

2. The apparatus of claim 1, further comprising at least a third network interface configured to a third network, and at least a third processor,

wherein the third processor is configured to:

obtain data received from the third network via the third network interface, perform network address translation on the data such that the data has an address on the second network, and communicate to the scheduling processor the data the third processor has translated from the third network to the second network, and

obtain from the scheduling processor data the scheduling processor has received from the second network via the second network interface, perform network address translation on the data such that the data has an address on the third network, and communicate the data to the third network via the third network interface,

wherein the third processor obtains data received from the third network work via the third network interface without use of the scheduling processor, and communicates data to the third network via the third network interface without use of the scheduling processor; and

wherein the scheduling processor is further configured to:

obtain from the third processor the data the third processor has translated from the third network to the second network and communicate the data to the second network via the second network interface; and

obtain data received from the second network via the second network interface and communicate the data to the third processor for the third processor to perform network address translation on such that the data has an address on the third network,

wherein the scheduling processor obtains data received from the second network work via the second network interface without use of the third processor, and communicates data to the second network via the second network interface without use of the third processor.

3. The apparatus of claim 2, wherein the apparatus is configured such that:

all data received by the first network interface from the first network is communicated from the first network interface only to the network processor; and

all data received by the third network interface from the third network is communicated from the third network interface only to the third processor.

4. The apparatus of claim 3, wherein the apparatus is re-configured such that:

all data received by the first network interface from the first network is communicated from the first network interface only to the third processor; and

all data received by the third network interface from the third network is communicated from the third network interface only to the network processor.

5. The apparatus of claim 1, wherein the network processor is further configured to perform port translation on data.
6. The apparatus of claim 1,
wherein a shared memory circular queue is associated with the network processor and the scheduling processor, the shared memory access queue having indexes associated therewith,
wherein to communicate by the network processor to the scheduling processor the data the network processor has translated from the first network to the second network includes to place by the network processor the data into the shared memory access queue,
wherein to obtain by the scheduling processor from the network processor the data the network processor has translated from the first network to the second network includes to read by the scheduling processor the data from shared memory access queue, and
wherein the network and scheduling processors update the indexes associated with the shared memory access queue to reflect data being placed into and read out of the shared memory access queue.
7. The apparatus of claim 6,
wherein another shared memory circular queue is associated with the network processor and the scheduling processor, the another shared memory access queue having indexes associated therewith,
wherein to communicate by the scheduling processor to the network processor the data the scheduling processor has received from the second network via the second network interface includes to place by the scheduling processor the data into the another shared memory access queue,
wherein to obtain by the network processor from the scheduling processor the data the scheduling processor has received from the second network via the second network interface includes to read by the network processor the data from the another shared memory access queue, and
wherein the network and scheduling processors update the indexes associated with the another shared memory access queue to reflect data being placed into and read out of the another shared memory access queue.
8. The apparatus of claim 3,
wherein for all data received by the first network interface from the first network to be communicated from the first network interface only to the network processor, and for all data received by the third network interface from the third network to be communicated from the third network interface only to the third processor includes to:
configure interrupts such that the network processor receives interrupts from the first network interface but not the third network interface, and
configure interrupts such that the third processor receives interrupts from the third network interface but not the first network interface.

9. The apparatus of claim 1, further comprising a semiconductor chip that has a plurality of cores, wherein a first of the cores comprises the network processor, and a second different one of the cores comprises the scheduling processor.

10. The apparatus of claim 1 further comprising:

- a third network interface configured to a third network,
- a fourth network interface configured to a fourth network,
- a third processor, and
- a fourth processor,

wherein the third processor is configured to:

- obtain data received from the third network via the third network interface, perform network address translation on the data such that the data has an address on the fourth network, and
- communicate to the fourth processor the data the third processor has translated from the third network to the fourth network, and

- obtain from the fourth processor data the fourth processor has received from the fourth network via the fourth network interface, perform network address translation on the data such that the data has an address on the third network, and communicate the data to the third network via the third network interface,

- wherein the third processor obtains data received from the third network work via the third network interface without use of the fourth processor, and communicates data to the third network via the third network interface without use of the fourth processor; and

wherein the fourth processor is configured to:

- obtain from the third processor the data the third processor has translated from the third network to the fourth network and communicate the data to the fourth network via the fourth network interface; and

- obtain data received from the fourth network via the fourth network interface and communicate the data to the third processor for the third processor to perform network address translation on such that the data has an address on the third network,

- wherein the fourth processor obtains data received from the fourth network work via the fourth network interface without use of the third processor, and communicates data to the fourth network via the fourth network interface without use of the third processor.

11. The apparatus of claim 10, further comprising a first and a second semiconductor chip each of which has a plurality of cores, wherein a first of the cores of the first chip comprises the network processor, a second different one of the cores of the first chip comprises the scheduling processor, a first of the cores of the second chip comprises the third processor, a second different one of the cores of the second chip comprises the fourth processor.

12. A non-transitory computer-readable medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:
a network processor to at least translate data and a scheduling processor to at least transmit the translated data by the network processor,

the network processor is configured to:

obtain data received from a first network via a first network interface of a plurality of network interfaces, perform network address translation on the data such that the data has an address on a second network, and communicate to the scheduling processor the data the network processor has translated from the first network to the second network, and

obtain from the scheduling processor data that the scheduling processor has received from the second network via a second network interface of the plurality of network interfaces, perform network address translation on the data such that the data has an address on the first network, and communicate the data to the first network via the first network interface,

wherein the network processor obtains data received from the first network via the first network interface without use of the scheduling processor, and communicates data to the first network via the first network interface without use of the scheduling processor;

the scheduling processor is configured to:

obtain from the network processor the data the network processor has translated from the first network to the second network and communicate the data to the second network via the second network interface; and

obtain data received from the second network via the second network interface and communicate the data to the network processor for the network processor to perform network address translation on such that the data has an address on the first network,

obtain data received from the second network work via the second network interface without use of the network processor, and communicates data to the second network via the second network interface without use of the network processor,

enable the scheduling processor to process interrupts from network interfaces that are configured to communicate with the scheduling processor, but disable the scheduling processor to process interrupts from network interfaces configured to communicate with the network processor,

wherein the first network interface is configured to be connected to a respective first network and the second network interface is configured to be connected to a respective second network.

13. The non-transitory computer-readable medium of claim 12, further comprising instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:

at least a third network interface configured to a third network, and at least a third processor,

wherein the third processor is configured to:

obtain data received from the third network via the third network interface, perform network address translation on the data such that the data has an address on the second network, and communicate to the scheduling processor the data the third processor has translated from the third network to the second network, and

obtain from the scheduling processor data the scheduling processor has received from the second network via the second network interface, perform network address translation on the data such that the data has an address on the third network, and communicate the data to the third network via the third network interface,

wherein the third processor obtains data received from the third network work via the third network interface without use of the scheduling processor, and communicates data to the third network via the third network interface without use of the scheduling processor; and

wherein the scheduling processor is further configured to:

obtain from the third processor the data the third processor has translated from the third network to the second network and communicate the data to the second network via the second network interface; and

obtain data received from the second network via the second network interface and communicate the data to the third processor for the third processor to perform network address translation on such that the data has an address on the third network,

wherein the scheduling processor obtains data received from the second network work via the second network interface without use of the third processor, and communicates data to the second network via the second network interface without use of the third processor.

14. The non-transitory computer-readable medium of claim 13, wherein the medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:

all data received by the first network interface from the first network is communicated from the first network interface only to the network processor; and

all data received by the third network interface from the third network is communicated from the third network interface only to the third processor.

15. The non-transitory computer-readable medium of claim 14, wherein the medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:

all data received by the first network interface from the first network is communicated from the first network interface only to the third processor; and

all data received by the third network interface from the third network is communicated from the third network interface only to the network processor.

16. The non-transitory computer-readable medium of claim 12, wherein the medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:

the network processor is further configured to perform port translation on data.

17. The non-transitory computer-readable medium of claim 12, wherein the medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:

a shared memory circular queue is associated with the network processor and the scheduling processor, the shared memory access queue having indexes associated therewith,

wherein to communicate by the network processor to the scheduling processor the data the network processor has translated from the first network to the second network includes to place by the network processor the data into the shared memory access queue,

wherein to obtain by the scheduling processor from the network processor the data the network processor has translated from the first network to the second network includes to read by the scheduling processor the data from shared memory access queue, and

wherein the network and scheduling processors update the indexes associated with the shared memory access queue to reflect data being placed into and read out of the shared memory access queue.

18. The non-transitory computer-readable medium of claim 17, the medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:

another shared memory circular queue is associated with the network processor and the scheduling processor, the another shared memory access queue having indexes associated therewith,

wherein to communicate by the scheduling processor to the network processor the data the scheduling processor has received from the second network via the second network interface includes to place by the scheduling processor the data into the another shared memory access queue,

wherein to obtain by the network processor from the scheduling processor the data the scheduling processor has received from the second network via the second network interface includes to read by the network processor the data from the another shared memory access queue, and

wherein the network and scheduling processors update the indexes associated with the another shared memory access queue to reflect data being placed into and read out of the another shared memory access queue.

19. The non-transitory computer-readable medium of claim 14, the medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:

wherein for all data received by the first network interface from the first network to be communicated from the first network interface only to the network processor, and for all data received by the third network interface from the third network to be communicated from the third network interface only to the third processor includes to:

configure interrupts such that the network processor receives interrupts from the first network interface but not the third network interface, and

configure interrupts such that the third processor receives interrupts from the third network interface but not the first network interface.

20. The non-transitory computer-readable medium of claim 12, the medium having instructions stored thereon that when executed by at least one processor comprising a scheduling processor and a network processor, direct:

- a third network interface configured to a third network,
- a fourth network interface configured to a fourth network,
- a third processor, and
- a fourth processor,

wherein the third processor is configured to:

obtain data received from the third network via the third network interface, perform network address translation on the data such that the data has an address on the fourth network, and communicate to the fourth processor the data the third processor has translated from the third network to the fourth network, and

obtain from the fourth processor data the fourth processor has received from the fourth network via the fourth network interface, perform network address translation on the data such that the data has an address on the third network, and communicate the data to the third network via the third network interface,

wherein the third processor obtains data received from the third network via the third network interface without use of the fourth processor, and communicates data to the third network via the third network interface without use of the fourth processor; and

wherein the fourth processor is configured to:

obtain from the third processor the data the third processor has translated from the third network to the fourth network and communicate the data to the fourth network via the fourth network interface; and

obtain data received from the fourth network via the fourth network interface and communicate the data to the third processor for the third processor to perform network address translation on such that the data has an address on the third network,

wherein the fourth processor obtains data received from the fourth network work via the fourth network interface without use of the third processor, and communicates data to the fourth network via the fourth network interface without use of the third processor.