



- (51) International Patent Classification:  
H04L 43/20 (2022.01) H04L 9/40 (2022.01)  
G06F 9/455 (2018.01)
- (21) International Application Number:  
PCT/US2022/020673
- (22) International Filing Date:  
17 March 2022 (17.03.2022)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
 

63/173,334	09 April 2021 (09.04.2021)	US
63/173,352	09 April 2021 (09.04.2021)	US
63/173,336	09 April 2021 (09.04.2021)	US
63/173,348	09 April 2021 (09.04.2021)	US
17/334,999	31 May 2021 (31.05.2021)	US

- (71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) Inventors: DEGRACE, Gerald Roy; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). BANSAL, Deepak; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). TEWARI, Rishabh; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). ZYGMUNT, Michal Czeslaw; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). JAGASIA, Deven; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (74) Agent: CHATTERJEE, Aaron C. et al.; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(54) Title: ARCHITECTURES FOR DISAGGREGATING SDN FROM THE HOST

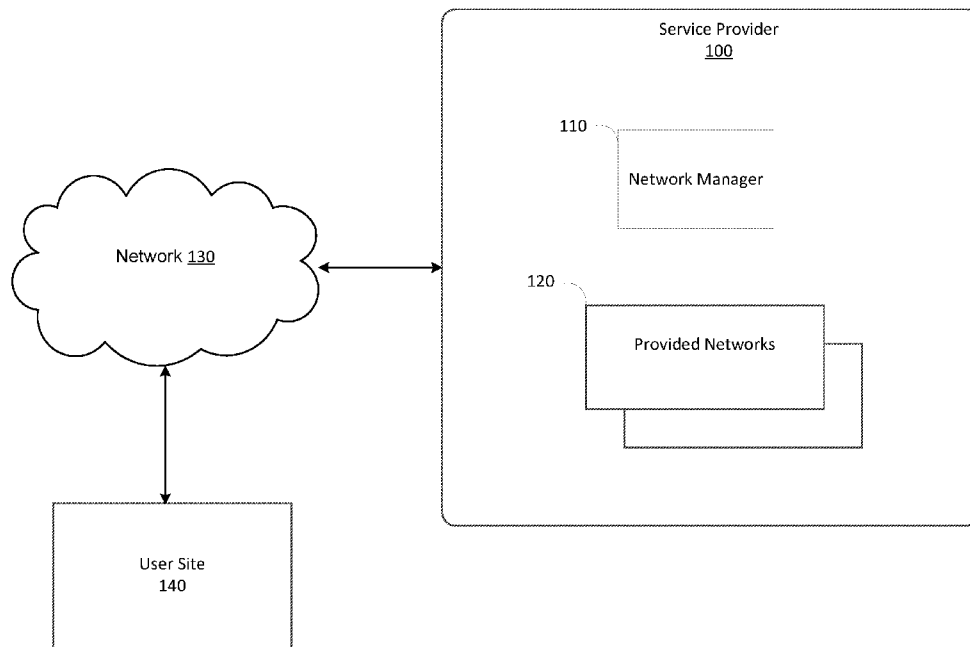


FIG. 1

(57) Abstract: Techniques are disclosed for processing data packets and implementing policies in a software defined network (SDN) of a virtual computing environment. At least one SDN appliance is configured to disaggregate enforcement of policies of the SDN from hosts of the virtual computing environment. The servers are communicatively coupled to network interfaces of the SDN appliance. The servers host a plurality of virtual machines. The SDN appliance comprises a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance.

WO 2022/216432 A1

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

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

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

**Published:**

- *with international search report (Art. 21(3))*

## ARCHITECTURES FOR DISAGGREGATING SDN FROM THE HOST

### BACKGROUND

5 [0001] A data center may house computer systems and various networking, storage, and other related components. Data centers may, for example, be used by service providers to provide computing services to businesses and individuals as a remote computing service or provide “software as a service” (e.g., cloud computing). Software defined networking (SDN) enables centralized configuration and management of physical and virtual network devices as well as dynamic and scalable implementation of network policies. The efficient  
10 processing of data traffic is important for maintaining scalability and efficient operation in such networks.

[0002] It is with respect to these considerations and others that the disclosure made herein is presented.

### SUMMARY

15 [0003] In some embodiments, the scenarios described above may benefit from the implementation of systems and methods for disaggregating policy processing off of the host machines. Disaggregation may also enable greater networking scale in order to match increasing demands from customers. In some embodiments, the SDN may implement a middle appliance, which may be referred to herein as an SDN appliance. The SDN appliance  
20 may incorporate some of the functionality of US Application Number xx/xxxxxx entitled “HARDWARE-BASED PACKET FLOW PROCESSING” filed on May 31, 2021.

[0004] In some embodiments, the SDN appliance may enable the use of the SDN control plane to manage network devices while providing high availability and fault tolerance, as further described herein. The SDN appliance provides a model to separate the application  
25 of SDN policies and configurations into a different computation environment. The operation of the SDN appliance is transparent to the virtual networks. Furthermore, the SDN appliance provides an opportunity to amortize the capabilities of the computation environment over many more virtual networks than was previously possible.

[0005] In the illustrated example embodiments, SDN capabilities may be enhanced by  
30 disaggregating policy enforcement from the host and moving it onto an SDN appliance strategically placed in the network. The SDN appliance may be configured to enforce SDN policies, perform associated transforms, and implement load balancer policies. In some embodiments, in order to move host SDN policy enforcement completely off the host, an SDN appliance including an FPGA may be used to move SDN policy enforcement off the

host. Implementation of the SDN appliance can free up work/compute capability for customer workloads and enable more predictable performance. The SDN appliance can be placed in datacenters to dynamically provide for a scale or feature that may not be possible or available on the host. Such features may include VM scaling, offloading of packet processing, and flexible SDN policy application.

[0006] Disclosed herein is a data center smart rack, which may also be referred to herein as a Cloud Smart Rack, that disaggregates SDN from the host. The data center smart rack distributes connection-based management functionality into the rack in a way that maximizes network resource utilization by distributing the network resources in efficient ways. The functionality may include SDN and offloaded storage scenarios.

[0007] Disclosed herein are methods for high availability (HA) for policy-based flow forwarding. High availability (HA) schemes may be implemented for a rack-based networking appliance. The HA methods may address failure modes including losing a smartNIC in the SDN appliance and losing the appliance altogether. The techniques may include the process of synchronizing to a new appliance or smartNIC once a failure occurs.

[0008] Disclosed herein are methods for scaling host policies via distribution across multiple SDN appliances. The disclosure provides for managing oversubscription of a rack-based networking appliance. The techniques may include spreading virtual machines across the appliance so that the capabilities of the appliance can be oversubscribed efficiently.

[0009] The described techniques can allow for virtual computing environments to support a variety of configurations including custom hardware and hybrid architectures while maintaining efficient use of computing resources such as processor cycles, memory, network bandwidth, and power. This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended that this Summary be used to limit the scope of the claimed subject matter. Furthermore, the claimed subject matter is not limited to implementations that solve any or all disadvantages noted in any part of this disclosure.

## DRAWINGS

[0010] The Detailed Description is described with reference to the accompanying figures. In the description detailed herein, references are made to the accompanying drawings that form a part hereof, and that show, by way of illustration, specific embodiments or examples. The drawings herein are not drawn to scale. Like numerals represent like elements throughout the several figures.

[0011] FIG. 1 is a diagram illustrating an example architecture in accordance with the present disclosure;

[0012] FIG. 2 is a diagram illustrating a data center in accordance with the present disclosure;

5 [0013] FIG. 3 is a diagram illustrating an example architecture in accordance with the present disclosure;

[0014] FIG. 4 is a diagram illustrating injected instances in accordance with the present disclosure;

[0015] FIG. 5 is a diagram illustrating an architecture for implementing virtual services  
10 in accordance with the present disclosure;

[0016] FIG. 6 is a diagram illustrating an architecture for implementing virtual services in accordance with the present disclosure;

[0017] FIG. 7 is a diagram illustrating an example architecture in accordance with the present disclosure;

15 [0018] FIG. 8 is a diagram illustrating an example architecture in accordance with the present disclosure;

[0019] FIG. 9 is a diagram illustrating an example architecture in accordance with the present disclosure;

[0020] FIG. 10 is a diagram illustrating an example architecture virtual network in  
20 accordance with the present disclosure;

[0021] FIG. 11 is a diagram illustrating an example architecture virtual network in accordance with the present disclosure;

[0022] FIG. 12A is a diagram illustrating an example architecture in accordance with the present disclosure;

25 [0023] FIG. 12B is a diagram illustrating an example architecture in accordance with the present disclosure;

[0024] FIG. 13 is a diagram illustrating an example architecture in accordance with the present disclosure;

[0025] FIG. 14 is a diagram illustrating an example architecture in accordance with the  
30 present disclosure;

[0026] FIG. 15 is a diagram illustrating an example architecture in accordance with the present disclosure;

[0027] FIG. 16 is a flowchart depicting an example procedure in accordance with the present disclosure;

[0028] FIG. 17 is an example computing system in accordance with the present disclosure.

#### DETAILED DESCRIPTION

[0029] The disclosed embodiments enable datacenters to provide services in a manner that can reduce the cost and complexity of their networks, allowing for more efficient use of computing, storage, and network resources. Efficient implementation of the end-to-end service by a cloud service provider can enable an experience that is seamless and more consistent across various footprints. The integration of multi-tenant and single-tenant resources with a comprehensive resource management approach can minimize the overhead for the user, who will not need to address policy enforcement issues and perform other complex management tasks. The effective distribution of the described disaggregation functions can be determined based on the implications for various performance and security implications such as latency and data security.

[0030] Disclosed herein is a data center smart rack, which may also be referred to herein as a Cloud Smart Rack, that disaggregates SDN from the host. The data center smart rack may be configured to distribute connection-based management functionality into the rack in a way that maximizes network resource utilization by distributing the network resources in efficient ways. The functionality may include software defined networking (SDN) and offloaded storage scenarios. In some embodiments, an SDN appliance may be implemented that is configured to enforce SDN policies, perform associated transforms, and implement load balancer policies. In one embodiment the appliance can be located on top of the data center smart rack rather than in other locations in the datacenter.

[0031] In an embodiment, the SDN appliance may be enabled with SmartNICs. In an embodiment, the servers in the data center smart rack may deploy “skinny” NICs which may be standard NICs without SmartNIC functionality. As used herein, the term hardware acceleration device may also encompass other ways of leveraging a hardware acceleration and offloading techniques to perform a function, such as, for example, a) a case in which at least some tasks are implemented in hard ASIC logic or the like; b) a case in which at least some tasks are implemented in soft (configurable) FPGA logic or the like; c) a case in which at least some tasks run as software on FPGA software processor overlays or the like; d) a case in which at least some tasks run as software on hard ASIC processors or the like, etc., or any combination thereof. In some embodiments, the peripheral device may be a network communications device, such as a network interface card (NIC). Such a NIC may be referred to herein as a smartNIC or sNIC.

[0032] In some embodiments disclosed herein, a flexible network interface may be

implemented. As used herein, such a flexible network interface may be referred to as a flexible network interface card, a floating network interface card, or fNIC, or more generally as a virtual port (vport). An fNIC may be inserted on a path from the host to the destination and may be configured to apply SDN policies before arriving at the destination. In some  
5 embodiments, one or more fNICs may be implemented on an SDN appliance. The point at which the SDN policy is implemented can float between the host and the SDN appliance as appropriate to the flow.

**[0033]** In conventional SDN applications, application of SDN policy may be based on tuple processing. When implemented as a middle box technology such as with the SDN  
10 appliance, traffic destinations behind the middle box in need of SDN operations may be pre-programmed as a fNIC that parses traffic comprised of a combination of custom defined identifiers such as VLAN, MAC, IP, and other information to uniquely identify flows and apply appropriate policy. This layer of programmability can provide flexibility for applying policies in different networking environments and scenarios.

**[0034]** An fNIC associated with a virtual machine (VM) in a cloud computing network may be configured to be elastically attached and detached from a parent NIC to thereby enable the virtual machine to simultaneously be connected to multiple different virtual networks (VNet) and/or subnets that are associated with the same or different subscriptions. The fNIC may, for example, enable a service provider to inject compute instances into an  
15 existing VNet in which the data plane uses a dedicated network interface to connect the customer's VNet, while another dedicated network interface provides management plane connectivity to the service provider. Such a configuration provides data plane isolation for the customer's VNet to comply with applicable security policies without disrupting management traffic between the injected resources and the service provider. Using a cross-  
20 subscription architecture, the parent NIC may be associated with a service subscription for management traffic to the injected compute instances, for example, and an attached fNIC associated with a customer subscription for data traffic.

**[0035]** In addition to the isolation provided between data and management traffic to the injected compute instances, utilization of the fNIC provides additional flexibility for cloud  
30 computing customers and service providers. For example, compute instances can be simultaneously connected to different subnets (which may have different security policies) in a customer's VNet. Such capabilities provided by the fNIC may advantageously promote efficient organization and consumption of resources in the customer's enterprise.

**[0036]** Utilization of the fNIC can support implementation of a multi-tenant architecture

to provide access by multiple tenants to a single shared VM. Each fNIC attached to a parent NIC associated with a service provider may use a unique network partition identifier (NPI) for each tenant subscription. The fNIC provides flexible implementation of multi-tenancy while enabling granular networking policies to be enforced to a particular discrete computing workload, rather than across the entire VM. A virtual filtering platform extension underlying the parent NIC may be configured to enforce specific networking policies that are tied to each fNIC including, for example, bandwidth metering, access control, VNet data encapsulation and addressing, etc. The data planes for each tenant on the VM may be operated concurrently yet are isolated from each other to ensure that data processing for one tenant has no impact on others.

**[0037]** The fNIC can also provide pre-provisioning of additional computing resources with associated policies that can be rapidly deployed on demand while reducing the time that is conventionally needed to inject resources into a customer's VNet. In such a "hot attach" architecture, a service provider can, for example, have a pool of already active compute instances on standby in a VM. The service can attach an fNIC to a parent NIC and associate it with the customer's subscription to provide access to the customer's VNet. Management and data planes operate independently to prevent disruption while providing conformance with applicable networking and security policies.

**[0038]** Referring to the appended drawings, in which like numerals represent like elements throughout the several FIGURES, aspects of various technologies for network disaggregation techniques and supporting technologies will be described. In the following detailed description, references are made to the accompanying drawings that form a part hereof, and which are shown by way of illustration specific configurations or examples.

**[0039]** FIG. 1 illustrates an example computing environment in which the embodiments described herein may be implemented. FIG. 1 illustrates a service provider 100 that is configured to provide computing resources to users at user site 140. The user site 140 may have user computers that may access services provided by service provider 100 via a network 130. The computing resources provided by the service provider 100 may include various types of resources, such as computing resources, data storage resources, data communication resources, and the like. For example, computing resources may be available as virtual machines. The virtual machines may be configured to execute applications, including Web servers, application servers, media servers, database servers, and the like. Data storage resources may include file storage devices, block storage devices, and the like. Networking resources may include virtual networking, software load balancer, and the like.

[0040] Service provider 100 may have various computing resources including servers, routers, and other devices that may provide remotely accessible computing and network resources using, for example, virtual machines. Other resources that may be provided include data storage resources. Service provider 100 may also execute functions that manage  
5 and control allocation of network resources, such as a network manager 110.

[0041] Network 130 may, for example, be a publicly accessible network of linked networks and may be operated by various entities, such as the Internet. In other embodiments, network 130 may be a private network, such as a dedicated network that is wholly or partially inaccessible to the public. Network 130 may provide access to computers  
10 and other devices at the user site 140.

[0042] FIG. 2 illustrates an example computing environment in which the embodiments described herein may be implemented. FIG. 2 illustrates a data center 200 that is configured to provide computing resources to users 200a, 200b, or 200c (which may be referred herein singularly as "a user 200" or in the plural as "the users 200") via user computers 202a, 202b,  
15 and 202c (which may be referred herein singularly as "a computer 202" or in the plural as "the computers 202") via a communications network 220. The computing resources provided by the data center 200 may include various types of resources, such as computing resources, data storage resources, data communication resources, and the like. Each type of computing resource may be general-purpose or may be available in a number of specific  
20 configurations. For example, computing resources may be available as virtual machines. The virtual machines may be configured to execute applications, including Web servers, application servers, media servers, database servers, and the like. Data storage resources may include file storage devices, block storage devices, and the like. Each type or configuration of computing resource may be available in different configurations, such as  
25 the number of processors, and size of memory and/or storage capacity. The resources may in some embodiments be offered to clients in units referred to as instances, such as virtual machine instances or storage instances. A virtual computing instance may be referred to as a virtual machine and may, for example, comprise one or more servers with a specified computational capacity (which may be specified by indicating the type and number of CPUs,  
30 the main memory size and so on) and a specified software stack (e.g., a particular version of an operating system, which may in turn run on top of a hypervisor).

[0043] Data center 200 may correspond to service provider 100 in FIGs. 1 and 2, or edge site 150 of FIG. 2. Data center 200 may include servers 226a, 226b, and 226c (which may be referred to herein singularly as "a server 226" or in the plural as "the servers 226") that

may be standalone or installed in server racks, and provide computing resources available as virtual machines 228a and 228b (which may be referred to herein singularly as "a virtual machine 228" or in the plural as "the virtual machines 228"). The virtual machines 228 may be configured to execute applications such as Web servers, application servers, media servers, database servers, and the like. Other resources that may be provided include data storage resources (not shown on FIG. 2) and may include file storage devices, block storage devices, and the like. Servers 226 may also execute functions that manage and control allocation of resources in the data center, such as a controller 225. Controller 225 may be a fabric controller or another type of program configured to manage the allocation of virtual machines on servers 226.

**[0044]** Referring to FIG. 2, communications network 220 may, for example, be a publicly accessible network of linked networks and may be operated by various entities, such as the Internet. In other embodiments, communications network 220 may be a private network, such as a corporate network that is wholly or partially inaccessible to the public.

**[0045]** Communications network 220 may provide access to computers 202. Computers 202 may be computers utilized by users 200. Computer 202a, 202b or 202c may be a server, a desktop or laptop personal computer, a tablet computer, a smartphone, a set-top box, or any other computing device capable of accessing data center 200. User computer 202a or 202b may connect directly to the Internet (e.g., via a cable modem). User computer 202c may be internal to the data center 200 and may connect directly to the resources in the data center 200 via internal networks. Although only three user computers 202a, 202b, and 202c are depicted, it should be appreciated that there may be multiple user computers.

**[0046]** Computers 202 may also be utilized to configure aspects of the computing resources provided by data center 200. For example, data center 200 may provide a Web interface through which aspects of its operation may be configured through the use of a Web browser application program executing on user computer 202. Alternatively, a stand-alone application program executing on user computer 202 may be used to access an application programming interface (API) exposed by data center 200 for performing the configuration operations.

**[0047]** Servers 226 may be configured to provide the computing resources described above. One or more of the servers 226 may be configured to execute a manager 230a or 230b (which may be referred herein singularly as "a manager 230" or in the plural as "the managers 230") configured to execute the virtual machines. The managers 230 may be a virtual machine monitor (VMM), fabric controller, or another type of program configured

to enable the execution of virtual machines 228 on servers 226, for example.

[0048] It should be appreciated that although the embodiments disclosed above are discussed in the context of virtual machines, other types of implementations can be utilized with the concepts and technologies disclosed herein.

5 [0049] In the example data center 200 shown in FIG. 2, a network device 222 may be utilized to interconnect the servers 226a and 226b. Network device 222 may comprise one or more switches, routers, or other network devices. Network device 222 may also be connected to gateway 240, which is connected to communications network 220. Network device 222 may facilitate communications within networks in data center 200, for example,  
10 by forwarding packets or other data communications as appropriate based on characteristics of such communications (e.g., header information including source and/or destination addresses, protocol identifiers, etc.) and/or the characteristics of the private network (e.g., routes based on network topology, etc.). It will be appreciated that, for the sake of simplicity, various aspects of the computing systems and other devices of this example are illustrated  
15 without showing certain conventional details. Additional computing systems and other devices may be interconnected in other embodiments and may be interconnected in different ways.

[0050] It should be appreciated that the network topology illustrated in FIG. 2 has been greatly simplified and that many more networks and networking devices may be utilized to  
20 interconnect the various computing systems disclosed herein. These network topologies and devices should be apparent to those skilled in the art.

[0051] It should also be appreciated that data center 200 described in FIG. 2 is merely illustrative and that other implementations might be utilized. Additionally, it should be appreciated that the functionality disclosed herein might be implemented in software,  
25 hardware or a combination of software and hardware. Other implementations should be apparent to those skilled in the art. It should also be appreciated that a server, gateway, or other computing device may comprise any combination of hardware or software that can interact and perform the described types of functionality, including without limitation  
30 desktop or other computers, database servers, network storage devices and other network devices, PDAs, tablets, smartphone, Internet appliances, television-based systems (e.g., using set top boxes and/or personal/digital video recorders), and various other consumer products that include appropriate communication capabilities. In addition, the functionality provided by the illustrated modules may in some embodiments be combined in fewer modules or distributed in additional modules. Similarly, in some embodiments the

functionality of some of the illustrated modules may not be provided and/or other additional functionality may be available.

**[0052]** In some embodiments, aspects of the present disclosure may be implemented in a mobile edge computing (MEC) environment implemented in conjunction with a 4G, 5G, or other cellular network. MEC is a type of edge computing that uses cellular networks and 5G and enables a data center to extend cloud services to local deployments using a distributed architecture that provide federated options for local and remote data and control management. MEC architectures may be implemented at cellular base stations or other edge nodes and enable operators to host content closer to the edge of the network, delivering high-bandwidth, low-latency applications to end users. For example, the cloud provider's footprint may be co-located at a carrier site (e.g., carrier data center), allowing for the edge infrastructure and applications to run closer to the end user via the 5G network.

**[0053]** FIG. 3 shows an illustrative cloud computing environment 300 in which a customer network 305 includes multiple portions including an on-premises network 330 and a virtual network (VNet) 335. The customer network in this example is a hybrid network but other network configurations may also be utilized depending on the particular requirements of the user scenario. The VNet may be physically implemented using one or more host machines 320 that are operated by a cloud service provider 325. It is noted that the diagram in FIG. 3 is simplified for clarity in exposition and typical networking equipment such as firewalls, routers, and the like are not shown.

**[0054]** The on-premises network and VNet are typically operatively coupled using instances of gateways 330, or other networking devices, over a communication network 335 which may include, for example, private and/or public networking infrastructure using various combinations of connectivity services. The VNet may include multiple subnets 340 that each include one or more instances of virtual machines 345 that are typically connected using load balancers 350 and/or other networking devices. Security and other networking policies (collectively indicated by reference numeral 355) are typically applicable to each subnet. The networking policies are typically different for each subnet, but they can be the same and/or overlap in some cases.

**[0055]** FIG. 4 shows an illustrative service from a service resource provider 405 that injects dedicated compute instances 410 into the customer's VNet 415. For example, in some implementations, such service resource providers may be implemented using a Platform as a Service (PaaS) to provide search, content delivery, etc. The service resource provider may be associated with the cloud service provider 405 or be a third party in some

cases. Service resource providers may inject compute instances or other resources into a VNet when provisioning a given cloud computing service that interacts with customer data traffic 415 that enters and exits from the gateway 330. As shown in FIG. 4, security or networking policies 455 implemented by a customer to protect a VNet or subnet are utilized to filter traffic and provide end node control at the VM/VNet/subnet for all network traffic flows.

[0056] The networking policies may block management traffic 440 by preventing management plane access from the NRP 405 which causes service disruptions. The service disruption may be addressed by an fNIC that is attachable and detachable from a parent network interface controller fNIC in an elastic manner as discussed below. The separation of management and data planes may be one characteristic of a software defined network (SDN). Thus, the customer's network 405 may be based in whole or part on SDN technologies, in some implementations, as described below.

[0057] FIG 5 shows an illustrative fNIC 505 that may be utilized to support a variety of scenarios that is attached to a parent NIC 510. The parent NIC may be implemented in this example as virtualization of a network interface at the host supporting the VM 345 using a container model, although physical embodiments may be utilized in some scenarios. An fNIC may be implemented using a compartment of the container as a child network interface configuration. Essentially, a compute instance may be created with a placeholder network interface such that multiple fNICs can be dynamically put up and taken down by respectively being attached and detached from the instance.

[0058] The parent NICs and fNICs provide identity, connectivity, and discoverability for the VMs in the customer's VNet. An fNIC enables flexibility for various VM deployment scenarios by its capabilities for attachment and detachment from the parent NIC. The flexibility enables rapid provisioning of a variety of cloud-computing features and services on an on-demand basis without needing to alter the fundamental workflow in a given VM/VNet/subnet while conforming with applicable networking policies.

[0059] As shown in FIG. 5, the use scenarios illustratively include, for example, cross-subscriptions and multi-VNet homing (indicated by reference numeral 515), multi-tenancy and subnet sharing 520, and pre-provisioning of resources or "hot attach" 525.

[0060] FIG. 6 shows an example virtual filtering platform (VFP) 600 extension to a VM switch 605 that enables data path isolation in the multi-tenant architecture discussed herein by enforcing specific networking policies that are tied to each container 610 that is used to implement a parent NIC to which an fNIC is attached based on the NPI. The VM switch

may logically underly the parent NIC 310 and may provide a port 612 to each VM supported by the NIC. The VFP may divide networking policies applicable to the port into layers that include rules that govern SDN behaviors and characteristics. The virtual filtering platform may provide capabilities to enforce policies and transform or tunnel data packets in a given  
5 computing workload that are entering and leaving the VM 145. The virtual filtering platform may include a central data packet processor (not shown) that performs the processing of data packets.

**[0061]** The networking policy layers may include those, in this example, relating to metering 615, access control lists (ACLs) 620, VNet addressing/routing 625, and other  
10 various SDN functions or features 630 which may include, for example, those pertaining to routing, tunneling, filtering, address translation, encryption, decryption, encapsulation, de-encapsulation, or quality of service (QoS). The packet processor in the VFP 600 may evaluate the packets of data traffic 635 as they traverse the networking policy layers, matching rules in each layer based on a state of the packet after an action is performed in  
15 the preceding layer. Returning packets may traverse the layers in the opposite direction and may be processed by the VFP to match the applicable rules. The rules used to express the networking policies may be entities that perform actions on matching packets (e. g., using a match action table model) as the computing workload is processed by the VFP.

**[0062]** In the illustrated example scenarios, SDN capabilities may be enhanced by  
20 disaggregating policy enforcement from the host and moving it onto SDN appliance strategically placed in the network.

**[0063]** Software defined networking (SDN) is conventionally implemented on a general-purpose compute node. The SDN control plane may program the host to provide core network functions such as security, virtual network, and load balancer policies.

**[0064]** Referring to FIG. 7, illustrated is an example of an SDN appliance 710 that can enable disaggregation according to some embodiments. In some embodiments, the SDN appliance 710 may enable the use of the SDN control plane to manage network devices while providing high availability and fault tolerance, as further described herein. FIG. 7 illustrates one example of a network optimized chassis including SDN agents 720, a network  
25 driver capable of performing network transforms such as a virtual filtering platform (VFP) 730, policies 740, and cards (e.g., FPGAs) 750. The SKU can be change, and hosts 750 may be used as the SDN appliance if needed. The various embodiments described herein show the use of the SDN appliance as a general concept.

**[0065]** As described above, various embodiments include datacenter networking models

that include selectively placed network hops that can apply software defined networking (SDN) policy at various points in a data center (i.e., network traversal point) before data traffic reaches their destination. This can provide improved overall performance, disaggregation from the host, and application of other SDN capabilities before the destination receives data traffic.

**[0066]** In an embodiment, criteria to determine where to implement the SDN policy for a flow can include, for example: age of the flow, rate of the flow, total bytes transferred on the flow, total number of flows in use at the correspondent host, and the like. Since the SDN policy applied to a flow can comprise multiple aspects, different aspects of the policy can be implemented at different locations.

**[0067]** In some implementations that use a rack level switch such as a top-of-rack (ToR) switch, such devices typically do not have the capability to perform transforms. An SDN appliance gateway can be used to host these agents and provide switch functionality, and can further provide transformations and connectivity. The SDN appliance can accept policies that perform transformations. In some embodiments, an agent can be implemented that programs the drivers that run on the SDN appliance. The traffic sent by workloads can be directed through the SDN appliance, which can apply policies and perform transformations on the traffic and send the traffic to the destination. In some configurations, the SDN appliance may include a virtual switch such as a virtual filtering platform.

**[0068]** SDN appliances can become a single point of failure for software defined networks. Mitigation of faults for SDN appliances must take into account the preservation of transient states (for example, TCP flow state) as well as the locality of the state within the individual SDN appliances. If two appliances are cross-wired to two switches, the single point of failure can be avoided from a physical device standpoint but may introduce state management and expected connectivity issues.

**[0069]** The described embodiments may support, for example, connected devices such as FPGAs on SDN appliances in multiple different network and physical topologies.

**[0070]** While FPGAs are used to illustrate the described techniques, it should be understood that the techniques may be applied to other types of connected devices such as a GPU.

**[0071]** The various aspects of the disclosure are described herein with regard to certain examples and embodiments, which are intended to illustrate but not to limit the disclosure. It should be appreciated that the subject matter presented herein may be implemented as a computer process, a computer-controlled apparatus, a computing system, an article of

manufacture, such as a computer-readable storage medium, or a component including hardware logic for implementing functions, such as a field-programmable gate array (FPGA) device, a massively parallel processor array (MPPA) device, a graphics processing unit (GPU), an application-specific integrated circuit (ASIC), a multiprocessor System-on-Chip (MPSoC), etc.

[072] A component may also encompass other ways of leveraging a device to perform a function, such as, for example, a) a case in which at least some tasks are implemented in hard ASIC logic or the like; b) a case in which at least some tasks are implemented in soft (configurable) FPGA logic or the like; c) a case in which at least some tasks run as software on FPGA software processor overlays or the like; d) a case in which at least some tasks run as software on hard ASIC processors or the like, etc., or any combination thereof. A component may represent a homogeneous collection of hardware acceleration devices, such as, for example, FPGA devices. On the other hand, a component may represent a heterogeneous collection of different types of hardware acceleration devices including different types of FPGA devices having different respective processing capabilities and architectures, a mixture of FPGA devices and other types hardware acceleration devices, etc.

[0073] With reference to FIG. 8, illustrated is an example of a data center smart rack 800 with two SDN appliances 810 having one or more smart NICs 820 and a plurality of compute rows 830 having servers. Any virtual machine 840 running on any server in the data center smart rack can utilize the SDN appliance 810. For example, virtual machines with a high connections per second (CPS) or flow scale needs can send flows through the SDN appliance. The appliances may be configured to perform SDN data path functions at a significantly faster rate as compared to conventional methods.

[0074] With reference to FIG. 9, illustrated is an example of implementing the appliance, or its components, in a virtual or distributed fashion to provide a flexible implementation approach. In other words, the components of the SDN appliance may be distributed in servers 910 rather than housed in a separate physical assembly. For example, each server 910 in the data center smart rack 900 may contain at least 2 x SmartNICs 930. In an embodiment, a grouping of the SmartNICs 930 may be logically combined to form a virtual appliance 940.

[0075] With reference to FIG. 10, illustrated is an example of SDN disaggregation where non-compute functionality is removed off the compute host. In one implementation, smartNICs 1010 may be pooled into shared appliances 1000. In an embodiment, the

appliance 1000 may be cost optimized. The appliance 1000 may be configured to perform all SDN data path functions. In this and other figures herein, the dashed line indicates skinny NICs 1040 and the solid line indicates smartNICs 1010. The skinny NICs 1040 may be implemented on the servers 1030 for low function, cost, and low power.

5 [0076] In an embodiment, storage traffic may bypass the appliance. FIG. 10 illustrates an example of a fault tolerant scheme that is resilient to ToR or appliance failure. Each appliance 1000 is cross-connected to each ToR 1020. In an embodiment, two appliances may use connection state replication. In one example, 2-4 SmartNICs 1010 per appliance may be implemented depending on the load. In an embodiment, each ToR 1020 may be fully  
10 connected to every SmartNIC 1010. For example, each SmartNIC 1010 may provide two redundant 100G bump-in-wire SDNs. The illustrated example allows for hot insertion/deletion/RMA of NICs. The servers 1030 may have dual ported skinny NICs 1040. In this embodiment, storage traffic does not transit the SmartNIC (storage bypass). In an embodiment, SDN agents 1050 may execute on the server or ToRs.

15 [0077] With reference to FIG. 11, illustrated is an embodiment of a data center smart rack 1100 where smartNICs 1130 are distributed into smart TORs 1120 (e.g., SONiC-based ToRs). The illustrated example may provide the same data center smart rack appliance functionality with storage bypass. In an embodiment, the smartToR 1120 may include a switch and SmartNICs 1130. In one implementation, SmartToRs 1130 may be  
20 interconnected similar to the data center smart rack shown in FIG. 10 by exposing NIC ports. In this implementation, the rack may be mechanically and board optimized with fewer connectors and boards. In one embodiment, servers 1110 may each have one skinny NIC 1140. In an embodiment, SDN agents 1150 may execute on the server or ToRs.

[0078] With reference to FIG. 12A, illustrated is another embodiment of a data center  
25 smart rack 1200 with distribution of smartNIC's 1240 among the servers 1210 to create a distributed and virtual appliance. This example may be similar to the SDN functionality with storage bypass. This embodiment eliminates a physically separate SDN appliance by adding SmartNICs 1240 inside the servers 1210, creating virtual appliances 1250 with  $N$  number of smartNICs (2 in this example). Most or all of the servers 1210 may further have  
30 one skinny NIC 1230. The ratio may be determined based on performance needs (e.g., a ratio of 3:1). In an embodiment, SDN agents 1260 may execute on the server or ToRs.

[0079] With reference to FIG. 12B, illustrated is another embodiment of the data center smart rack 1200 with distribution of smartNIC's 1280 among the servers 1270 to create a distributed and virtual appliance. In some embodiments, a virtual appliance may be

implemented in a bare metal server scenario. This embodiment also eliminates a separate SDN appliance by adding SmartNICs 1280 inside the servers 1270, where  $N$  servers are each configured with a smartNIC to form a virtual appliance with  $N$  number of smartNICs. In some embodiments, such a configuration may be used to support 5G edge applications.

5 [0080] With reference to FIG. 13, the left side of the figure shows a skinny or standard NIC 1320 that provides RDMA offload functionality 1325 and connectivity to the appliance. Networking traffic may be tunneled to the appliance 1330. The right side of the figure shows stateful network policy-based forwarding and security 1350 performed on SmartNICs 1340.

[0081] Disclosed herein are techniques for high availability (HA) for policy-based flow forwarding. High availability (HA) schemes may be implemented for a rack-based networking appliance as disclosed herein. Failure modes that may be addressed by the HA schemes include losing a smartNIC in the appliance and losing the appliance altogether. The techniques include the process of synchronizing to a new appliance or smartNIC once a failure occurs.

15 [0082] In an embodiment, high availability may be achieved via the following concepts:

- Active-Passive architecture with two (2) SDN appliances
- Overprovisioning – same ENI handled by multiple appliance smartNICs
- Pairing – each sNIC from one SDN appliance will have “paired” card in second SDN appliance, with shared/same VIP

20 [0083] Disclosed herein is a scaling host policy via distribution across multiple SDN appliances. Techniques are described for managing oversubscription of a rack-based networking appliance. The techniques include spreading virtual machines across the appliance so that the capabilities of the appliance can be oversubscribed efficiently.

[0084] In an embodiment, scaling may be achieved via the following concepts:

25

- Overprovisioning – the same ENI handled by multiple appliance sNICs, each sNIC with a different VIP
- Flow splitting – TOR (or source node) splitting traffic equally across different VIPs

[0085] For high availability, one goal may be to have zero downtime planned failover, and <2 sec downtime unplanned failover. Additionally, another goal can be the ability to resume connections in the event of both unplanned and planned failover. The high availability scheme may be implemented such that if the appliance receives a valid packet, the appliance does not drop the packet due to flow replication delays. A further objective

30

can be to ensure that both inbound and outbound packets take the same appliance for a given flow.

5 [0086] FIG. 14 shows an example of a high-level architecture of rack design for high availability and scale. This example illustrates active/active hardware design and active/passive ENI design (software). FIG. 14 illustrates an embodiment with two TORs with a cross-wire design. Each sNIC may be connected to both TORs for availability and scale.

10 [0087] Both SDN appliances may be available in case of a single TOR failure. Some ENIs may be “Active” on SDN Appliance 1, and “Passive” on SDN Appliance 2. Other ENIs are “Active” on SDN Appliance 2 and “Passive” on SDN Appliance 1

[0088] Both SDN Appliances may be used for traffic. In an embodiment, each SDN Appliance may be provisioned up to a set threshold of capacity (to allow for failovers). For example, 80% may be used. If full throughput redundancy is desired, then the appliances may be run at 50%.

15 [0089] The described techniques provide a scheme where a single TOR failure does not impact the overall connection rate.

[0090] FIG. 15 shows an example of traffic flow in a highly available and scalable appliance set. In an embodiment, ENIs from a single VM may be provisioned on multiple sNICs on the same SDN appliance. For example, a sNIC from SDN Appliance 1 may be 20 paired with a sNIC from SDN Appliance 2 (where each sNIC participates only in single pairing relationship). Pairwise flow replication may be provided. A single sNIC may handle multiple ENIs.

[0091] Some ENIs on the same sNIC may execute operate in an active mode, and some other ENIs on the same sNIC may operate in a passive mode.

25 [0092] Each sNIC may have two VIPs: one VIP announced with short path thru BGP (used by ENIs in “Active” mode), and a second VIP announced with longer path thru BGP (used by ENIs in “Passive” mode). Paired sNICs may announce the same set of VIPs.

[0093] The configuration may include ENI-based (not card based) flow replication (direction of flow replication: “Active ENI” → “Passive ENI”).

30 [0094] A single ENI may be programmed on multiple sNICs (each sNIC with a different VIP).

[0095] The TOR (or source node where VM is located) may perform traffic load balancing/splitting/sharding for a selected ENI across VIPs of the sNICs on which this ENI is provisioned.

[0096] High Availability Architecture

[0097] Active-Passive model

[0098] The Active-Passive model may ensure high availability in case either (1) one of the TOR/SDN Appliance fails or otherwise becomes unavailable, or (2) a single sNIC fails  
5 or otherwise becomes unavailable.

[0099] Setup details

[00100] Each sNIC from “SDN Appliance 1” may have a “paired” sNIC from “SDN Appliance 2”

[00101] “Paired” sNICs can serve the same ENI with the exact same policies setup on each

10 [00102] “Paired” sNICs may continuously replicate active flows from the Active sNIC to Passive sNIC

[00103] Both sNICs may announce the same VIP via BGP

[00104] “SDN Appliance 1” may be in Active mode (announcing preferred, shorter path to itself thru BGP)

15 [00105] “SDN Appliance 2” may be in Passive mode (announcing less preferred, longer path to itself through Border Gateway Protocol (BGP))

[00106] Normal traffic pattern

[00107] Normal traffic pattern for ENIs handled by sNICs may go through “SDN Appliance 1” (Active).

20 [00108] Failure mode

[00109] In case of failure, BGP may route from “SDN Appliance 1” (previously active) which may be withdrawn and the TOR may switch to “SDN Appliance 2” and redirect traffic there, ensuring continuous traffic and an uninterrupted customer experience.

[00110] Pairing

25 [00111] sNICs between “SDN Appliance 1” and “SDN Appliance 2” may be paired with each other to create the described “Active-Passive” model.

[00112] In an embodiment, the control plane may be responsible for creating “pairing” relationship between sNICs – select which sNICs create a pair.

[00113] ENI Policy configuration

30 [00114] “Paired” sNICs may be configured (by the control plane) with the same ENI and same policy.

[00115] The control plane may be responsible for configuring the same ENI and same policy across both paired sNICs. Replication of ENI policy may not be required by the sNIC as it is handled by the control plane.

[00116] Flow replication

[00117] Once the “pairing” relationship is established, flows may start being replicated and synced. sNICs may be responsible for replicating and syncing flows across sNICs in the “paired” relationship.

5 [00118] In case of outage (e.g., the entire SDN appliance is not available for a longer period of time), the Pairing” relationship might be changed by the control plane.

[00119] In this case the control plane can:

- Withdraw the “pairing” relationship from sNIC (unpair the sNIC)
- Establish a new “pairing” relationship to another sNIC (pair different sNIC)

10 [00120] The original sNIC (which is currently in the active state) may continue to receive traffic when the pairing relationship changes.

[00121] Once the new pairing is established, the flow transfer/sync can start. The new sNIC may become “passive” from the point of view of traffic and to ensure that no outage happens, the new sNIC should not become “active” until all flows are fully synced with the  
15 original sNIC.

[00122] Overprovisioning

[00123] The same ENI may be handled by multiple sNICs

[00124] Scalability

[00125] Overprovisioning and flow splitting can allow for high CPS as well as high  
20 bandwidth for customers, as all the connections will be distributed across multiple sNICs. In various embodiments, a different number of sNICs may be provisioned depending on customer needs for scale.

[00126] Overprovisioning

[00127] A single ENI may be provisioned on multiple sNICs in a single SDN appliance.  
25 The same policy (with exemption of sNIC VIP which will be different) may be setup on each sNIC.

[00128] Each sNIC may announce a different VIP.

[00129] The TOR (or source side) may be configured for splitting (spreading) traffic going through SDN appliances across multiple VIPs to ensure traffic is equally distributed across  
30 all the overprovisioned sNICs.

[00130] In addition, for the purpose of high availability (as described above), the same ENI may also be setup on “paired” sNICs on the secondary SDN appliance.

[00131] Flow splitting

**[00132]** The objective of flow splitting is to ensure that ECMP or any other mechanism will ensure that any set of flows that were active and synced actually end up on the passive node, and to avoid ECMP or other mechanisms from landing a different set of flows that are already synced.

5 **[00133]** Flow splitting may be performed by an intelligent TOR, directly on a source based on stable hashing, or directly on the source node (where the VM is).

**[00134]** As a single ENI will be handled by multiple VIPs (overprovisioned) – ex. 23.0.0.1, 23.0.0.2, 23.0.0.3, the TOR may equally rewrite the destination address to ensure a similar outcome as in the “ECMP” protocol (with additional explicit destination address rewrite).

10 **[00135]** Scenarios

**[00136]** Single TOR Failure - Single TOR becomes inaccessible

**[00137]** How traffic pattern changes

**[00138]** The SDN appliance behind this TOR is still accessible through the second TOR. There may be a loss of 50% bandwidth, and no loss of CPS. The second TOR must now  
15 handle double the bandwidth and double the CPS. Assuming the sNIC is actually the bottleneck for CPS (not TOR), there is no CPS loss, the only impact is on bandwidth.

**[00139]** The TOR becomes the bottleneck for bandwidth and CPS

**[00140]** Both SDN Appliances are still operating at 100% capacity

**[00141]** By splitting the load across multiple sNICs, there is only a loss of 50% of the  
20 connections from the sNIC that failed, and not the entire load of the VM.

**[00142]** Single link failure - Single link between TOR and single sNIC becomes unavailable

**[00143]** How traffic pattern changes

**[00144]** The sNIC is still being served by the link to the second TOR.

25 **[00145]** There may be a loss of 50% bandwidth to that sNIC as a single connection is used for both ingress and egress

**[00146]** No impact on CPS

**[00147]** By splitting the load across multiple sNICs there is only a loss of 50% of the connections from the sNIC that failed, and not the entire load of the VM.

30 **[00148]** Single sNIC failure

- Single sNIC becomes unavailable
- This sNIC was set up as “active” for some ENIs and “passive” for other ENIs
- This sNIC no longer serves traffic for “active” ENIs

[00149] How traffic pattern changes

[00150] This sNIC no longer announces its own VIPs via BGP

[00151] Longer route for same VIP is used by TORs

5 [00152] “Paired” sNIC becomes “active” for all ENIs (it was already “active” for some ENIs, and “passive” for other ENIs, now the “passive” ENIs are becoming “active”)

[00153] ENIs served by this sNIC reduces utilization from 80% each -> 50% each (loss of 3/8th or 37.5% capacity per sNIC). This assumes that the original sNIC was allocated up to only 80% (to allow for failover). This number can be adjusted.

10 [00154] Considering that a single ENI is load balanced across multiple sNICs, other sNICs are not affected and the actual capacity reduction (bandwidth + CPS) is much lower than 30%. Assuming 5 sNICs are allocated per ENI, loss of a single sNIC reduces capacity by 7.5%.

[00155] Previously: 100%, 100%, 100%, 100%, 100% = 100% capacity

[00156] Now: 62.5%, 100%, 100%, 100%, 100% = 92.5% capacity

15 [00157] What happens next

[00158] If the outage of the single sNIC persists, the control plane may select a different sNIC on another SDN Appliance (SDN Appliance 3) and initiate pairing with that sNIC. A backup and “empty” SDN Appliance may be provided to handle this scenario.

20 [00159] Flow sync may start between sNIC currently processing traffic (half of “active” ENIs) and a new sNIC may be added to the pairing relationship. Once flow sync completes for all ENIs, the newly paired sNIC may start to announce its VIP as “active” (with a shorter BGP path).

[00160] Traffic switches again to the new sNIC, as it is primary.

[00161] Capacity is fully restored.

25 [00162] Single SDN Appliance failure (all sNICs on that appliance) - All sNICs on that appliance become inaccessible

[00163] How traffic pattern changes

- Each sNIC was paired with other sNIC from the second SDN Appliance
- The second SDN appliance becomes active for all ENIs
- 30 • Capacity is reduced (both bandwidth and CPS)

[00164] What happens next

[00165] If the outage persists, the control plane removes the existing pairing relationship, allocates a new SDN appliance, and creates a pairing relationship with that new SDN

appliance.

**[00166]** Flow Replication using Perfect Sync

**[00167]** Perfect Sync provides consistent synchronization of flows between paired sNICs as those sNICs are active and receiving new connections is important.

5 **[00168]** The algorithm below allows for consistent state replication between a pair of the sNICs during the sync process after a pairing relationship is established, re-established, or recovered.

**[00169]** Prerequisites

10 **[00170]** The following examples use colors to represent connections/flows when a pairing fails. However, it should be understood that other methods can be used, such as using time stamps or another indication that can be used to associate flows with pairings such as flags, bit patterns, and so forth.

**[00171]** There are at least 2 colors (suggested: 8 colors represented by 3 bits)

**[00172]** All connections/entries in the flow table are colored

15 **[00173]** A pairing relationship is established between two sNICs (primary sNIC and secondary sNIC)

**[00174]** There exists a way to replicate connection (entry in a flow table) to a paired device.

**[00175]** Algorithm:

- When a sNIC comes online, it chooses a color (“Color A”)
- 20 • As the sNIC starts receiving traffic it creates new flows (creates entries in flow table)
- New flows are added to the flow table with the chosen color (“Color A”) and immediately replicated to the paired device (colors do not get replicated)
- All connections/flows use the same color until the pairing fails, and continues
- 25 using the color if there is no active pairing
- When pairing is re-established, the device “changes” color to new color (“Color B”)
- New flows are added to the flow table using the new color (“Color B”) and immediately replicated to paired device (colors do not get replicated).
- 30 • Device starts replicating (synchronizing) existing connections (the new device to which either a new pairing was created or the existing device to which pairing was re-established)
- Synchronization method replicates only connections (entries in the flow table)

which have a color that is different the actively used color (different than “Color B”)

- Entries with the currently active color (“Color B”) are not replicated via the synchronization algorithm, as they are replicated in real-time (immediately as they are created).

5       • Algorithm finishes. Flow replication has been completed.

**[00176]** The above algorithm ensures two synchronizations are occurring in parallel:

- Real-time immediate synchronization of new changes (which will remain active after “Perfect Sync” is completed)

- Synchronization of existing connections (occurring during “Perfect Sync”)

10 **[00177]** As the synchronization is occurring, new connections as well as changes in the state of existing connections (irrespective of color) are replicated immediately in real-time (outside of the synchronization algorithm).

**[00178]** It is possible that a connection will end (FIN) and will result in the primary sNIC removing a flow and immediately sending the new connection state change to close the connection to a paired device (before the existing connection was even journaled/synchronized to the paired device). To deal with this possibility on the paired device side, if the connection does not already exist in its table, then this update message should be ignored.

15 **[00179]** When pairing is re-established, the secondary sNIC may empty the entire state of the flow table to allow it to receive a clean state.

**[0180]** Turning now to FIG. 16, illustrated is an example operational procedure for processing data packets and implementing policies in a software defined network (SDN) of a virtual computing environment, by at least one SDN appliance configured to disaggregate enforcement of policies of the SDN from hosts of the virtual computing environment. In an embodiment, the servers are communicatively coupled to network interfaces of the SDN appliance. In an embodiment, the servers host a plurality of virtual machines. In an embodiment, the SDN appliance comprises a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance. Such an operational procedure can be provided by one or more components illustrated in FIGs. 1 through 15.

25

30 The operational procedure may be implemented in a system comprising one or more computing devices. It should be understood by those of ordinary skill in the art that the operations of the methods disclosed herein are not necessarily presented in any particular order and that performance of some or all of the operations in an alternative order(s) is

possible and is contemplated. The operations have been presented in the demonstrated order for ease of description and illustration. Operations may be added, omitted, performed together, and/or performed simultaneously, without departing from the scope of the appended claims.

5 **[0181]** It should also be understood that the illustrated methods can end at any time and need not be performed in their entirety. Some or all operations of the methods, and/or substantially equivalent operations, can be performed by execution of computer-readable instructions included on a computer-storage media, as defined herein. The term “computer-readable instructions,” and variants thereof, as used in the description and claims, is used  
10 expansively herein to include routines, applications, application modules, program modules, programs, components, data structures, algorithms, and the like. Computer-readable instructions can be implemented on various system configurations, including single-processor or multiprocessor systems, minicomputers, mainframe computers, personal computers, hand-held computing devices, microprocessor-based, programmable consumer  
15 electronics, combinations thereof, and the like.

**[0182]** It should be appreciated that the logical operations described herein are implemented (1) as a sequence of computer implemented acts or program modules running on a computing system such as those described herein) and/or (2) as interconnected machine logic circuits or circuit modules within the computing system. The implementation is a  
20 matter of choice dependent on the performance and other requirements of the computing system. Accordingly, the logical operations may be implemented in software, in firmware, in special purpose digital logic, and any combination thereof. Thus, although the routine 1600 is described as running on a system, it can be appreciated that the routine 1600 and other operations described herein can be executed on an individual computing device or  
25 several devices.

**[0183]** Referring to FIG. 16, operation 1601 illustrates receiving, at the SDN appliance from a device that is remote from the virtual computing environment, a data packet addressed to an endpoint in a virtual network hosted by one of the virtual machines. In an embodiment, the data packet comprises an identifier indicative of the remote device.

30 **[0184]** Operation 1601 may be followed by operation 1603. Operation 1603 illustrates based on the identifier, determining that the data packet is associated with the virtual network.

**[0185]** Operation 1603 may be followed by operation 1605. Operation 1605 illustrates based on the determined association, mapping one of a plurality of policies to a data flow

of the virtual network.

[0186] Operation 1605 may be followed by operation 1607. Operation 1607 illustrates modifying, by the SDN appliance, the packet in accordance with the mapped policy; wherein the mapped policy is dynamically adjustable based on the data flow.

5 [0187] Operation 1607 may be followed by operation 1609. Operation 1609 illustrates forwarding the modified packet to the endpoint in the virtual network.

[0188] FIG. 17 illustrates a general-purpose computing device 1700. In the illustrated embodiment, computing device 1700 includes one or more processors 1710a, 1710b, and/or 1710n (which may be referred herein singularly as "a processor 1710" or in the plural as  
10 "the processors 1710") coupled to a system memory 1717 via an input/output (I/O) interface 1730. Computing device 1700 further includes a network interface 1740 coupled to I/O interface 1730.

[0189] In various embodiments, computing device 1700 may be a uniprocessor system including one processor 1710 or a multiprocessor system including several processors 1710  
15 (e.g., two, four, eight, or another suitable number). Processors 1710 may be any suitable processors capable of executing instructions. For example, in various embodiments, processors 1710 may be general-purpose or embedded processors implementing any of a variety of instruction set architectures (ISAs), such as the x1717, PowerPC, SPARC, or MIPS ISAs, or any other suitable ISA. In multiprocessor systems, each of processors 1710  
20 may commonly, but not necessarily, implement the same ISA.

[0190] System memory 1717 may be configured to store instructions and data accessible by processor(s) 1710. In various embodiments, system memory 1717 may be implemented using any suitable memory technology, such as static random access memory (SRAM), synchronous dynamic RAM (SDRAM), nonvolatile/Flash-type memory, or any other type  
25 of memory. In the illustrated embodiment, program instructions and data implementing one or more desired functions, such as those methods, techniques and data described above, are shown stored within system memory 1720 as code 1725 and data 17217.

[0191] In one embodiment, I/O interface 1730 may be configured to coordinate I/O traffic between the processor 1710, system memory 1717, and any peripheral devices in the device,  
30 including network interface 1740 or other peripheral interfaces. In some embodiments, I/O interface 1730 may perform any necessary protocol, timing, or other data transformations to convert data signals from one component (e.g., system memory 1720) into a format suitable for use by another component (e.g., processor 1710). In some embodiments, I/O interface 1730 may include support for devices attached through various types of peripheral

buses, such as a variant of the Peripheral Component Interconnect (PCI) bus standard or the Universal Serial Bus (USB) standard, for example. In some embodiments, the function of I/O interface 1730 may be split into two or more separate components. Also, in some embodiments some or all of the functionality of I/O interface 1730, such as an interface to system memory 1720, may be incorporated directly into processor 1710.

**[0192]** Network interface 1740 may be configured to allow data to be exchanged between computing device 1700 and other device or devices 17170 attached to a network or network(s) 17170, such as other computer systems or devices as illustrated in FIGS. 1 through 5, for example. In various embodiments, network interface 1740 may support communication via any suitable wired or wireless general data networks, such as types of Ethernet networks, for example. Additionally, network interface 1740 may support communication via telecommunications/telephony networks such as analog voice networks or digital fiber communications networks, via storage area networks such as Fibre Channel SANs or via any other suitable type of network and/or protocol.

**[0193]** In some embodiments, system memory 1720 may be one embodiment of a computer-accessible medium configured to store program instructions and data as described above for the Figures for implementing embodiments of the corresponding methods and apparatus. However, in other embodiments, program instructions and/or data may be received, sent or stored upon different types of computer-accessible media. A computer-accessible medium may include non-transitory storage media or memory media, such as magnetic or optical media, e.g., disk or DVD/CD coupled to computing device 1700 via I/O interface 1730. A non-transitory computer-accessible storage medium may also include any volatile or non-volatile media, such as RAM (e.g. SDRAM, DDR SDRAM, RDRAM, SRAM, etc.), ROM, etc., that may be included in some embodiments of computing device 1700 as system memory 1720 or another type of memory. Further, a computer-accessible medium may include transmission media or signals such as electrical, electromagnetic or digital signals, conveyed via a communication medium such as a network and/or a wireless link, such as may be implemented via network interface 1740. Portions or all of multiple computing devices, such as those illustrated in FIG. 17, may be used to implement the described functionality in various embodiments; for example, software components running on a variety of different devices and servers may collaborate to provide the functionality. In some embodiments, portions of the described functionality may be implemented using storage devices, network devices, or special-purpose computer systems, in addition to or instead of being implemented using general-purpose computer systems. The term

"computing device," as used herein, refers to at least all these types of devices and is not limited to these types of devices.

5 [0194] Various storage devices and their associated computer-readable media provide non-volatile storage for the computing devices described herein. Computer-readable media as discussed herein may refer to a mass storage device, such as a solid-state drive, a hard disk or CD-ROM drive. However, it should be appreciated by those skilled in the art that computer-readable media can be any available computer storage media that can be accessed by a computing device.

10 [0195] By way of example, and not limitation, computer storage media may include volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules or other data. For example, computer media includes, but is not limited to, RAM, ROM, EPROM, EEPROM, flash memory or other solid state memory technology, CD-ROM, digital versatile disks ("DVD"), HD-DVD, BLU-RAY, or other  
15 optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computing devices discussed herein. For purposes of the claims, the phrase "computer storage medium," "computer-readable storage medium" and variations thereof, does not include waves, signals, and/or other transitory and/or intangible  
20 communication media, per se.

[0196] Encoding the software modules presented herein also may transform the physical structure of the computer-readable media presented herein. The specific transformation of physical structure may depend on various factors, in different implementations of this description. Examples of such factors may include, but are not limited to, the technology  
25 used to implement the computer-readable media, whether the computer-readable media is characterized as primary or secondary storage, and the like. For example, if the computer-readable media is implemented as semiconductor-based memory, the software disclosed herein may be encoded on the computer-readable media by transforming the physical state of the semiconductor memory. For example, the software may transform the state of  
30 transistors, capacitors, or other discrete circuit elements constituting the semiconductor memory. The software also may transform the physical state of such components in order to store data thereupon.

[0197] As another example, the computer-readable media disclosed herein may be implemented using magnetic or optical technology. In such implementations, the software

presented herein may transform the physical state of magnetic or optical media, when the software is encoded therein. These transformations may include altering the magnetic characteristics of particular locations within given magnetic media. These transformations also may include altering the physical features or characteristics of particular locations within given optical media, to change the optical characteristics of those locations. Other transformations of physical media are possible without departing from the scope and spirit of the present description, with the foregoing examples provided only to facilitate this discussion.

**[0198]** In light of the above, it should be appreciated that many types of physical transformations take place in the disclosed computing devices in order to store and execute the software components and/or functionality presented herein. It is also contemplated that the disclosed computing devices may not include all of the illustrated components shown in FIG. 17, may include other components that are not explicitly shown in FIG. 17, or may utilize an architecture completely different than that shown in FIG. 17.

**[0199]** Although the various configurations have been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended representations is not necessarily limited to the specific features or acts described. Rather, the specific features and acts are disclosed as example forms of implementing the claimed subject matter.

**[0200]** Conditional language used herein, such as, among others, "can," "could," "might," "may," "e.g.," and the like, unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements, and/or steps. Thus, such conditional language is not generally intended to imply that features, elements, and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without author input or prompting, whether these features, elements, and/or steps are included or are to be performed in any particular embodiment. The terms "comprising," "including," "having," and the like are synonymous and are used inclusively, in an open-ended fashion, and do not exclude additional elements, features, acts, operations, and so forth. Also, the term "or" is used in its inclusive sense (and not in its exclusive sense) so that when used, for example, to connect a list of elements, the term "or" means one, some, or all of the elements in the list.

**[0201]** While certain example embodiments have been described, these embodiments have been presented by way of example only, and are not intended to limit the scope of the

inventions disclosed herein. Thus, nothing in the foregoing description is intended to imply that any particular feature, characteristic, step, module, or block is necessary or indispensable. Indeed, the novel methods and systems described herein may be embodied in a variety of other forms; furthermore, various omissions, substitutions and changes in the form of the methods and systems described herein may be made without departing from the spirit of the inventions disclosed herein. The accompanying claims and their equivalents are intended to cover such forms or modifications as would fall within the scope and spirit of certain of the inventions disclosed herein.

**[0202]** It should be appreciated any reference to “first,” “second,” etc. items and/or abstract concepts within the description is not intended to and should not be construed to necessarily correspond to any reference of “first,” “second,” etc. elements of the claims. In particular, within this Summary and/or the following Detailed Description, items and/or abstract concepts such as, for example, individual computing devices and/or operational states of the computing cluster may be distinguished by numerical designations without such designations corresponding to the claims or even other paragraphs of the Summary and/or Detailed Description. For example, any designation of a “first operational state” and “second operational state” of the computing cluster within a paragraph of this disclosure is used solely to distinguish two different operational states of the computing cluster within that specific paragraph – not any other paragraph and particularly not the claims.

**[0203]** Although the various techniques have been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended representations is not necessarily limited to the specific features or acts described. Rather, the specific features and acts are disclosed as example forms of implementing the claimed subject matter. The disclosure presented herein also encompasses the subject matter set forth in the following clauses:

**[0205]** Clause 1: A method for processing data packets and implementing policies in a software defined network (SDN) of a virtual computing environment, by at least one SDN appliance configured to disaggregate enforcement of policies of the SDN from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance, the method comprising:

**[0206]** receiving, at the SDN appliance from a device that is remote from the virtual

computing environment, a data packet addressed to an endpoint in a virtual network hosted by one of the virtual machines, the data packet comprising an identifier indicative of the remote device;

[0207] based on the identifier:

5 [0208] determining that the data packet is associated with the virtual network; and

[0209] based on the determined association, mapping one of a plurality of policies to a data flow of the virtual network;

[0210] modifying, by the SDN appliance, the packet in accordance with the mapped policy; wherein the mapped policy is dynamically adjustable based on the data flow; and

10 [0211] forwarding the modified packet to the endpoint in the virtual network.

[0212] Clause 2: The method of clause 1, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among the servers.

[0213] Clause 3: The method of any of clauses 1-2, wherein the SDN appliance is implemented as two physical devices that are communicatively coupled to at least two top-  
15 of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the physical devices.

[0214] Clause 4: The method of any of clauses 1-3, wherein storage traffic bypasses the sNICs.

20 [0215] Clause 5: The method of any of clauses 1-4, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among one or more top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the top-of-rack switches.

25 [0216] Clause 6: The method of any of clauses 1-5, wherein at least two top-of-rack switches are configured with SDN agents configured to manage functionality of the SDN appliance.

[0217] Clause 7: A system comprising:

[0218] a plurality of servers communicatively coupled to at least one software defined  
30 network (SDN) appliance configured to disaggregate enforcement of policies of a SDN of a virtual computing environment from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality

of the SDN appliance;

[0219] the system configured to:

[0220] receive, from a device that is remote from the virtual computing environment, a data packet addressed to an endpoint in a virtual network hosted by one of the virtual machines, the data packet comprising an identifier indicative of the remote device;

[0221] based on the identifier:

[0222] determining that the data packet is associated with the virtual network; and

[0223] based on the determined association, mapping one of a plurality of policies to a data flow of the virtual network;

10 [0224] modifying, by the SDN appliance, the packet in accordance with the mapped policy; wherein the mapped policy is dynamically adjustable based on the data flow; and

[0225] forwarding the modified packet to the endpoint in the virtual network.

[0226] Clause 8: The system of clause 7, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among the plurality of servers.

15 [0227] Clause 9: The system of any of clauses 7 and 8, wherein the SDN appliance is implemented as two physical devices that are communicatively coupled to at least two top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the physical devices.

20 [0228] Clause 10: The system of any clauses 7-9, wherein storage traffic bypasses the sNICs.

[0229] Clause 11: The system of any clauses 7-10, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among one or more top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the top-of-rack switches.

25 [0230] Clause 12: The system of any clauses 7-11, wherein at least two top-of-rack switches are configured with SDN agents configured to manage functionality of the SDN appliance.

30 [0231] Clause 13: The system of any clauses 7-12, wherein the policy is applied at selectively placed network hops in the virtual network.

[0232] Clause 14: The system of any clauses 7-13, wherein the SDN appliance is configured to apply policies of the virtual computing environment to data traffic on the

virtual network after the data traffic leaves its source and before the data traffic reaches its destination.

**[0233]** Clause 15: A data center rack comprising:

**[0234]** a plurality of servers; the servers communicatively coupled to at least one software defined network (SDN) appliance configured to disaggregate enforcement of policies of a SDN of a virtual computing environment from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance;

**[0235]** the data center rack configured to:

**[0236]** receive, from a device that is remote from a software defined network (SDN) of a virtual computing environment, a data packet addressed to an endpoint in a virtual network hosted by one of the virtual machines, the data packet comprising an identifier indicative of the remote device;

**[0237]** based on the identifier:

**[0238]** determining that the data packet is associated with the virtual network; and

**[0239]** based on the determined association, mapping one of a plurality of policies to a data flow of the virtual network;

**[0240]** modifying, by the SDN appliance, the packet in accordance with the mapped policy; wherein the mapped policy is dynamically adjustable based on the data flow; and

**[0241]** forwarding the modified packet to the endpoint in the virtual network.

**[0242]** Clause 16: The computing environment of clause 15, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among the plurality of servers.

**[0243]** Clause 17: The computing environment of any of clauses 15 and 16, wherein the SDN appliance is implemented as two physical devices that are communicatively coupled to at least two top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the physical devices.

**[0244]** Clause 18: The computing environment of any of the clauses 15-17, wherein storage traffic bypasses the sNICs.

**[0245]** Clause 19: The computing environment of any of the clauses 15-18, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically

distributed among one or more top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the top-of-rack switches.

5 [0246] Clause 20: The computing environment of any of the clauses 15-19, wherein at least two top-of-rack switches are configured with SDN agents configured to manage functionality of the SDN appliance.

[0247] The disclosure presented herein also encompasses the subject matter set forth in the following clauses:

10 [0248] Clause 1: A method for processing data packets in a virtualized computing environment by a hardware-based networking device configured to disaggregate processing of data packets from hosts of the virtualized computing environment, the hardware-based networking device comprising a hardware-based component implementing a plurality of behavioral models indicative of packet processing graphs for data flows in the virtualized computing environment, the method comprising:

15 [0249] receiving, by the hardware-based networking device, a data packet having a source from or destination to an endpoint in a virtual network of the virtualized computing environment;

[0250] determining that the data packet is a first packet of a data flow to or from the endpoint;

20 [0251] based on the determination, mapping one of the behavioral models to the data flow;

[0252] modifying the packet in accordance with the mapped behavioral model, wherein the mapped behavioral model is processed in the hardware-based component;

[0253] storing a state of the data flow; and

[0254] processing subsequent data packets of the data flow based on the stored state.

25 [0255] Clause 2: The method of clause 1, further comprising:

[0256] determining that a received subsequent data packet is a last packet of the data flow; and

[0257] in response to the determination, removing the state of the data flow.

30 [0258] Clause 3: The method of any of clauses 1-2, wherein the plurality of behavioral models comprises a threshold number or percentage of data flows of the virtualized computing environment.

[0259] Clause 4: The method of any of clauses 1-3, wherein the hardware-based component is one of a flexible hardware engine, specialized processor, or programmable circuit.

[0260] Clause 5: The method of any of clauses 1-4, wherein the hardware-based component comprises an inbound packet processing pipeline and outbound packet processing pipeline.

5 [0261] Clause 6: The method of any of clauses 1-5, wherein the flows are stored in the hardware-based component in a GFT.

[0262] Clause 7: The method of clauses 1-6, wherein the hardware-based component comprises one or more of a NAT layer, VNET layer, a decapsulation layer, or transposition layer.

10 [0263] Clause 8: The method of any of clauses 1-7, wherein the inbound packet processing pipeline and outbound packet processing pipeline comprise an ACL pipeline.

[0264] Clause 9: A hardware-based networking device configured to disaggregate processing of data packets from hosts of a virtualized computing environment, the hardware-based networking device comprising a hardware-based component implementing a plurality of behavioral models indicative of packet processing graphs for data flows in the  
15 virtualized computing environment, the hardware-based networking device configured to:

[0265] receive a data packet having a source from or destination to an endpoint in a virtual network of a virtualized computing environment;

[0266] determine that the data packet is a first packet of a data flow to or from the endpoint;

20 [0267] based on the determination, map one of the behavioral models to the data flow;

[0268] modify the packet in accordance with the mapped behavioral model, wherein the mapped behavioral model is processed in the hardware-based component;

[0269] store a state of the data flow; and

[0270] process subsequent data packets of the data flow based on the stored state.

25 [0271] Clause 10: The hardware-based networking device of clause 9, wherein the hardware-based networking device is further configured to perform packet transformation for VNET to VNET data traffic.

[0272] Clause 11: The hardware-based networking device of any of clauses 9 and 10, wherein the hardware-based networking device is further configured to perform packet  
30 transformation for load balancer traffic.

[0273] Clause 12: The hardware-based networking device of any clauses 9-11, wherein the hardware-based networking device is further configured to perform packet transformation for private link traffic.

[0274] Clause 13: The hardware-based networking device of any clauses 9-12, wherein

the hardware-based networking device is further configured to perform packet transformation for service tunneling.

5 [0275] Clause 14: The hardware-based networking device of any clauses 9-13, wherein the hardware-based networking device is further configured to perform packet transformation for one or more of VNET encryption, telemetry, BGP, debugging, or flow replication.

[0276] Clause 15: The hardware-based networking device of any clauses 9-14, wherein the hardware-based networking device is further configured to:

10 [0277] determine that a received subsequent data packet is a last packet of the data flow; and

[0278] in response to the determination, remove the state of the data flow.

[0279] Clause 16: A computing environment comprising a plurality of computing devices and one or more hardware-based networking devices configured to disaggregate processing of data packets from the plurality of computing devices, the hardware-based networking device comprising a hardware-based component implementing a plurality of behavioral models indicative of packet processing graphs for data flows in the computing environment, the hardware-based networking device configured to:

15 [0280] receive a data packet having a source from or destination to an endpoint in a virtual network of a virtualized computing environment;

20 [0281] determine that the data packet is a first packet of a data flow to or from the endpoint;

[0282] based on the determination, map one of the behavioral models to the data flow;

[0283] modify the packet in accordance with the mapped behavioral model, wherein the mapped behavioral model is processed in the hardware-based component;

25 [0284] store a state of the data flow; and

[0285] process subsequent data packets of the data flow based on the stored state.

[0286] Clause 17: The computing environment of clause 16, wherein the plurality of behavioral models comprises a threshold number or percentage of data flows of the virtualized computing environment.

30 [0287] Clause 18: The computing environment of any of clauses 16 and 17, wherein the hardware-based component is one of a flexible hardware engine, specialized processor, or programmable circuit.

[0288] Clause 19: The computing environment of any of the clauses 16-18, wherein the hardware-based component comprises an inbound packet processing pipeline and outbound

packet processing pipeline.

[0289] Clause 20: The computing environment of any of the clauses 16-19, wherein the flows are stored in the hardware-based component in a GFT.

[0290] The disclosure presented herein also encompasses the subject matter set forth in the following clauses:

[0291] Clause 1: A method for processing data packets and implementing policies in a software defined network (SDN) of a virtual computing environment, by at least two SDN appliances configured to disaggregate enforcement of policies of the SDN from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the servers communicatively coupled to network interfaces of at least two top-of-rack switches (ToRs), the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance, the sNICs having a floating network interface configured to provide a virtual port connection to an endpoint within a virtual network of the virtual computing environment, wherein:

[0292] each sNIC is associated with the first SDN appliance is paired with an sNIC associated with the second SDN appliance;

[0293] each paired sNIC is configured to serve a same floating network interface;

[0294] the method comprising:

[0295] selecting one of the SDN appliances as an active SDN appliance, wherein the other SDN appliance is a passive SDN appliance;

[0296] announcing a same virtual IP (VIP) via border gateway protocol (BGP) for each paired sNIC;

[0297] for each paired sNIC, replicating active flows from an active sNIC to a passive sNIC; and

[0298] in response to a failure of the active SDN appliance:

[0299] withdrawing routes from the active SDN appliance; and

[0300] causing the ToRs to redirect data traffic to the passive SDN appliance.

[0301] Clause 2: The method of clause 1, further comprising:

[0302] receiving, from a device that is remote from the virtual computing environment, a data packet addressed to an endpoint in a virtual network hosted by one of the virtual machines, the data packet comprising an identifier indicative of the remote device;

[0303] based on the identifier:

[0304] determining that the data packet is associated with the virtual network; and

[0305] based on the determined association, mapping one of a plurality of policies to a data flow of the virtual network;

[0306] modifying the packet in accordance with the mapped policy; wherein the mapped policy is dynamically adjustable based on the data flow; and

5 [0307] forwarding the modified packet to the endpoint in the virtual network.

[0308] Clause 3: The method of any of clauses 1-2, wherein the VIP is announced via short path using BGP, further comprising announcing a second VIP announced with a long path using BGP, wherein the second VIP is used by the floating network interfaces when in passive mode.

10 [0309] Clause 4: The method of any of clauses 1-3, wherein paired sNICs announce the same set of VIPs.

[0310] Clause 5: The method of any of clauses 1-4, wherein flow replication is based on the floating network interfaces, and the direction of flow replication is active to passive.

15 [0311] Clause 6: The method of any of clauses 1-5, wherein a single floating network interface is programmed on multiple sNICs.

[0312] Clause 7: A system comprising:

[0313] a plurality of servers communicatively coupled to at least two software defined network (SDN) appliances configured to disaggregate enforcement of policies of a SDN of a virtual computing environment from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the servers communicatively coupled to network interfaces of at least two top-of-rack switches (ToRs), the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance, the sNICs having a floating network interface configured to provide a virtual port connection to an endpoint within a virtual network of the virtual computing environment,

[0314] wherein:

[0315] each sNIC is associated with the first SDN appliance is paired with an sNIC associated with the second SDN appliance;

30 [0316] each paired sNIC is configured to serve a same floating network interface;

[0317] the system configured to:

[0318] select one of the SDN appliances as an active SDN appliance, wherein the other SDN appliance is a passive SDN appliance;

[0319] announce a same virtual IP (VIP) via border gateway protocol (BGP) for each

paired sNIC;

[0320] for each paired sNIC, replicate active flows from an active sNIC to a passive sNIC;  
and

5 [0321] in response to a failure of the active SDN appliance, withdraw routes from the  
active SDN appliance and cause the ToRs to redirect data traffic to the passive SDN  
appliance.

[0322] Clause 8: The system of clause 7, wherein the ToRs are configured to perform  
traffic load balancing, splitting, or sharding for selected floating network interfaces across  
VIPs of cards on which the selected floating network interface is provisioned.

10 [0323] Clause 9: The system of any of clauses 7 and 8, wherein a control plane of the  
virtual computing environment is configured to create pairing relationships between sNICs,  
wherein the paired sNICs are configured with a same ENI and policy.

[0324] Clause 10: The system of any clauses 7-9, wherein in response to a failure of an  
SDN appliance not available for a threshold period of time:

15 [0325] withdrawing the pairing relationship from an associated sNIC; and

[0326] establishing a new pairing to another sNIC.

[0327] Clause 11: The system of any clauses 7-10, wherein the SDN appliance is  
implemented as a distributed appliance where the sNICs are physically distributed among  
one or more top-of-rack switches, the servers communicatively coupled to network  
20 interfaces of the top-of-rack switches so that each of the servers have a switchable  
communications path to each sNIC of the top-of-rack switches.

[0328] Clause 12: The system of any clauses 7-11, wherein the VIP is announced via short  
path using BGP, further comprising announcing a second VIP announced with a long path  
using BGP, wherein the second VIP is used by the floating network interfaces when in  
25 passive mode.

[0329] Clause 13: The system of any clauses 7-12, wherein paired sNICs announce the  
same set of VIPs.

[0330] Clause 14: The system of any clauses 7-13, wherein flow replication is based on  
the floating network interfaces, and the direction of flow replication is active to passive.

30 [0331] Clause 15: A data center rack comprising:

[0332] a plurality of servers communicatively coupled to at least two software defined  
network (SDN) appliances configured to disaggregate enforcement of policies of a SDN of  
a virtual computing environment from hosts of the virtual computing environment, the hosts  
implemented on servers communicatively coupled to network interfaces of the SDN

appliance, the servers hosting a plurality of virtual machines, the servers communicatively coupled to network interfaces of at least two top-of-rack switches (ToRs), the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance, the sNICs having a floating network interface configured to provide a virtual port connection to an endpoint within a virtual network of the virtual computing environment,

5 [0333] wherein:

[0334] each sNIC is associated with the first SDN appliance is paired with an sNIC associated with the second SDN appliance;

10 [0335] each paired sNIC is configured to serve a same floating network interface;

[0336] the data center rack configured to:

[0337] select one of the SDN appliances as an active SDN appliance, wherein the other SDN appliance is a passive SDN appliance;

15 [0338] announce a same virtual IP (VIP) via border gateway protocol (BGP) for each paired sNIC;

[0339] for each paired sNIC, replicate active flows from an active sNIC to a passive sNIC; and

20 [0340] in response to a failure of the active SDN appliance, withdraw routes from the active SDN appliance and cause the ToRs to redirect data traffic to the passive SDN appliance.

[0341] Clause 16: The computing environment of clause 15, wherein the VIP is announced via short path using BGP, further comprising announcing a second VIP announced with a long path using BGP, wherein the second VIP is used by the floating network interfaces when in passive mode.

25 [0342] Clause 17: The computing environment of any of clauses 15 and 16, wherein paired sNICs announce the same set of VIPs.

[0343] Clause 18: The computing environment of any of the clauses 15-17, wherein flow replication is based on the floating network interfaces, and the direction of flow replication is active to passive.

30 [0344] Clause 19: The computing environment of any of the clauses 15-18, wherein a single floating network interface is programmed on multiple sNICs.

[0345] Clause 20: The computing environment of any of the clauses 15-19, wherein the ToRs are configured to perform traffic load balancing, splitting, or sharding for selected floating network interfaces across VIPs of cards on which the selected floating network

interface is provisioned.

[0346] The disclosure presented herein also encompasses the subject matter set forth in the following clauses:

[0347] Clause 1: A method for processing data packets and implementing policies in a software defined network (SDN) of a virtual computing environment, by at least two SDN appliances configured to disaggregate enforcement of policies of the SDN from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the servers communicatively coupled to network interfaces of at least two top-of-rack switches (ToRs), the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance, the sNICs having a floating network interface configured to provide a virtual port connection to an endpoint within a virtual network of the virtual computing environment, wherein:

[0348] each sNIC that is associated with the first SDN appliance is paired with an sNIC associated with the second SDN appliance;

[0349] each paired sNIC is configured to serve multiple floating network interfaces;

[0350] each floating network interface is serviced by multiple sNICs;

[0351] the method comprising:

[0352] selecting one of the SDN appliances as an active SDN appliance, wherein the other SDN appliance is a passive SDN appliance;

[0353] announcing a different virtual IP (VIP) via border gateway protocol (BGP) for each paired sNIC;

[0354] configuring each floating network interface to be serviced by multiple VIPs;

[0355] splitting, by the ToRs, data traffic equally across different VIPs; and

[0356] performing a flow synchronization process between paired sNICs as passive sNICs become active.

[0357] Clause 2: The method of clause 1, wherein the flow synchronization process comprises:

[0358] implementing at least two colors;

[0359] associating each connection and entry in a flow table with a color;

[0360] establishing a pairing relationship between a primary sNIC and a secondary sNIC; wherein a connection to a paired device can be replicated;

[0361] assigning a color to an sNIC that comes online;

[0362] as the online sNIC receives traffic, adding new flows to the flow table with the

assigned color and replicated new flows to paired sNICs, wherein all connections use the same color until the pairing fails, and continues using the same color if there is no active pairing;

[0363] when pairing is re-established, changing the assigned color to a new color;

5 [0364] using the new color for new flows that are added to the flow table and replicating new flows to the paired sNIC; and

[0365] replicating existing connections to a new device to which a new pairing was created or an existing device to which pairing was re-established, wherein:

10 [0366] only connections which have different color than an actively used color is replicated;

[0367] entries with a currently active color are not replicated.

[0368] Clause 3: The method of any of clauses 1-2, further comprising clearing a state of a flow table of a secondary sNIC when pairing is re-established.

15 [0369] Clause 4: The method of any of clauses 1-3, further comprising ignoring update messages when a connection does not already exist in a flow table.

[0370] Clause 5: The method of any of clauses 1-4, wherein in response to a single card sNIC:

[0371] if outage of the single sNIC persists, selecting a different sNIC on another SDN appliance and initiating pairing with an sNIC on the another appliance; and

20 [0372] announce a VIP of the selected sNIC as active when flow synchronization completes for all floating network interfaces.

[0373] Clause 6: The method of any of clauses 1-5, wherein in response to a single SDN appliance failure: if outage of the single SDN appliance persists, removing existing pairing relationship, allocating a new SDN appliance, and creating a pairing relationship with the  
25 new SDN appliance.

[0374] Clause 7: A system comprising:

30 [0375] a plurality of servers communicatively coupled to at least two software defined network (SDN) appliances configured to disaggregate enforcement of policies of a SDN of a virtual computing environment from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the servers communicatively coupled to network interfaces of at least two top-of-rack switches (ToRs), the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance, the sNICs having a floating network

interface configured to provide a virtual port connection to an endpoint within a virtual network of the virtual computing environment, wherein:

[0376] each sNIC that is associated with the first SDN appliance is paired with an sNIC associated with the second SDN appliance;

5 [0377] each paired sNIC is configured to serve multiple floating network interfaces;

[0378] each floating network interface is serviced by multiple sNICs;

[0379] the system configured to:

[0380] selecting one of the SDN appliances as an active SDN appliance, wherein the other SDN appliance is a passive SDN appliance;

10 [0381] announcing a different virtual IP (VIP) via border gateway protocol (BGP) for each paired sNIC;

[0382] configuring each floating network interface to be serviced by multiple VIPs;

[0383] splitting, by the ToRs, data traffic equally across different VIPs; and

[0384] performing a flow synchronization process between paired sNICs as passive sNICs

15 become active.

[0385] Clause 8: The system of clause 7, wherein the flow synchronization process comprises:

[0386] implementing at least two colors;

[0387] associating each connection and entry in a flow table with a color;

20 [0388] establishing a pairing relationship between a primary sNIC and a secondary sNIC; wherein a connection to a paired device can be replicated;

[0389] assigning a color to an sNIC that comes online; and

[0390] as the online SNIC receives traffic, adding new flows to the flow table with the assigned color and replicated new flows to paired sNICs, wherein all connections use the

25 same color until the pairing fails, and continues using the same color if there is no active pairing;

[0391] when pairing is re-established, changing the assigned color to a new color;

[0392] using the new color for new flows that are added to the flow table and replicating new flows to the paired sNIC;

30 [0393] replicating existing connections to a new device to which a new pairing was created or an existing device to which pairing was re-established, wherein:

[0394] only connections which have different color than an actively used color is replicated;

[0395] entries with a currently active color are not replicated.

[0396] Clause 9: The system of any of clauses 7 and 8, further comprising clearing a state of a flow table of a secondary sNIC when pairing is re-established.

[0397] Clause 10: The system of any clauses 7-9, further comprising ignoring update messages when a connection does not already exist in a flow table.

5 [0398] Clause 11: The system of any clauses 7-10, wherein in response to a single card sNIC:

[0399] if outage of the single sNIC persists, selecting a different sNIC on another SDN appliance and initiating pairing with an sNIC on the another appliance; and

10 [0400] announce a VIP of the selected sNIC as active when flow synchronization completes for all floating network interfaces.

[0401] Clause 12: The system of any clauses 7-11, wherein in response to a single SDN appliance failure: if outage of the single SDN appliance persists, removing existing pairing relationship, allocating a new SDN appliance, and creating a pairing relationship with the new SDN appliance.

15 [0402] Clause 13: The system of any clauses 7-12, wherein 8 colors are implemented that are represented by 3 bits.

[0403] Clause 14: The system of any clauses 7-13, wherein flow splitting is performed by an intelligent ToR, on a source based on stable hashing, or directly on a source node.

[0404] Clause 15: A data center rack comprising:

20 [0405] a plurality of servers communicatively coupled to at least two software defined network (SDN) appliances configured to disaggregate enforcement of policies of a SDN of a virtual computing environment from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the servers communicatively  
25 coupled to network interfaces of at least two top-of-rack switches (ToRs), the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance, the sNICs having a floating network interface configured to provide a virtual port connection to an endpoint within a virtual network of the virtual computing environment, wherein:

30 [0406] each sNIC that is associated with the first SDN appliance is paired with an sNIC associated with the second SDN appliance;

[0407] each paired sNIC is configured to serve multiple floating network interfaces;

[0408] each floating network interface is serviced by multiple sNICs;

[0409] the data center rack configured to:

[0410] selecting one of the SDN appliances as an active SDN appliance, wherein the other SDN appliance is a passive SDN appliance;

[0411] announcing a different virtual IP (VIP) via border gateway protocol (BGP) for each paired sNIC;

5 [0412] configuring each floating network interface to be serviced by multiple VIPs;

[0413] splitting, by the ToRs, data traffic equally across different VIPs; and

[0414] performing a flow synchronization process between paired sNICs as passive sNICs become active.

[0415] Clause 16: The computing environment of clause 15, wherein the flow  
10 synchronization process comprises:

[0416] implementing at least two colors;

[0417] associating each connection and entry in a flow table with a color;

[0418] establishing a pairing relationship between a primary sNIC and a secondary sNIC; wherein a connection to a paired device can be replicated;

15 [0419] assigning a color to an sNIC that comes online;

[0420] as the online sNIC receives traffic, adding new flows to the flow table with the assigned color and replicating new flows to paired sNICs, wherein all connections use the same color until the pairing fails, and continues using the same color if there is no active pairing;

20 [0421] when pairing is re-established, changing the assigned color to a new color;

[0422] using the new color for new flows that are added to the flow table and replicating new flows to the paired sNIC; and

[0423] replicating existing connections to a new device to which a new pairing was created or an existing device to which pairing was re-established, wherein:

25 [0424] only connections which have different color than an actively used color is replicated;

[0425] entries with a currently active color are not replicated.

[0426] Clause 17: The computing environment of any of clauses 15 and 16, further comprising clearing a state of a flow table of a secondary sNIC when pairing is re-  
30 established.

[0427] Clause 18: The computing environment of any of the clauses 15-17, further comprising ignoring update messages when a connection does not already exist in the flow table.

[0428] Clause 19: The computing environment of any of the clauses 15-18, wherein in

response to a single card sNIC:

**[0429]** if outage of a single sNIC persists, selecting a different sNIC on another SDN appliance and initiating pairing with an sNIC on the another appliance; and

**[0430]** announce a VIP of the selected sNIC as active when flow synchronization  
5 completes for all floating network interfaces.

**[0431]** Clause 20: The computing environment of any of the clauses 15-19, wherein in response to a single SDN appliance failure: if outage of the single SDN appliance persists, removing existing pairing relationship, allocating a new SDN appliance, and creating a pairing relationship with the new SDN appliance.

## CLAIMS

1. A method for processing data packets and implementing policies in a software defined network (SDN) of a virtual computing environment, by at least one SDN appliance configured to disaggregate enforcement of policies of the SDN from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance, the method comprising:

receiving, at the SDN appliance from a device that is remote from the virtual computing environment, a data packet addressed to an endpoint in a virtual network hosted by one of the virtual machines, the data packet comprising an identifier indicative of the remote device;

based on the identifier:

determining that the data packet is associated with the virtual network; and

based on the determined association, mapping one of a plurality of policies to a data flow of the virtual network;

modifying, by the SDN appliance, the packet in accordance with the mapped policy; wherein the mapped policy is dynamically adjustable based on the data flow; and

forwarding the modified packet to the endpoint in the virtual network.

2. The method of claim 1, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among the servers.

3. The method of claim 1, wherein the SDN appliance is implemented as two physical devices that are communicatively coupled to at least two top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the physical devices.

4. The method of claim 3, wherein storage traffic bypasses the sNICs.

5. The method of claim 1, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among one or more top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the top-of-rack switches.

6. The method of claim 2, wherein at least two top-of-rack switches are

configured with SDN agents configured to manage functionality of the SDN appliance.

7. A system comprising:

a plurality of servers communicatively coupled to at least one software defined network (SDN) appliance configured to disaggregate enforcement of policies of a SDN of a virtual computing environment from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance;

the system configured to:

receive, from a device that is remote from the virtual computing environment, a data packet addressed to an endpoint in a virtual network hosted by one of the virtual machines, the data packet comprising an identifier indicative of the remote device;

based on the identifier:

determining that the data packet is associated with the virtual network; and

based on the determined association, mapping one of a plurality of policies to a data flow of the virtual network;

modifying, by the SDN appliance, the packet in accordance with the mapped policy; wherein the mapped policy is dynamically adjustable based on the data flow; and

forwarding the modified packet to the endpoint in the virtual network.

8. The system of claim 7, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among the plurality of servers.

9. The system of claim 7, wherein the SDN appliance is implemented as two physical devices that are communicatively coupled to at least two top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the physical devices.

10. The system of claim 9, wherein storage traffic bypasses the sNICs.

11. The system of claim 7, wherein the SDN appliance is implemented as a distributed appliance where the sNICs are physically distributed among one or more top-of-rack switches, the servers communicatively coupled to network interfaces of the top-of-rack switches so that each of the servers have a switchable communications path to each sNIC of the top-of-rack switches.

12. The system of claim 9, wherein at least two top-of-rack switches are configured with SDN agents configured to manage functionality of the SDN appliance.

13. The system of claim 7, wherein the policy is applied at selectively placed network hops in the virtual network.

14. The system of claim 7, wherein the SDN appliance is configured to apply policies of the virtual computing environment to data traffic on the virtual network after the data traffic leaves its source and before the data traffic reaches its destination.

15. A data center rack comprising:

a plurality of servers; the servers communicatively coupled to at least one software defined network (SDN) appliance configured to disaggregate enforcement of policies of a SDN of a virtual computing environment from hosts of the virtual computing environment, the hosts implemented on servers communicatively coupled to network interfaces of the SDN appliance, the servers hosting a plurality of virtual machines, the SDN appliance comprising a plurality of smart network interface cards (sNICs) configured to implement functionality of the SDN appliance;

the data center rack configured to:

receive, from a device that is remote from a software defined network (SDN) of a virtual computing environment, a data packet addressed to an endpoint in a virtual network hosted by one of the virtual machines, the data packet comprising an identifier indicative of the remote device;

based on the identifier:

determining that the data packet is associated with the virtual network; and

based on the determined association, mapping one of a plurality of policies to a data flow of the virtual network;

modifying, by the SDN appliance, the packet in accordance with the mapped policy;

wherein the mapped policy is dynamically adjustable based on the data flow; and

forwarding the modified packet to the endpoint in the virtual network.

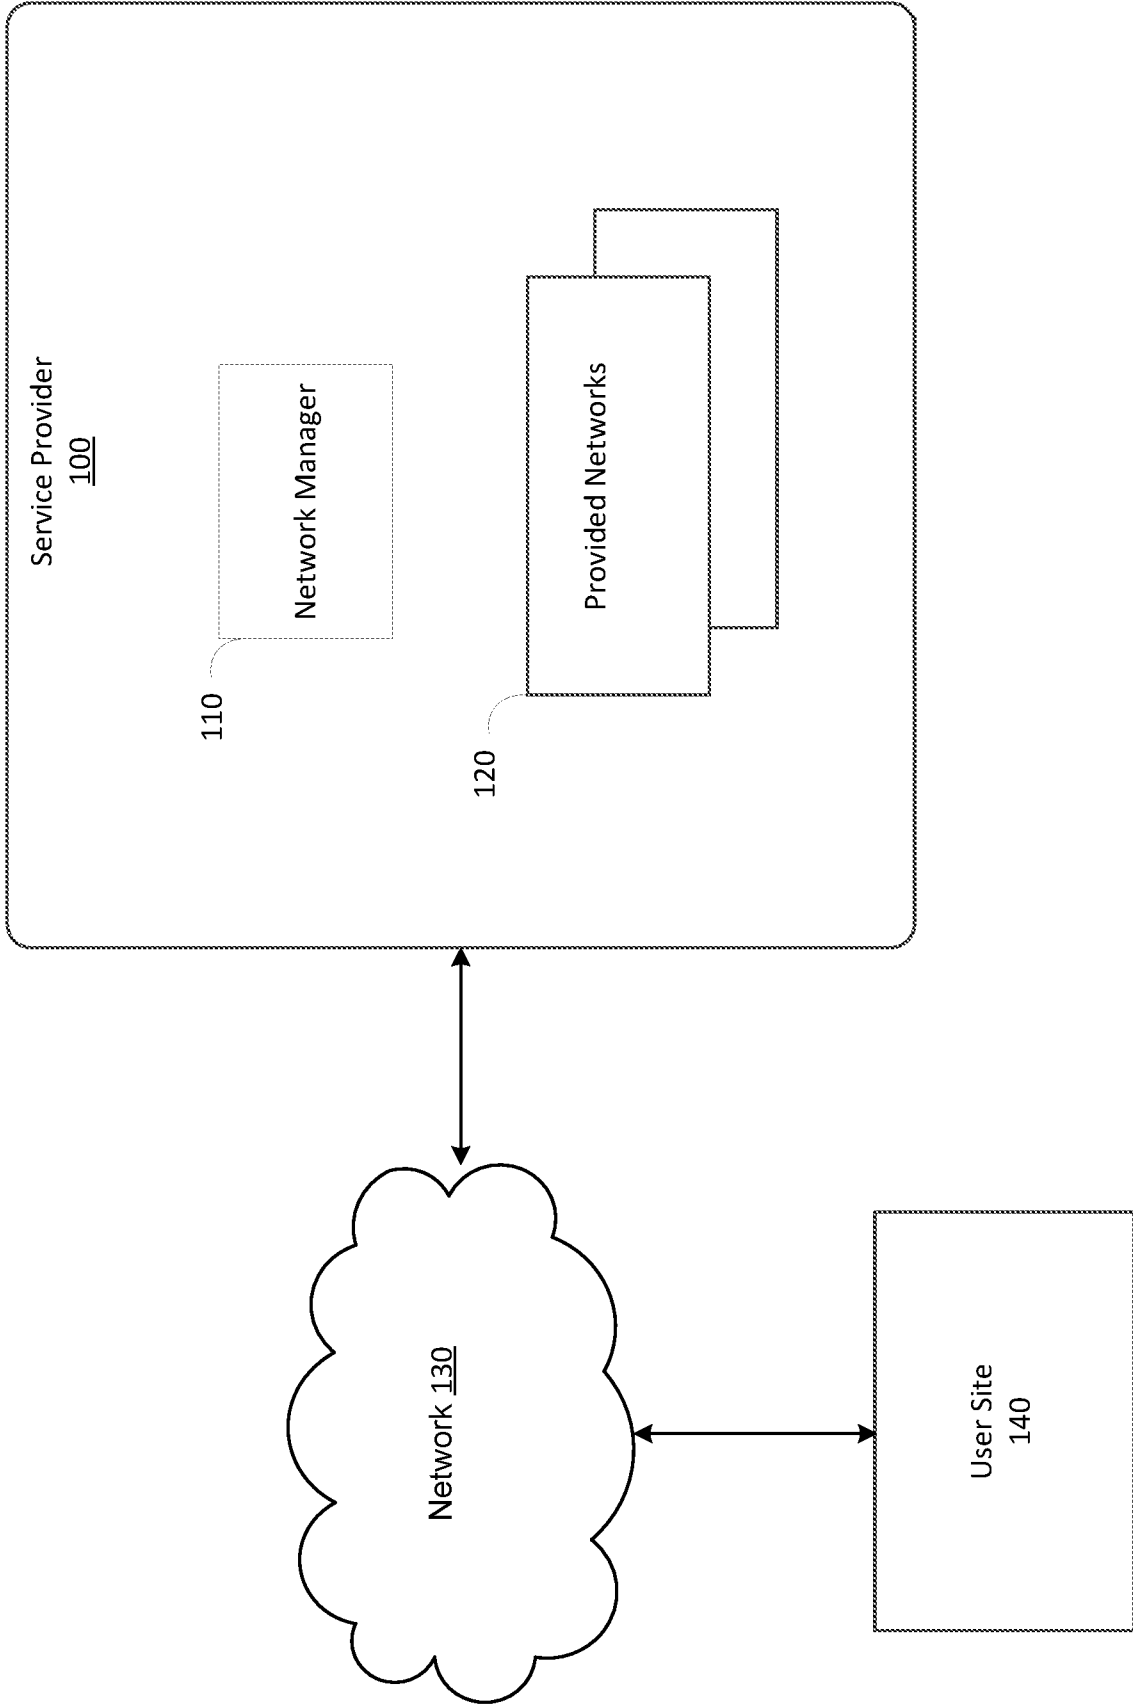


FIG. 1




FIG. 2

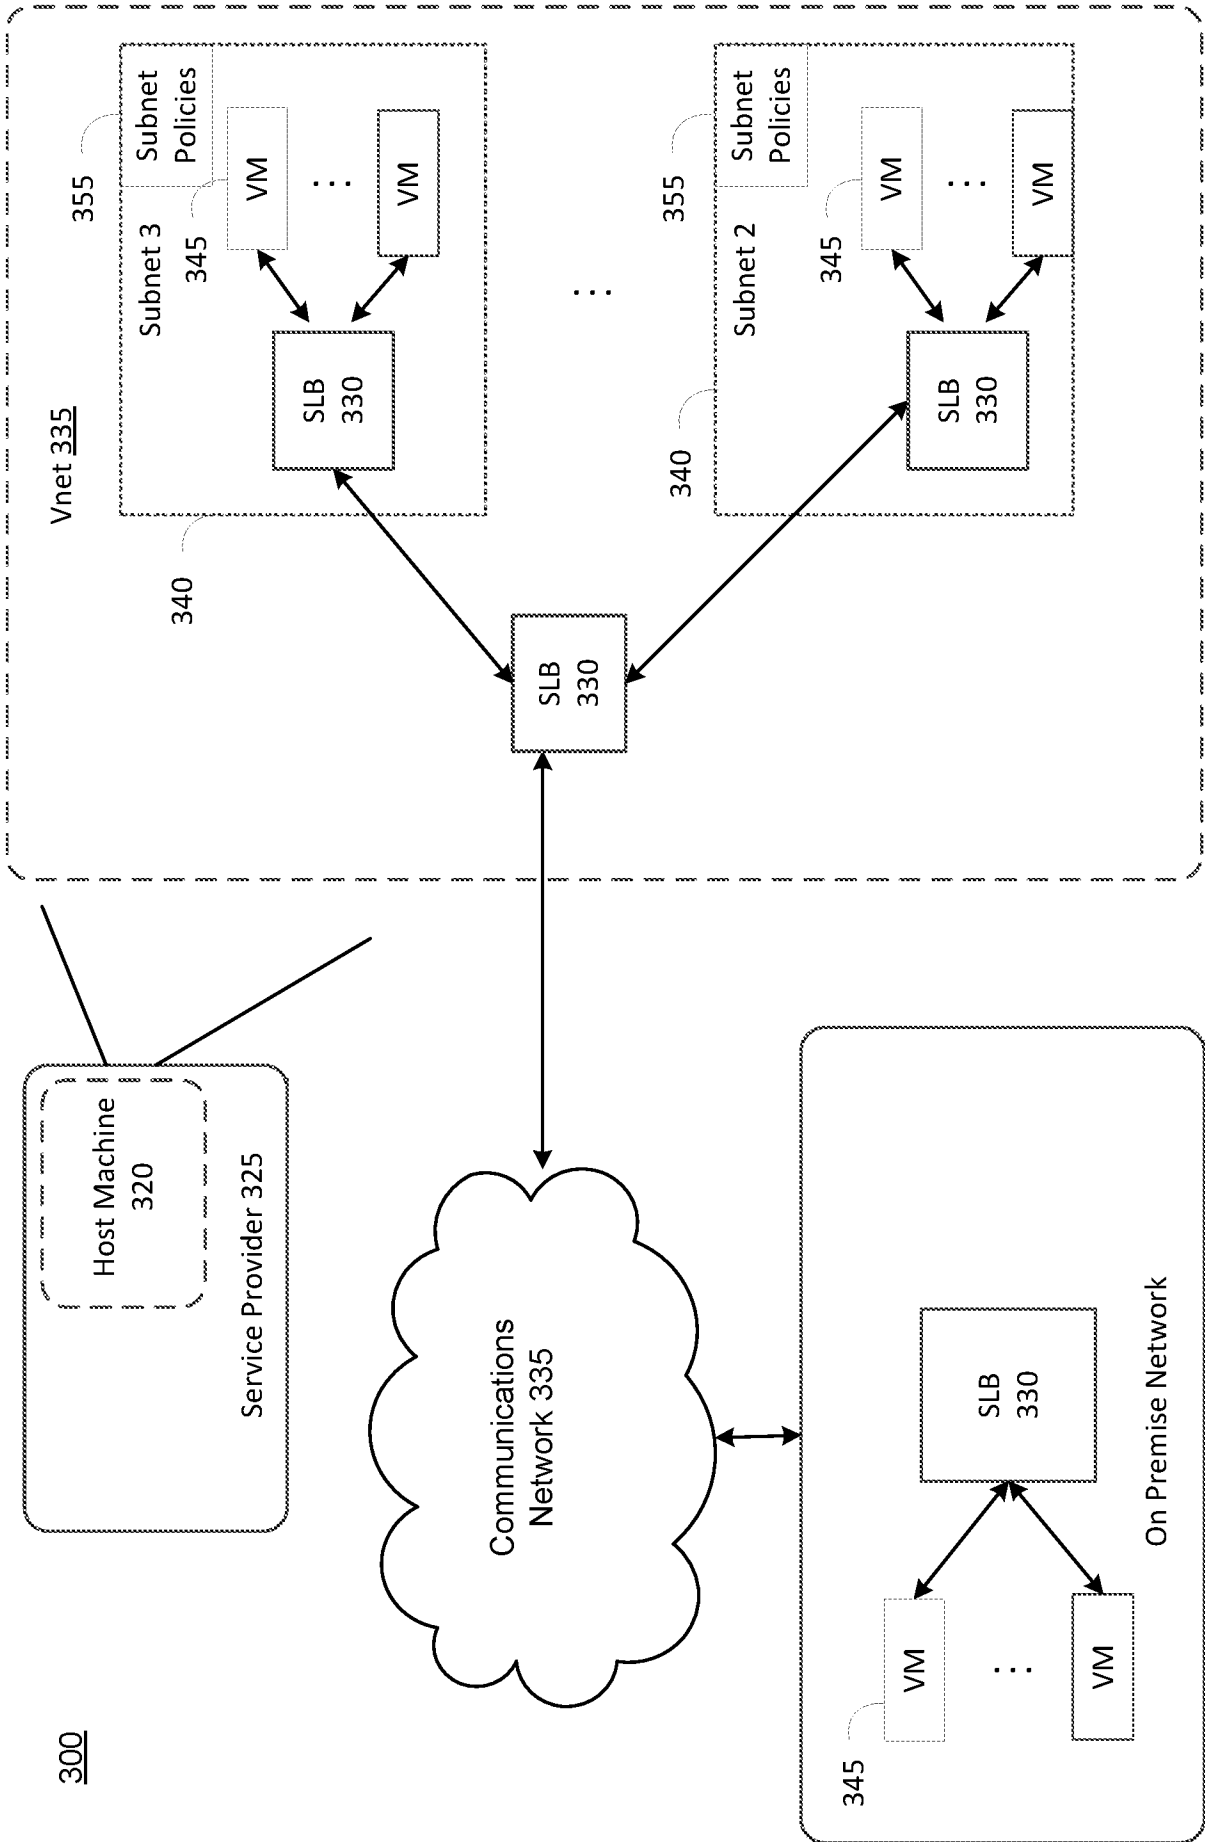


FIG. 3

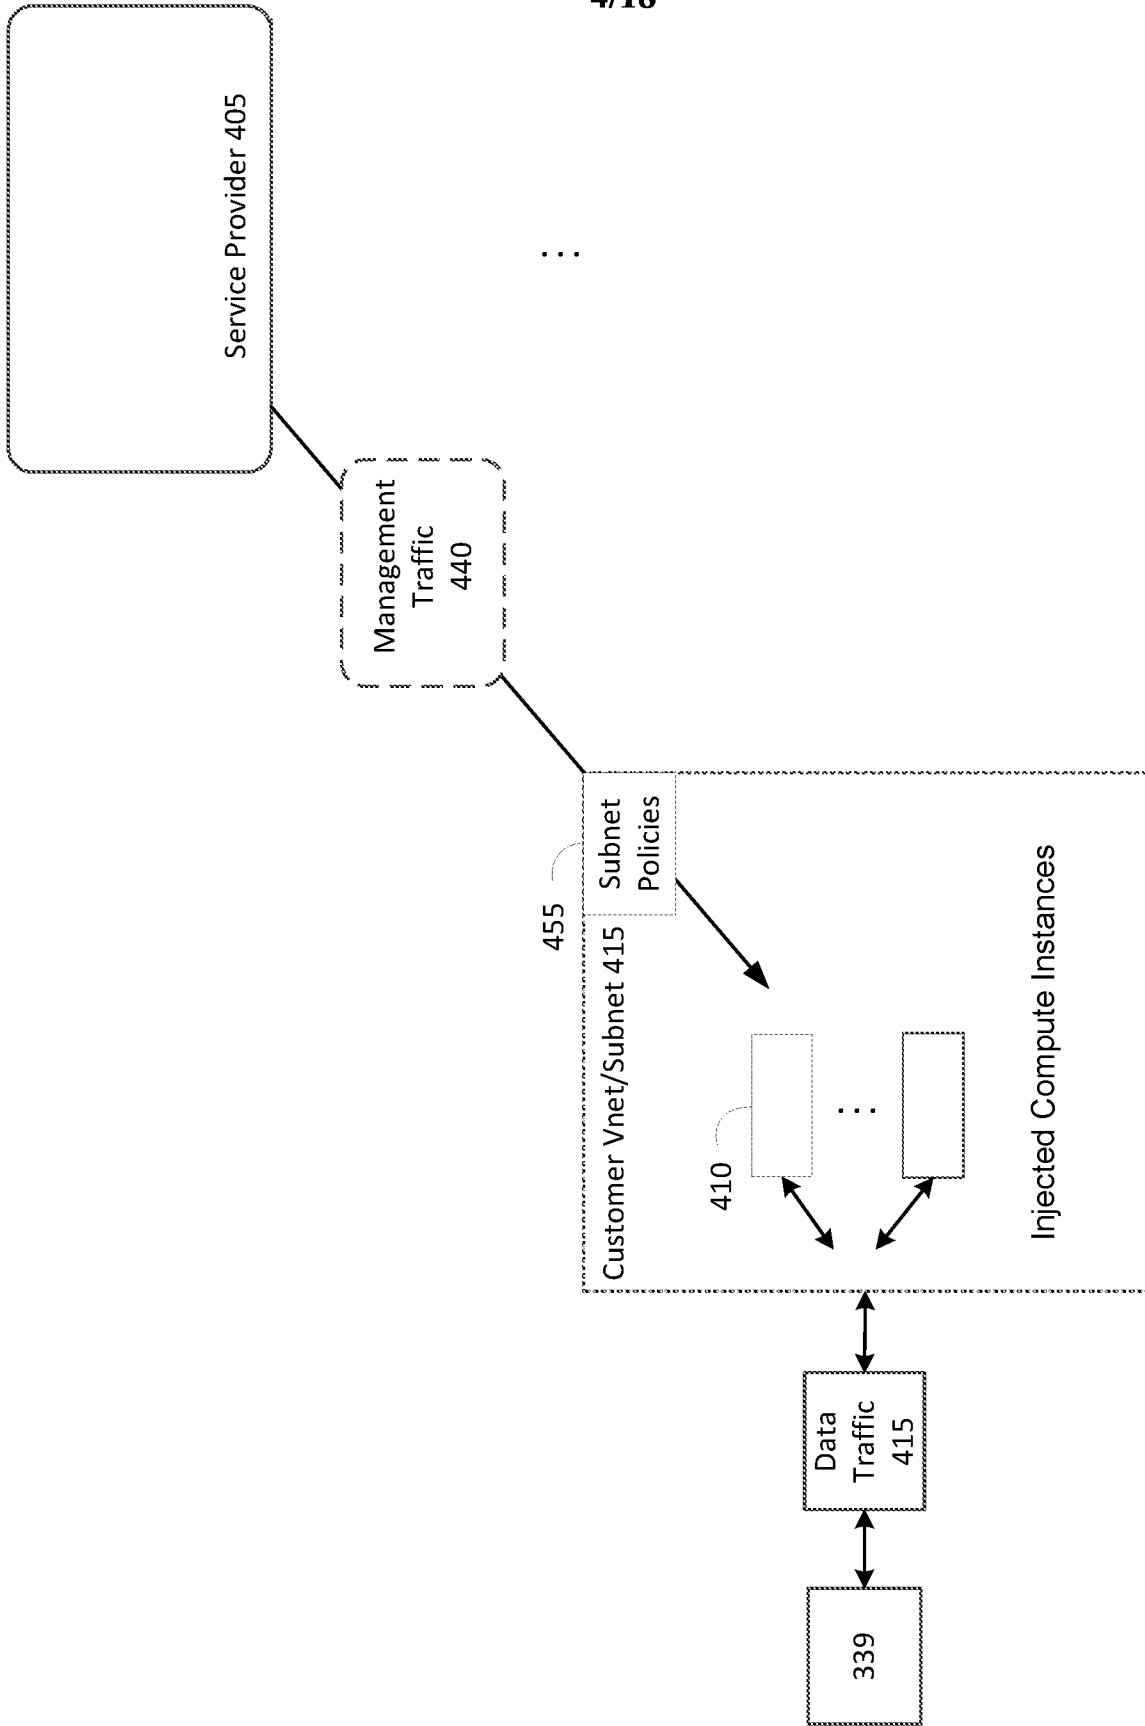


FIG. 4

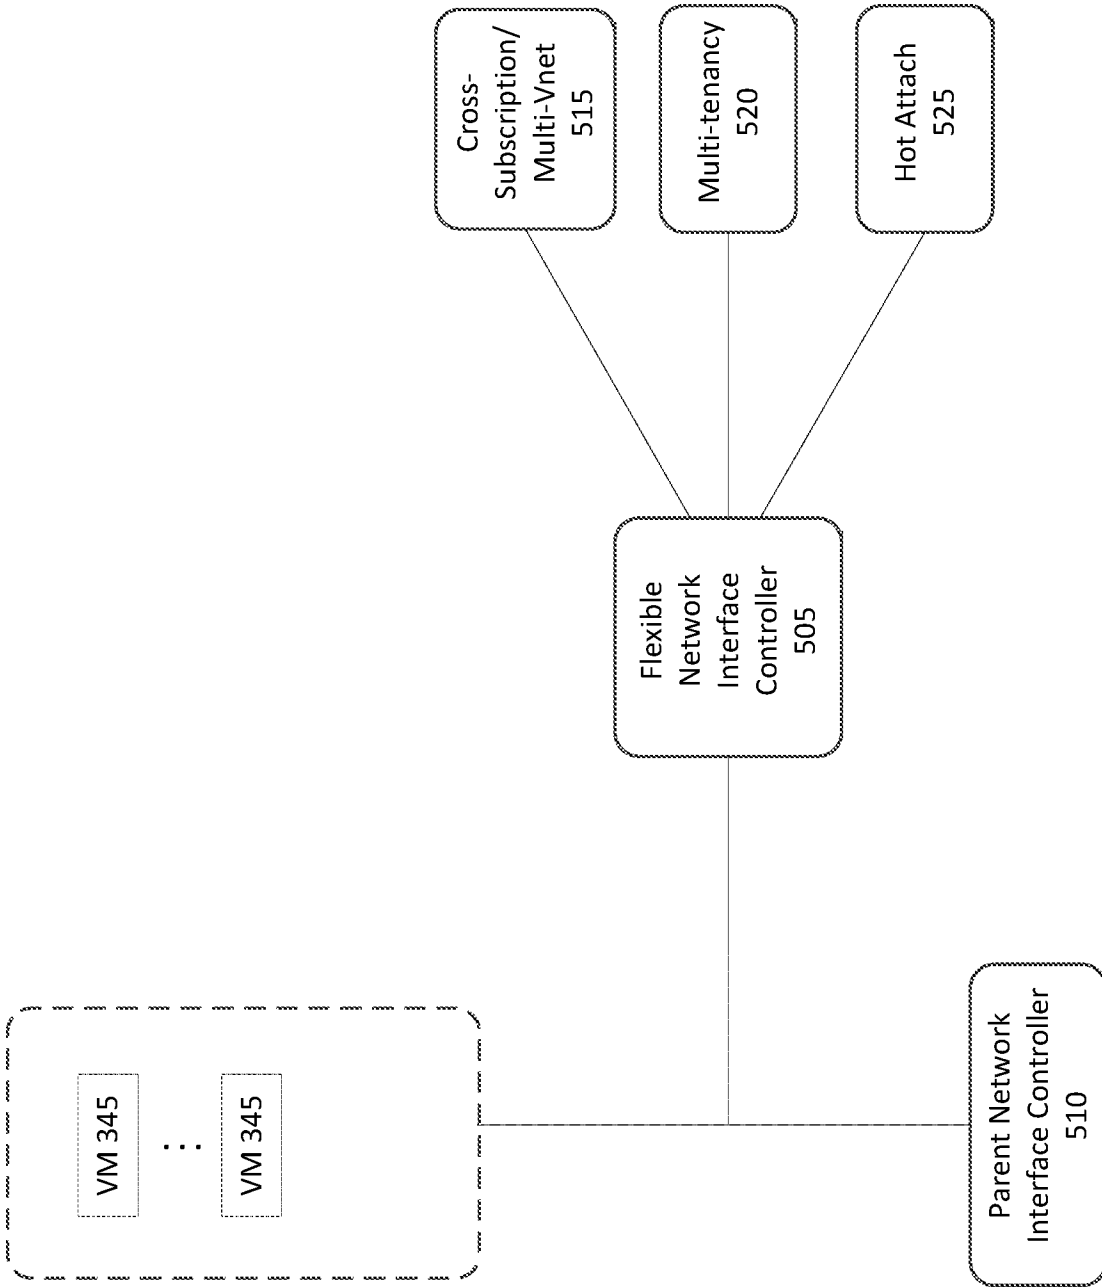


FIG. 5

6/18

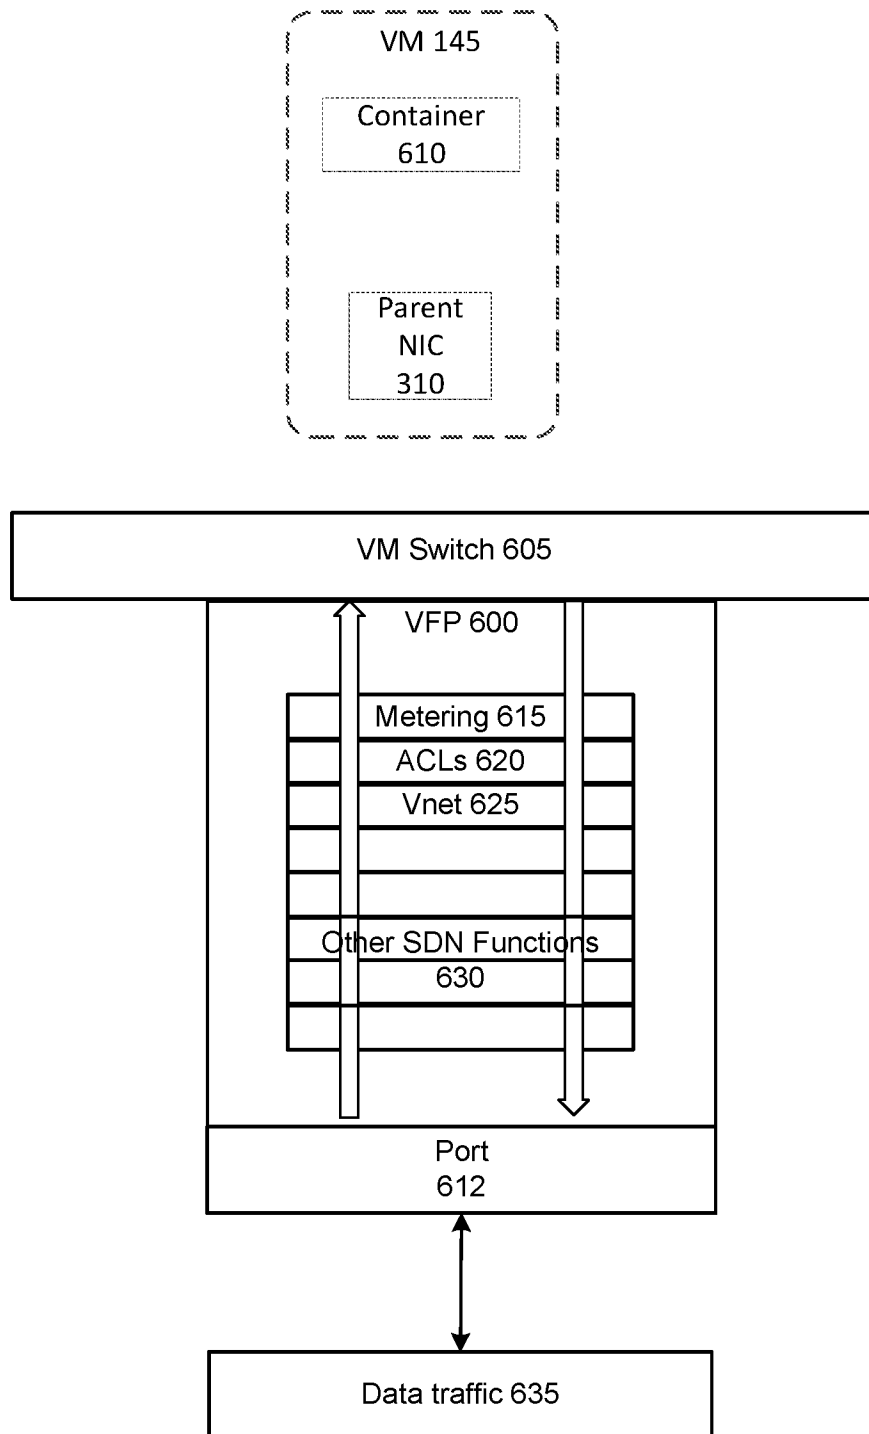


FIG. 6




FIG. 7

800

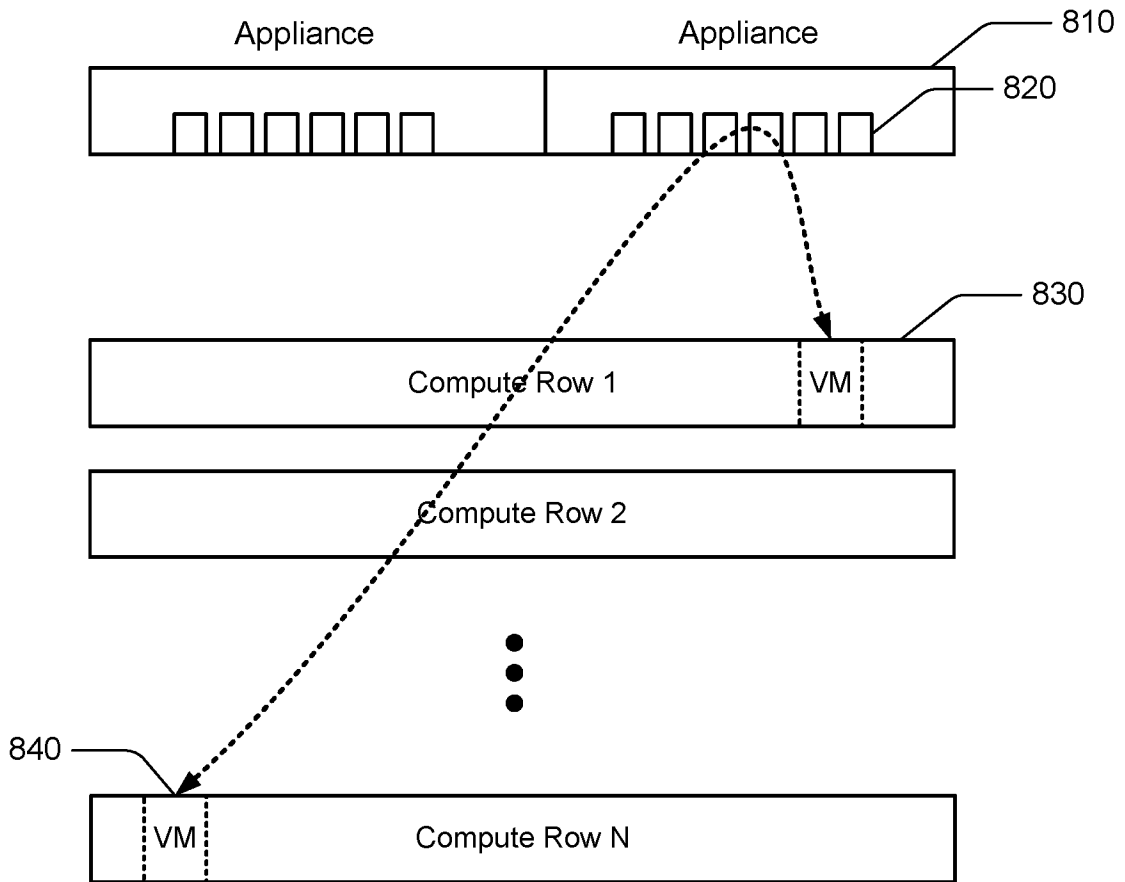


FIG. 8

9/18

900




FIG. 9




FIG. 10

11/18

1100




FIG. 11

12/18

1200




FIG. 12A

13/18

1200




FIG. 12B




FIG. 13





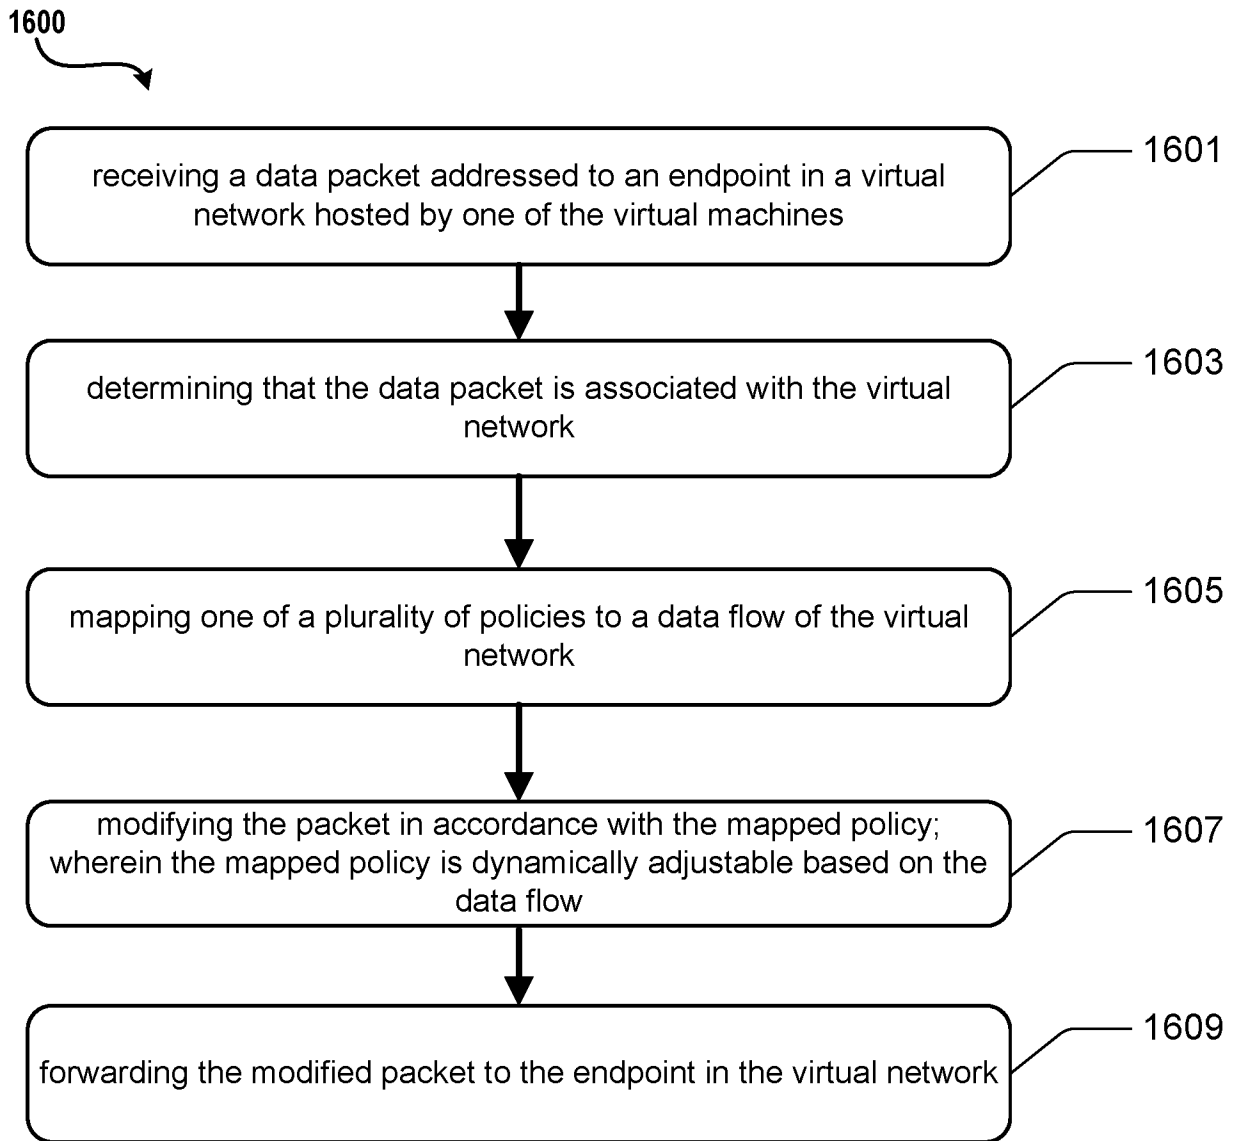
FIG. 14



Where X = Special switch , Source FPGA , CX7

ENI on multiple cards gives bigger pool, but makes solution complex.

FIG. 15



**FIG. 16**

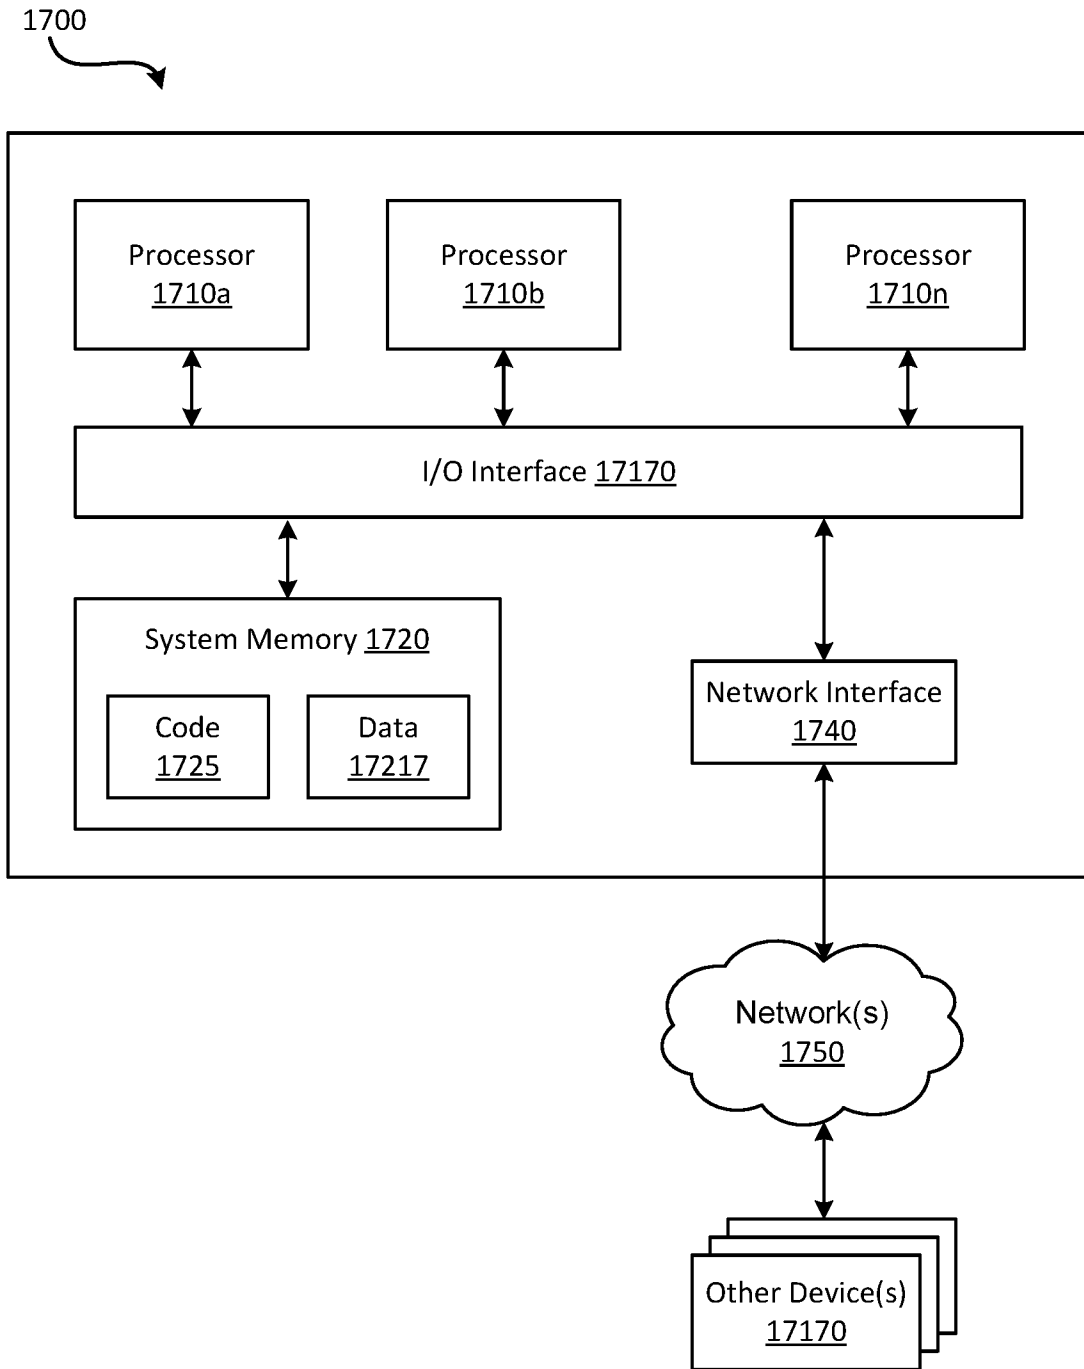


FIG. 17

# INTERNATIONAL SEARCH REPORT

International application No  
**PCT/US2022/020673**

**A. CLASSIFICATION OF SUBJECT MATTER**  
**INV. H04L43/20 G06F9/455 H04L9/40**  
**ADD.**

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

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
**H04L G06F**

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

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

**EPO-Internal**

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
<b>X</b>	<b>US 2020/371842 A1 (TEWARI RISHABH [US] ET AL) 26 November 2020 (2020-11-26)</b>	<b>1, 7, 14, 15</b>
<b>Y</b>	<b>paragraphs [0003], [0006], [0020] - [0026]; figures 1, 2a, 2b, 3</b> -----	<b>2-6, 8-13</b>
<b>Y</b>	<b>US 2020/278892 A1 (NAINAR NAGENDRA KUMAR [US] ET AL) 3 September 2020 (2020-09-03)</b> <b>paragraphs [0027], [0071] - [0078]; figures 1, 7</b> -----	<b>2-6, 8-13</b>

Further documents are listed in the continuation of Box C.

See patent family annex.

\* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search  
  
**24 May 2022**

Date of mailing of the international search report  
  
**02/06/2022**

Name and mailing address of the ISA/  
 European Patent Office, P.B. 5818 Patentlaan 2  
 NL - 2280 HV Rijswijk  
 Tel. (+31-70) 340-2040,  
 Fax: (+31-70) 340-3016

Authorized officer  
  
**Cicalese, Danilo**

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2022/020673

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
<b>US 2020371842 A1</b>	<b>26-11-2020</b>	<b>EP 3977276 A1</b>	<b>06-04-2022</b>
		<b>US 2020371842 A1</b>	<b>26-11-2020</b>
		<b>WO 2020242957 A1</b>	<b>03-12-2020</b>
-----			
<b>US 2020278892 A1</b>	<b>03-09-2020</b>	<b>CN 113454971 A</b>	<b>28-09-2021</b>
		<b>EP 3932041 A1</b>	<b>05-01-2022</b>
		<b>US 2020278892 A1</b>	<b>03-09-2020</b>
		<b>WO 2020176673 A1</b>	<b>03-09-2020</b>
-----			