

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2017/132399 A1

(43) International Publication Date

3 August 2017 (03.08.2017)

(51) International Patent Classification:

G06F 9/455 (2006.01) *H04L 29/12* (2006.01)
H04L 12/931 (2013.01)

(74) Agents: MEYER, Sheldon, R. et al.; Tucker Ellis LLP, One Market Plaza, Steuart Tower, Suite 700, San Francisco, California 94105 (US).

(21) International Application Number:

PCT/US2017/015169

(22) International Filing Date:

26 January 2017 (26.01.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/287,712 27 January 2016 (27.01.2016) US
15/415,668 25 January 2017 (25.01.2017) US

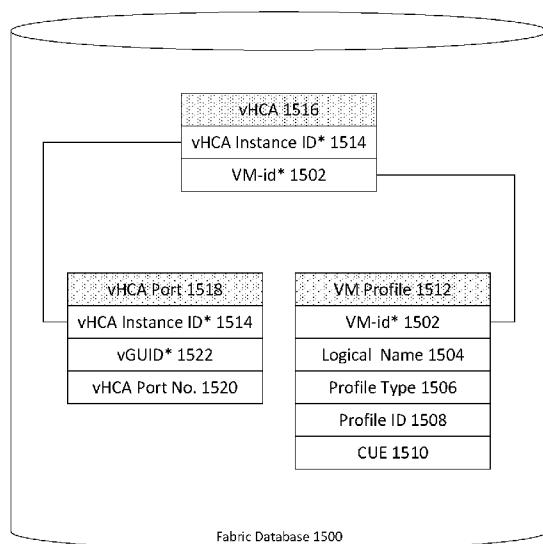
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, 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, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).

(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).

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR DEFINING VIRTUAL MACHINE FABRIC PROFILES OF VIRTUAL MACHINES IN A HIGH-PERFORMANCE COMPUTING ENVIRONMENT



(57) Abstract: Systems and methods for defining virtual machine fabric profiles of virtual machines which can contain a virtual machine identifier, a virtual host channel adapter instance ID, and a virtual globally unique identifier. The virtual machine identifier, virtual host channel adapter instance ID, and virtual globally unique identifier can be mapped to each other so that the virtual host channel adapter instance ID, and virtual globally unique identifier can be retrieved through access to the virtual machine identifier. Further, a relationship can be created between a P_Key that defines an admin partition and the virtual globally unique identifier, where the relationship between the P_Key and the virtual globally unique identifier defines the virtual globally unique identifier as a member of the admin partition defined by the P_Key.

FIGURE 15



Published:

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

**SYSTEM AND METHOD FOR DEFINING VIRTUAL MACHINE FABRIC PROFILES OF
VIRTUAL MACHINES IN A HIGH-PERFORMANCE COMPUTING ENVIRONMENT**

COPYRIGHT NOTICE

5 A portion of the disclosure of this patent document contains
material which is subject to copyright protection. The copyright
owner has no objection to the facsimile reproduction by anyone
of the patent document or the patent disclosure, as it appears
in the Patent and Trademark Office patent file or records, but
10 otherwise reserves all copyright rights whatsoever.

Field of Invention

[001] The present invention is generally related to computer systems, and is particularly
related to defining virtual machine fabric profiles of virtual machines.

15 **Background**

[002] As larger cloud computing architectures are introduced, the performance and
administrative bottlenecks associated with the traditional network and storage have become
a significant problem. There has been an increased interest in using high performance
20 lossless interconnects such as InfiniBand™ (IB) technology as the foundation for a cloud
computing fabric. This is the general area that embodiments of the invention are intended to
address.

Summary

[003] Described herein are systems and methods for defining virtual machine fabric
profiles of virtual machines. An exemplary embodiment can provide a virtual machine
identifier, a virtual host channel adapter instance ID, and a virtual globally unique identifier.
The virtual machine identifier, virtual host channel adapter instance ID, and virtual globally
unique identifier can be mapped to each other so that the virtual host channel adapter
30 instance ID, and virtual globally unique identifier can be retrieved through access to the
virtual machine identifier. Further, a relationship can be created between a P_Key that
defines an admin partition and the virtual globally unique identifier, where the relationship
between the P_Key and the virtual globally unique identifier defines the virtual globally
unique identifier as a member of the admin partition defined by the P_Key.

35

Brief Description of the Figures

[004] Figure 1 shows an illustration of an InfiniBand™ environment, in accordance with

an embodiment.

[005] **Figure 2** shows an illustration of a partitioned cluster environment, in accordance with an embodiment

[006] **Figure 3** shows an illustration of a tree topology in a network environment, in accordance with an embodiment.

[007] **Figure 4** shows an exemplary shared port architecture, in accordance with an embodiment.

[008] **Figure 5** shows an exemplary vSwitch architecture, in accordance with an embodiment.

10 [009] **Figure 6** shows an exemplary vPort architecture, in accordance with an embodiment.

[0010] **Figure 7** shows an exemplary vSwitch architecture with prepopulated LIDs, in accordance with an embodiment.

15 [0011] **Figure 8** shows an exemplary vSwitch architecture with dynamic LID assignment, in accordance with an embodiment.

[0012] **Figure 9** shows an exemplary vSwitch architecture with vSwitch with dynamic LID assignment and prepopulated LIDs, in accordance with an embodiment.

[0013] **Figure 10** shows an exemplary multi-subnet InfiniBand™ fabric, in accordance with an embodiment.

20 [0014] **Figure 11** shows an exemplary InfiniBand™ fabric and subnet including exemplary physical and logical subnet resources, in accordance with an embodiment.

[0015] **Figure 12** shows an exemplary InfiniBand™ fabric and subnet including exemplary subnet resources as members of different admin partitions, in accordance with an embodiment.

25 [0016] **Figure 13** shows an exemplary InfiniBand™ fabric and subnet including exemplary subnet resources as members of a hierarchical management scheme, including both admin partitions and resource domains, in accordance with an embodiment.

[0017] **Figure 14** is a flowchart for a method for assigning mutual access rights to members and associated resources of an admin partition associated with a fabric-level resource domain, in accordance with an embodiment.

30 [0018] **Figure 15** shows an exemplary database structure for storing VM fabric profile information, in accordance with an embodiment.

[0019] **Figure 16** is a flow chart for making a VM fabric profile available to subnet resources, in accordance with an embodiment.

35 [0020] **Figure 17** is a flow chart for creating a virtual machine fabric profile for a virtual machine, in accordance with an embodiment.

Detailed Description

[0021] The invention is illustrated, by way of example and not by way of limitation, in the figures of the accompanying drawings in which like references indicate similar elements. It should be noted that references to “an” or “one” or “some” embodiment(s) in this disclosure are not necessarily to the same embodiment, and such references mean at least one. While specific implementations are discussed, it is understood that the specific implementations are provided for illustrative purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without departing from the scope and spirit of the invention.

5 [0022] Common reference numerals can be used to indicate like elements throughout the drawings and detailed description; therefore, reference numerals used in a figure may or may not be referenced in the detailed description specific to such figure if the element is described elsewhere.

10 [0023] Described herein are systems and methods for defining virtual machine fabric profiles of virtual machines.

15 [0024] The following description of the invention uses an InfiniBand™ (IB) network as an example for a high performance network. Throughout the following description, reference can be made to the InfiniBand™ specification (also referred to variously as the InfiniBand specification, IB specification, or the legacy IB specification). Such reference is understood to refer to the InfiniBand® Trade Association Architecture Specification, Volume 1, Version 20 1.3, released March, 2015, available at <http://www.infinibandta.org>, which is herein incorporated by reference in its entirety. It will be apparent to those skilled in the art that other types of high performance networks can be used without limitation. The following description also uses the fat-tree topology as an example for a fabric topology. It will be apparent to those skilled in the art that other types of fabric topologies can be used without limitation.

InfiniBand™

30 [0025] InfiniBand™ (IB) is an open standard lossless network technology developed by the InfiniBand™ Trade Association. The technology is based on a serial point-to-point full-duplex interconnect that offers high throughput and low latency communication, geared particularly towards high-performance computing (HPC) applications and datacenters.

35 [0026] The InfiniBand™ Architecture (IBA) supports a two-layer topological division. At the lower layer, IB networks are referred to as subnets, where a subnet can include a set of hosts interconnected using switches and point-to-point links. At the higher level, an IB fabric constitutes one or more subnets, which can be interconnected using routers.

[0027] Within a subnet, hosts can be connected using switches and point-to-point links.

Additionally, there can be a master management entity, the subnet manager (SM), which resides on a designated device in the subnet. The subnet manager is responsible for configuring, activating and maintaining the IB subnet. Additionally, the subnet manager (SM) can be responsible for performing routing table calculations in an IB fabric. Here, for 5 example, the routing of the IB network aims at proper load balancing between all source and destination pairs in the local subnet.

[0028] Through the subnet management interface, the subnet manager exchanges control packets, which are referred to as subnet management packets (SMPs), with subnet management agents (SMAs). The subnet management agents reside on every IB subnet 10 device. By using SMPs, the subnet manager is able to discover the fabric, configure end-nodes and switches, and receive notifications from SMAs.

[0029] In accordance with an embodiment, intra-subnet routing in an IB network can be based on linear forwarding tables (LFTs) stored in the switches. The LFTs are calculated by the SM according to the routing mechanism in use. In a subnet, Host Channel Adapter (HCA) 15 ports on the end-nodes and switches are addressed using local identifiers (LIDs). Each entry in an LFT consists of a destination LID (DLID) and an output port. Only one entry per LID in the table is supported. When a packet arrives at a switch, its output port is determined by looking up the DLID in the forwarding table of the switch. The routing is deterministic as packets take the same path in the network between a given source-destination pair (LID 20 pair).

[0030] Generally, all other subnet managers, excepting the master subnet manager, act in standby mode for fault-tolerance. In a situation where a master subnet manager fails, however, a new master subnet manager is negotiated by the standby subnet managers. The master subnet manager also performs periodic sweeps of the subnet to detect any topology 25 changes and reconfigures the network accordingly.

[0031] Furthermore, hosts and switches within a subnet can be addressed using local identifiers (LIDs), and a single subnet can be limited to 49151 unicast LIDs. Besides the LIDs, which are the local addresses that are valid within a subnet, each IB device can have a 64-bit global unique identifier (GUID). A GUID can be used to form a global identifier (GID), 30 which is an IB layer three (L3) address.

[0032] The SM can calculate routing tables (i.e., the connections/routes between each pair of nodes within the subnet) at network initialization time. Furthermore, the routing tables can be updated whenever the topology changes, in order to ensure connectivity and optimal performance. During normal operations, the SM can perform periodic light sweeps of the 35 network to check for topology changes. If a change is discovered during a light sweep or if a message (trap) signaling a network change is received by the SM, the SM can reconfigure the network according to the discovered changes.

[0033] For example, the SM can reconfigure the network when the network topology changes, such as when a link goes down, when a device is added, or when a link is removed. The reconfiguration steps can include the steps performed during the network initialization. Furthermore, the reconfigurations can have a local scope that is limited to the 5 subnets in which the network changes occurred. Also, the segmenting of a large fabric with routers may limit the reconfiguration scope.

[0034] An example InfiniBand™ fabric is shown in **Figure 1**, which shows an illustration of an InfiniBand™ environment 100, in accordance with an embodiment. In the example shown in Figure 1, nodes A-E, 101-105, use the InfiniBand™ fabric 120 to communicate, via the 10 respective host channel adapters 111-115. In accordance with an embodiment, the various nodes, e.g., nodes A-E 101-105, can be represented by various physical devices. In accordance with an embodiment, the various nodes, e.g., nodes A-E 101-105, can be represented by various virtual devices, such as virtual machines.

15 Data Partitions in InfiniBand™

[0035] In accordance with an embodiment, IB networks can support partitioning as a security mechanism to provide for isolation of logical groups of systems sharing a network fabric. Each HCA port on a node in the fabric can be a member of one or more partitions. In accordance with an embodiment, the present disclosure provides for two types of partitions 20 that can be defined within an IB subnet – data partitions (discussed in detail in the following paragraphs) and admin partitions (discussed in detail later in the disclosure).

[0036] Data partition memberships are managed by a centralized partition manager, which can be part of the SM. The SM can configure data partition membership information on each 25 port as a table of 16-bit partition keys (P_Keys). The SM can also configure switch and router ports with the data partition enforcement tables containing P_Key information associated with the end-nodes that send or receive data traffic through these ports. Additionally, in a general case, data partition membership of a switch port can represent a union of all membership indirectly associated with LIDs routed via the port in an egress (towards the link) direction.

[0037] In accordance with an embodiment, data partitions are logical groups of ports such 30 that the members of a group can only communicate to other members of the same logical group. At host channel adapters (HCAs) and switches, packets can be filtered using the data partition membership information to enforce isolation. Packets with invalid partitioning information can be dropped as soon as the packets reaches an incoming port. In partitioned IB systems, data partitions can be used to create tenant clusters. With data partition enforcement in place, a node cannot communicate with other nodes that belong to a different tenant cluster. In this way, the security of the system can be guaranteed even in the presence of compromised or malicious tenant nodes.

[0038] In accordance with an embodiment, for the communication between nodes, Queue Pairs (QPs) and End-to-End contexts (EECs) can be assigned to a particular data partition, except for the management Queue Pairs (QP0 and QP1). The P_Key information can then be added to every IB transport packet sent. When a packet arrives at an HCA port or a 5 switch, its P_Key value can be validated against a table configured by the SM. If an invalid P_Key value is found, the packet is discarded immediately. In this way, communication is allowed only between ports sharing a data partition.

[0039] An example of IB data partitions is shown in **Figure 2**, which shows an illustration of a data partitioned cluster environment, in accordance with an embodiment. In the example 10 shown in Figure 2, nodes A-E, 101-105, use the InfiniBand™ fabric, 120, to communicate, via the respective host channel adapters 111-115. The nodes A-E are arranged into data partitions, namely data partition 1, 130, data partition 2, 140, and data partition 3, 150. Data partition 1 comprises node A 101 and node D 104. Data partition 2 comprises node A 101, node B 102, and node C 103. Data partition 3 comprises node C 103 and node E 105. 15 Because of the arrangement of the data partitions, node D 104 and node E 105 are not allowed to communicate as these nodes do not share a data partition. Meanwhile, for example, node A 101 and node C 103 are allowed to communicate as these nodes are both members of data partition 2, 140.

20 Virtual Machines in InfiniBand™

[0040] During the last decade, the prospect of virtualized High Performance Computing (HPC) environments has improved considerably as CPU overhead has been practically removed through hardware virtualization support; memory overhead has been significantly reduced by virtualizing the Memory Management Unit; storage overhead has been reduced 25 by the use of fast SAN storages or distributed networked file systems; and network I/O overhead has been reduced by the use of device passthrough techniques like Single Root Input/Output Virtualization (SR-IOV). It is now possible for clouds to accommodate virtual HPC (vHPC) clusters using high performance interconnect solutions and deliver the necessary performance.

30 [0041] However, when coupled with lossless networks, such as InfiniBand™ (IB), certain cloud functionality, such as live migration of virtual machines (VMs), still remains an issue due to the complicated addressing and routing schemes used in these solutions. IB is an interconnection network technology offering high bandwidth and low latency, thus, is very well suited for HPC and other communication intensive workloads.

35 [0042] The traditional approach for connecting IB devices to VMs is by utilizing SR-IOV with direct assignment. However, achieving live migration of VMs assigned with IB Host Channel Adapters (HCAs) using SR-IOV has proved to be challenging. Each IB connected

node has three different addresses: LID, GUID, and GID. When a live migration happens, one or more of these addresses change. Other nodes communicating with the VM-in-migration can lose connectivity. When this happens, the lost connection can be attempted to be renewed by locating the virtual machine's new address to reconnect to by sending Subnet

5 Administration (SA) path record queries to the IB Subnet Manager (SM).

[0043] IB uses three different types of addresses. A first type of address is the 16 bits Local Identifier (LID). At least one unique LID is assigned to each HCA port and each switch by the SM. The LIDs are used to route traffic within a subnet. Since the LID is 16 bits long, 65536 unique address combinations can be made, of which only 49151 (0x0001-0xBFFF) can be used as unicast addresses. Consequently, the number of available unicast addresses defines the maximum size of an IB subnet. A second type of address is the 64 bits Global Unique Identifier (GUID) assigned by the manufacturer to each device (e.g. HCAs and switches) and each HCA port. The SM may assign additional subnet unique GUIDs to an HCA port, which is useful when SR-IOV is used. A third type of address is the 128 bits Global Identifier (GID). The GID is a valid IPv6 unicast address, and at least one is assigned to each HCA port. The GID is formed by combining a globally unique 64 bits prefix assigned by the fabric administrator, and the GUID address of each HCA port.

Fat-Tree (FTree) Topologies and Routing

[0044] In accordance with an embodiment, some of the IB based HPC systems employ a fat-tree topology to take advantage of the useful properties fat-trees offer. These properties include full bisection-bandwidth and inherent fault-tolerance due to the availability of multiple paths between each source destination pair. The initial idea behind fat-trees was to employ fatter links between nodes, with more available bandwidth, as the tree moves towards the roots of the topology. The fatter links can help to avoid congestion in the upper-level switches and the bisection-bandwidth is maintained.

[0045] Figure 3 shows an illustration of a tree topology in a network environment, in accordance with an embodiment. As shown in Figure 3, one or more end-nodes 201-204 can be connected in a network fabric 200. The network fabric 200 can be based on a fat-tree topology, which includes a plurality of leaf switches 211-214, and multiple spine switches or root switches 231-234. Additionally, the network fabric 200 can include one or more intermediate switches, such as switches 221-224.

[0046] Also as shown in Figure 3, each of the end-nodes 201-204 can be a multi-homed node, i.e., a single node that is connected to two or more parts of the network fabric 200 through multiple ports. For example, the node 201 can include the ports H1 and H2, the node 202 can include the ports H3 and H4, the node 203 can include the ports H5 and H6, and the node 204 can include the ports H7 and H8.

[0047] Additionally, each switch can have multiple switch ports. For example, the root switch 231 can have the switch ports 1-2, the root switch 232 can have the switch ports 3-4, the root switch 233 can have the switch ports 5-6, and the root switch 234 can have the switch ports 7-8.

5 **[0048]** In accordance with an embodiment, the fat-tree routing mechanism is one of the most popular routing algorithm for IB based fat-tree topologies. The fat-tree routing mechanism is also implemented in the OFED (Open Fabric Enterprise Distribution – a standard software stack for building and deploying IB based applications) subnet manager, OpenSM.

10 **[0049]** The fat-tree routing mechanism aims to generate LFTs that evenly spread shortest-path routes across the links in the network fabric. The mechanism traverses the fabric in the indexing order and assigns target LIDs of the end-nodes, and thus the corresponding routes, to each switch port. For the end-nodes connected to the same leaf switch, the indexing order can depend on the switch port to which the end-node is connected (i.e., port numbering 15 sequence). For each port, the mechanism can maintain a port usage counter, and can use this port usage counter to select a least-used port each time a new route is added.

20 **[0050]** In accordance with an embodiment, in a partitioned subnet, nodes that are not members of a common data partition are not allowed to communicate. Practically, this means that some of the routes assigned by the fat-tree routing algorithm are not used for the user traffic. The problem arises when the fat tree routing mechanism generates LFTs for those routes the same way it does for the other functional paths. This behavior can result in degraded balancing on the links, as nodes are routed in the order of indexing. As routing can be performed oblivious to the data partitions, fat-tree routed subnets, in general, provide poor isolation among data partitions.

25 **[0051]** In accordance with an embodiment, a Fat-Tree is a hierarchical network topology that can scale with the available network resources. Moreover, Fat-Trees are easy to build using commodity switches placed on different levels of the hierarchy. Different variations of Fat-Trees are commonly available, including *k*-ary-*n*-trees, Extended Generalized Fat-Trees (XGFTs), Parallel Ports Generalized Fat-Trees (PGFTs) and Real Life Fat-Trees (RLFTs).

30 **[0052]** A *k*-ary-*n*-tree is an *n* level Fat-Tree with k^n end-nodes and $n \cdot k^{n-1}$ switches, each with $2k$ ports. Each switch has an equal number of up and down connections in the tree. XGFT Fat-Tree extends *k*-ary-*n*-trees by allowing both different number of up and down connections for the switches, and different number of connections at each level in the tree. The PGFT definition further broadens the XGFT topologies and permits multiple connections 35 between switches. A large variety of topologies can be defined using XGFTs and PGFTs. However, for practical purposes, RLFT, which is a restricted version of PGFT, is introduced to define Fat-Trees commonly found in today's HPC clusters. An RLFT uses the same port-

count switches at all levels in the Fat-Tree.

Input/Output (I/O) Virtualization

[0053] In accordance with an embodiment, I/O Virtualization (IOV) can provide availability

5 of I/O by allowing virtual machines (VMs) to access the underlying physical resources. The combination of storage traffic and inter-server communication impose an increased load that may overwhelm the I/O resources of a single server, leading to backlogs and idle processors as they are waiting for data. With the increase in number of I/O requests, IOV can provide availability; and can improve performance, scalability and flexibility of the (virtualized) I/O 10 resources to match the level of performance seen in modern CPU virtualization.

[0054] In accordance with an embodiment, IOV is desired as it can allow sharing of I/O

resources and provide protected access to the resources from the VMs. IOV decouples a logical device, which is exposed to a VM, from its physical implementation. Currently, there can be different types of IOV technologies, such as emulation, paravirtualization, direct 15 assignment (DA), and single root-I/O virtualization (SR-IOV).

[0055] In accordance with an embodiment, one type of IOV technology is software

emulation. Software emulation can allow for a decoupled front-end/back-end software architecture. The front-end can be a device driver placed in the VM, communicating with the back-end implemented by a hypervisor to provide I/O access. The physical device sharing 20 ratio is high and live migrations of VMs are possible with just a few milliseconds of network downtime. However, software emulation introduces additional, undesired computational overhead.

[0056] In accordance with an embodiment, another type of IOV technology is direct device

assignment. Direct device assignment involves a coupling of I/O devices to VMs, with no 25 device sharing between VMs. Direct assignment, or device passthrough, provides near to native performance with minimum overhead. The physical device bypasses the hypervisor and is directly attached to the VM. However, a downside of such direct device assignment is limited scalability, as there is no sharing among virtual machines – one physical network card is coupled with one VM.

[0057] In accordance with an embodiment, Single Root IOV (SR-IOV) can allow a physical

device to appear through hardware virtualization as multiple independent lightweight instances of the same device. These instances can be assigned to VMs as passthrough devices, and accessed as Virtual Functions (VFs). The hypervisor accesses the device through a unique (per device), fully featured Physical Function (PF). SR-IOV eases the 30 scalability issue of pure direct assignment. However, a problem presented by SR-IOV is that it can impair VM migration. Among these IOV technologies, SR-IOV can extend the PCI

Express (PCIe) specification with the means to allow direct access to a single physical 35

device from multiple VMs while maintaining near to native performance. Thus, SR-IOV can provide good performance and scalability.

[0058] SR-IOV allows a PCIe device to expose multiple virtual devices that can be shared between multiple guests by allocating one virtual device to each guest. Each SR-IOV device has at least one physical function (PF) and one or more associated virtual functions (VF). A PF is a normal PCIe function controlled by the virtual machine monitor (VMM), or hypervisor, whereas a VF is a light-weight PCIe function. Each VF has its own base address (BAR) and is assigned with a unique requester ID that enables I/O memory management unit (IOMMU) to differentiate between the traffic streams to/from different VFs. The IOMMU also applies memory and interrupt translations between the PF and the VFs.

[0059] Unfortunately, however, direct device assignment techniques pose a barrier for cloud providers in situations where transparent live migration of virtual machines is desired for data center optimization. The essence of live migration is that the memory contents of a VM are copied to a remote hypervisor. Then the VM is paused at the source hypervisor, and the VM's operation is resumed at the destination. When using software emulation methods, the network interfaces are virtual so their internal states are stored into the memory and get copied as well. Thus the downtime could be brought down to a few milliseconds.

[0060] However, migration becomes more difficult when direct device assignment techniques, such as SR-IOV, are used. In such situations, a complete internal state of the network interface cannot be copied as it is tied to the hardware. The SR-IOV VFs assigned to a VM are instead detached, the live migration will run, and a new VF will be attached at the destination. In the case of InfiniBand™ and SR-IOV, this process can introduce downtime in the order of seconds. Moreover, in an SR-IOV shared port model the addresses of the VM will change after the migration, causing additional overhead in the SM and a negative impact on the performance of the underlying network fabric.

InfiniBand™ SR-IOV Architecture – Shared Port

[0061] There can be different types of SR-IOV models, e.g. a shared port model, a virtual switch model, and a virtual port model.

[0062] **Figure 4** shows an exemplary shared port architecture, in accordance with an embodiment. As depicted in the figure, a host 300 (e.g., a host channel adapter) can interact with a hypervisor 310, which can assign the various virtual functions 330, 340, 350, to a number of virtual machines. As well, the physical function can be handled by the hypervisor 310.

[0063] In accordance with an embodiment, when using a shared port architecture, such as that depicted in Figure 4, the host, e.g., HCA, appears as a single port in the network with a single shared LID and shared Queue Pair (QP) space between the physical function 320 and

the virtual functions 330, 350, 350. However, each function (i.e., physical function and virtual functions) can have their own GID.

[0064] As shown in Figure 4, in accordance with an embodiment, different GIDs can be assigned to the virtual functions and the physical function, and the special queue pairs, QP0 and QP1 (i.e., special purpose queue pairs that are used for InfiniBand™ management packets), are owned by the physical function. These QPs are exposed to the VFs as well, but the VFs are not allowed to use QP0 (all SMPs coming from VFs towards QP0 are discarded), and QP1 can act as a proxy of the actual QP1 owned by the PF.

[0065] In accordance with an embodiment, the shared port architecture can allow for highly scalable data centers that are not limited by the number of VMs (which attach to the network by being assigned to the virtual functions), as the LID space is only consumed by physical machines and switches in the network.

[0066] However, a shortcoming of the shared port architecture is the inability to provide transparent live migration, hindering the potential for flexible VM placement. As each LID is associated with a specific hypervisor, and shared among all VMs residing on the hypervisor, a migrating VM (i.e., a virtual machine migrating to a destination hypervisor) has to have its LID changed to the LID of the destination hypervisor. Furthermore, as a consequence of the restricted QP0 access, a subnet manager cannot run inside a VM.

InfiniBand™ SR-IOV Architecture Models – Virtual Switch (vSwitch)

[0067] Figure 5 shows an exemplary vSwitch architecture, in accordance with an embodiment. As depicted in the figure, a host 400 (e.g., a host channel adapter) can interact with a hypervisor 410, which can assign the various virtual functions 430, 440, 450, to a number of virtual machines. As well, the physical function can be handled by the hypervisor 410. A virtual switch 415 can also be handled by the hypervisor 401.

[0068] In accordance with an embodiment, in a vSwitch architecture each virtual function 430, 440, 450 is a complete virtual Host Channel Adapter (vHCA), meaning that the VM assigned to a VF is assigned a complete set of IB addresses (e.g., GID, GUID, LID) and a dedicated QP space in the hardware. For the rest of the network and the SM, the HCA 400 looks like a switch, via the virtual switch 415, with additional nodes connected to it. The hypervisor 410 can use the PF 420, and the VMs (attached to the virtual functions) use the VFs.

[0069] In accordance with an embodiment, a vSwitch architecture provide transparent virtualization. However, because each virtual function is assigned a unique LID, the number of available LIDs gets consumed rapidly. As well, with many LID addresses in use (i.e., one each for each physical function and each virtual function), more communication paths have to be computed by the SM and more Subnet Management Packets (SMPs) have to be sent

to the switches in order to update their LFTs. For example, the computation of the communication paths might take several minutes in large networks. Because LID space is limited to 49151 unicast LIDs, and as each VM (via a VF), physical node, and switch occupies one LID each, the number of physical nodes and switches in the network limits the 5 number of active VMs, and vice versa.

InfiniBand™ SR-IOV Architecture Models – Virtual Port (vPort)

[0070] **Figure 6** shows an exemplary vPort concept, in accordance with an embodiment. As depicted in the figure, a host 300 (e.g., a host channel adapter) can interact with a 10 hypervisor 410, which can assign the various virtual functions 330, 340, 350, to a number of virtual machines. As well, the physical function can be handled by the hypervisor 310.

[0071] In accordance with an embodiment, the vPort concept is loosely defined in order to give freedom of implementation to vendors (e.g. the definition does not rule that the implementation has to be SRIOV specific), and a goal of the vPort is to standardize the way 15 VMs are handled in subnets. With the vPort concept, both SR-IOV Shared-Port-like and vSwitch-like architectures or a combination of both, that can be more scalable in both the space and performance domains, can be defined. A vPort supports optional LIDs, and unlike the Shared-Port, the SM is aware of all the vPorts available in a subnet even if a vPort is not using a dedicated LID.

InfiniBand™ SR-IOV Architecture Models – vSwitch with Prepopulated LIDs

[0072] In accordance with an embodiment, the present disclosure provides a system and method for providing a vSwitch architecture with prepopulated LIDs.

[0073] **Figure 7** shows an exemplary vSwitch architecture with prepopulated LIDs, in 25 accordance with an embodiment. As depicted in the figure, a number of switches 501-504 can provide communication within the network switched environment 600 (e.g., an IB subnet) between members of a fabric, such as an InfiniBand fabric. The fabric can include a number of hardware devices, such as host channel adapters 510, 520, 530. Each of the host channel adapters 510, 520, 530, can in turn interact with a hypervisor 511, 521, and 531, 30 respectively. Each hypervisor can, in turn, in conjunction with the host channel adapter it interacts with, setup and assign a number of virtual functions 514, 515, 516, 524, 525, 526, 534, 535, 536, to a number of virtual machines. For example, virtual machine 1 550 can be assigned by the hypervisor 511 to virtual function 1 514. Hypervisor 511 can additionally assign virtual machine 2 551 to virtual function 2 515, and virtual machine 3 552 to virtual 35 function 3 516. Hypervisor 531 can, in turn, assign virtual machine 4 553 to virtual function 1 534. The hypervisors can access the host channel adapters through a fully featured physical function 513, 523, 533, on each of the host channel adapters.

[0074] In accordance with an embodiment, each of the switches 501-504 can comprise a number of ports (not shown), which are used in setting a linear forwarding table in order to direct traffic within the network switched environment 600.

[0075] In accordance with an embodiment, the virtual switches 512, 522, and 532, can be

5 handled by their respective hypervisors 511, 521, 531. In such a vSwitch architecture each virtual function is a complete virtual Host Channel Adapter (vHCA), meaning that the VM assigned to a VF is assigned a complete set of IB addresses (e.g., GID, GUID, LID) and a dedicated QP space in the hardware. For the rest of the network and the SM (not shown), the HCAs 510, 520, and 530 look like a switch, via the virtual switches, with additional nodes 10 connected to them.

[0076] In accordance with an embodiment, the present disclosure provides a system and method for providing a vSwitch architecture with prepopulated LIDs. Referring to Figure 7, the LIDs are prepopulated to the various physical functions 513, 523, 533, as well as the virtual functions 514-516, 524-526, 534-536 (even those virtual functions not currently

15 associated with an active virtual machine). For example, physical function 513 is prepopulated with LID 1, while virtual function 1 534 is prepopulated with LID 10. The LIDs are prepopulated in an SR-IOV vSwitch-enabled subnet when the network is booted. Even when not all of the VFs are occupied by VMs in the network, the populated VFs are assigned with a LID as shown in Figure 7.

20 **[0077]** In accordance with an embodiment, much like physical host channel adapters can have more than one port (two ports are common for redundancy), virtual HCAs can also be represented with two ports and be connected via one, two or more virtual switches to the external IB subnet.

[0078] In accordance with an embodiment, in a vSwitch architecture with prepopulated

25 LIDs, each hypervisor can consume one LID for itself through the PF and one more LID for each additional VF. The sum of all the VFs available in all hypervisors in an IB subnet, gives the maximum amount of VMs that are allowed to run in the subnet. For example, in an IB subnet with 16 virtual functions per hypervisor in the subnet, then each hypervisor consumes 17 LIDs (one LID for each of the 16 virtual functions plus one LID for the physical function) in

30 the subnet. In such an IB subnet, the theoretical hypervisor limit for a single subnet is ruled by the number of available unicast LIDs and is: 2891 (49151 available LIDs divided by 17 LIDs per hypervisor), and the total number of VMs (i.e., the limit) is 46256 (2891 hypervisors times 16 VFs per hypervisor). (In actuality, these numbers are actually smaller since each switch, router, or dedicated SM node in the IB subnet consumes a LID as well). Note that

35 the vSwitch does not need to occupy an additional LID as it can share the LID with the PF

[0079] In accordance with an embodiment, in a vSwitch architecture with prepopulated LIDs, communication paths are computed for all the LIDs the first time the network is booted.

When a new VM needs to be started the system does not have to add a new LID in the subnet, an action that would otherwise cause a complete reconfiguration of the network, including path recalculation, which is the most time consuming part. Instead, an available port for a VM is located (i.e., an available virtual function) in one of the hypervisors and the 5 virtual machine is attached to the available virtual function.

[0080] In accordance with an embodiment, a vSwitch architecture with prepopulated LIDs also allows for the ability to calculate and use different paths to reach different VMs hosted by the same hypervisor. Essentially, this allows for such subnets and networks to use a LID Mask Control (LMC) like feature to provide alternative paths towards one physical machine, 10 without being bound by the limitation of the LMC that requires the LIDs to be sequential. The freedom to use non-sequential LIDs is particularly useful when a VM needs to be migrated and carry its associated LID to the destination.

[0081] In accordance with an embodiment, along with the benefits shown above of a vSwitch architecture with prepopulated LIDs, certain considerations can be taken into 15 account. For example, because the LIDs are prepopulated in an SR-IOV vSwitch-enabled subnet when the network is booted, the initial path computation (e.g., on boot-up) can take longer than if the LIDs were not pre-populated.

InfiniBand™ SR-IOV Architecture Models – vSwitch with Dynamic LID Assignment

[0082] In accordance with an embodiment, the present disclosure provides a system and 20 method for providing a vSwitch architecture with dynamic LID assignment.

[0083] **Figure 8** shows an exemplary vSwitch architecture with dynamic LID assignment, in accordance with an embodiment. As depicted in the figure, a number of switches 501-504 can provide communication within the network switched environment 700 (e.g., an IB subnet) 25 between members of a fabric, such as an InfiniBand fabric. The fabric can include a number of hardware devices, such as host channel adapters 510, 520, 530. Each of the host channel adapters 510, 520, 530, can in turn interact with a hypervisor 511, 521, 531, respectively. Each hypervisor can, in turn, in conjunction with the host channel adapter it interacts with, setup and assign a number of virtual functions 514, 515, 516, 524, 525, 526, 30 534, 535, 536, to a number of virtual machines. For example, virtual machine 1 550 can be assigned by the hypervisor 511 to virtual function 1 514. Hypervisor 511 can additionally assign virtual machine 2 551 to virtual function 2 515, and virtual machine 3 552 to virtual function 3 516. Hypervisor 531 can, in turn, assign virtual machine 4 553 to virtual function 1 35 534. The hypervisors can access the host channel adapters through a fully featured physical function 513, 523, 533, on each of the host channel adapters.

[0084] In accordance with an embodiment, each of the switches 501-504 can comprise a number of ports (not shown), which are used in setting a linear forwarding table in order to

direct traffic within the network switched environment 700.

[0085] In accordance with an embodiment, the virtual switches 512, 522, and 532, can be handled by their respective hypervisors 511, 521, 531. In such a vSwitch architecture each virtual function is a complete virtual Host Channel Adapter (vHCA), meaning that the VM assigned to a VF is assigned a complete set of IB addresses (e.g., GID, GUID, LID) and a dedicated QP space in the hardware. For the rest of the network and the SM (not shown), the HCAs 510, 520, and 530 look like a switch, via the virtual switches, with additional nodes connected to them.

[0086] In accordance with an embodiment, the present disclosure provides a system and method for providing a vSwitch architecture with dynamic LID assignment. Referring to Figure 8, the LIDs are dynamically assigned to the various physical functions 513, 523, 533, with physical function 513 receiving LID 1, physical function 523 receiving LID 2, and physical function 533 receiving LID 3. Those virtual functions that are associated with an active virtual machine can also receive a dynamically assigned LID. For example, because virtual machine 1 550 is active and associated with virtual function 1 514, virtual function 514 can be assigned LID 5. Likewise, virtual function 2 515, virtual function 3 516, and virtual function 1 534 are each associated with an active virtual function. Because of this, these virtual functions are assigned LIDs, with LID 7 being assigned to virtual function 2 515, LID 11 being assigned to virtual function 3 516, and LID 9 being assigned to virtual function 1 534. Unlike vSwitch with prepopulated LIDs, those virtual functions not currently associated with an active virtual machine do not receive a LID assignment.

[0087] In accordance with an embodiment, with the dynamic LID assignment, the initial path computation can be substantially reduced. When the network is booting for the first time and no VMs are present then a relatively small number of LIDs can be used for the initial path calculation and LFT distribution.

[0088] In accordance with an embodiment, much like physical host channel adapters can have more than one port (two ports are common for redundancy), virtual HCAs can also be represented with two ports and be connected via one, two or more virtual switches to the external IB subnet.

[0089] In accordance with an embodiment, when a new VM is created in a system utilizing vSwitch with dynamic LID assignment, a free VM slot is found in order to decide on which hypervisor to boot the newly added VM, and a unique non-used unicast LID is found as well. However, there are no known paths in the network and the LFTs of the switches for handling the newly added LID. Computing a new set of paths in order to handle the newly added VM is not desirable in a dynamic environment where several VMs may be booted every minute. In large IB subnets, computing a new set of routes can take several minutes, and this procedure would have to repeat each time a new VM is booted.

[0090] Advantageously, in accordance with an embodiment, because all the VFs in a hypervisor share the same uplink with the PF, there is no need to compute a new set of routes. It is only needed to iterate through the LFTs of all the physical switches in the network, copy the forwarding port from the LID entry that belongs to the PF of the hypervisor 5 —where the VM is created— to the newly added LID, and send a single SMP to update the corresponding LFT block of the particular switch. Thus the system and method avoids the need to compute a new set of routes.

[0091] In accordance with an embodiment, the LIDs assigned in the vSwitch with dynamic LID assignment architecture do not have to be sequential. When comparing the LIDs 10 assigned on VMs on each hypervisor in vSwitch with prepopulated LIDs versus vSwitch with dynamic LID assignment, it is notable that the LIDs assigned in the dynamic LID assignment architecture are non-sequential, while those prepopulated in are sequential in nature. In the vSwitch dynamic LID assignment architecture, when a new VM is created, the next available LID is used throughout the lifetime of the VM. Conversely, in a vSwitch with prepopulated 15 LIDs, each VM inherits the LID that is already assigned to the corresponding VF, and in a network without live migrations, VMs consecutively attached to a given VF get the same LID.

[0092] In accordance with an embodiment, the vSwitch with dynamic LID assignment 20 architecture can resolve the drawbacks of the vSwitch with prepopulated LIDs architecture model at a cost of some additional network and runtime SM overhead. Each time a VM is created, the LFTs of the physical switches in the subnet are updated with the newly added LID associated with the created VM. One subnet management packet (SMP) per switch is needed to be sent for this operation. The LMC-like functionality is also not available, because each VM is using the same path as its host hypervisor. However, there is no limitation on the 25 total amount of VFs present in all hypervisors, and the number of VFs may exceed that of the unicast LID limit. Of course, not all of the VFs are allowed to be attached on active VMs simultaneously if this is the case, but having more spare hypervisors and VFs adds flexibility for disaster recovery and optimization of fragmented networks when operating close to the unicast LID limit.

30 **InfiniBand™ SR-IOV Architecture Models – vSwitch with Dynamic LID Assignment and Prepopulated LIDs**

[0093] Figure 9 shows an exemplary vSwitch architecture with vSwitch with dynamic LID assignment and prepopulated LIDs, in accordance with an embodiment. As depicted in the figure, a number of switches 501-504 can provide communication within the network 35 switched environment 800 (e.g., an IB subnet) between members of a fabric, such as an InfiniBand fabric. The fabric can include a number of hardware devices, such as host channel adapters 510, 520, 530. Each of the host channel adapters 510, 520, 530, can in

turn interact with a hypervisor 511, 521, and 531, respectively. Each hypervisor can, in turn, in conjunction with the host channel adapter it interacts with, setup and assign a number of virtual functions 514, 515, 516, 524, 525, 526, 534, 535, 536, to a number of virtual machines. For example, virtual machine 1 550 can be assigned by the hypervisor 511 to 5 virtual function 1 514. Hypervisor 511 can additionally assign virtual machine 2 551 to virtual function 2 515. Hypervisor 521 can assign virtual machine 3 552 to virtual function 3 526. Hypervisor 531 can, in turn, assign virtual machine 4 553 to virtual function 2 535. The hypervisors can access the host channel adapters through a fully featured physical function 513, 523, 533, on each of the host channel adapters.

10 [0094] In accordance with an embodiment, each of the switches 501-504 can comprise a number of ports (not shown), which are used in setting a linear forwarding table in order to direct traffic within the network switched environment 800.

15 [0095] In accordance with an embodiment, the virtual switches 512, 522, and 532, can be handled by their respective hypervisors 511, 521, 531. In such a vSwitch architecture each virtual function is a complete virtual Host Channel Adapter (vHCA), meaning that the VM assigned to a VF is assigned a complete set of IB addresses (e.g., GID, GUID, LID) and a dedicated QP space in the hardware. For the rest of the network and the SM (not shown), the HCAs 510, 520, and 530 look like a switch, via the virtual switches, with additional nodes connected to them.

20 [0096] In accordance with an embodiment, the present disclosure provides a system and method for providing a hybrid vSwitch architecture with dynamic LID assignment and prepopulated LIDs. Referring to Figure 9, hypervisor 511 can be arranged with vSwitch with prepopulated LIDs architecture, while hypervisor 521 can be arranged with vSwitch with prepopulated LIDs and dynamic LID assignment. Hypervisor 531 can be arranged with 25 vSwitch with dynamic LID assignment. Thus, the physical function 513 and virtual functions 514-516 have their LIDs prepopulated (i.e., even those virtual functions not attached to an active virtual machine are assigned a LID). Physical function 523 and virtual function 1 524 can have their LIDs prepopulated, while virtual function 2 and 3, 525 and 526, have their LIDs dynamically assigned (i.e., virtual function 2 525 is available for dynamic LID assignment, and virtual function 3 526 has a LID of 11 dynamically assigned as virtual machine 3 552 is attached). Finally, the functions (physical function and virtual functions) 30 associated with hypervisor 3 531 can have their LIDs dynamically assigned. This results in virtual functions 1 and 3, 534 and 536, are available for dynamic LID assignment, while virtual function 2 535 has LID of 9 dynamically assigned as virtual machine 4 553 is attached there.

35 [0097] In accordance with an embodiment, such as that depicted in Figure 9, where both vSwitch with prepopulated LIDs and vSwitch with dynamic LID assignment are utilized

(independently or in combination within any given hypervisor), the number of prepopulated LIDs per host channel adapter can be defined by a fabric administrator and can be in the range of $0 \leq$ prepopulated VFs \leq Total VFs (per host channel adapter) , and the VFs available for dynamic LID assignment can be found by subtracting the number of 5 prepopulated VFs from the total number of VFs (per host channel adapter).

[0098] In accordance with an embodiment, much like physical host channel adapters can have more than one port (two ports are common for redundancy), virtual HCAs can also be represented with two ports and be connected via one, two or more virtual switches to the external IB subnet.

10

InfiniBand™ – Inter-Subnet Communication

[0099] In accordance with an embodiment, in addition to providing an InfiniBand™ fabric within a single subnet, embodiments of the current disclosure can also provide for an InfiniBand™ fabric that spans two or more subnets.

15

[0100] **Figure 10** shows an exemplary multi-subnet InfiniBand™ fabric, in accordance with an embodiment. As depicted in the figure, within subnet A 1000, a number of switches 1001-1004 can provide communication within subnet A 1000 (e.g., an IB subnet) between members of a fabric, such as an InfiniBand™ fabric. The fabric can include a number of hardware devices, such as, for example, channel adapter 1010. Host channel adapters 1010 can in turn interact with a hypervisor 1011. The hypervisor can, in turn, in conjunction with the host channel adapter it interacts with, setup a number of virtual functions 1014. The hypervisor can additionally assign virtual machines to each of the virtual functions, such as virtual machine 1 1015 being assigned to virtual function 1 1014. The hypervisor can access their associated host channel adapters through a fully featured physical function, such as physical function 1013, on each of the host channel adapters.

20

[0101] With further reference to Figure 10, and in accordance with an embodiment, a number of switches 1021-1024 can provide communication within subnet B 1040 (e.g., an IB subnet) between members of a fabric, such as an InfiniBand™ fabric. The fabric can include a number of hardware devices, such as, for example, host channel adapter 1030. Host channel adapter 1030 can in turn interact with a hypervisor 1031. The hypervisor can, in turn, in conjunction with the host channel adapter it interacts with, setup a number of virtual functions 1034. The hypervisors can additionally assign virtual machines to each of the virtual functions, such as virtual machine 2 1035 being assigned to virtual function 2 1034. The hypervisor can access their associated host channel adapters through a fully featured physical function, such as physical function 1033, on each of the host channel adapters. It is noted that, although only one host channel adapter is shown within each subnet (i.e., subnet 25 30 35

A and subnet B), it is to be understood that a plurality of host channel adapters, and their corresponding components, can be included within each subnet.

[0102] In accordance with an embodiment, each of the host channel adapters can additionally be associated with a virtual switch, such as virtual switch 1012 and virtual switch

5 1032, and each HCA can be set up with a different architecture model, as discussed above. Although both subnets within Figure 10 are shown as using a vSwitch with prepopulated LID architecture model, this is not meant to imply that all such subnet configurations must follow a similar architecture model.

[0103] In accordance with an embodiment, at least one switch within each subnet can be 10 associated with a router, such as switch 1002 within subnet A 1000 being associated with router 1005, and switch 1021 within subnet B 1040 being associated with router 1006.

[0104] In accordance with an embodiment, when traffic at an originating source, such as 15 virtual machine 1 within subnet A, is addressed to a destination at a different subnet, such as virtual machine 2 within subnet B, the traffic can be addressed to the router within subnet A,

i.e., router 1005, which can then pass the traffic to subnet B via its link with router 1006.

Fabric Manager

[0105] As discussed above, a network fabric, such as an InfiniBand™ fabric, can span a plurality of subnets through the use of interconnected routers in each subnet of the fabric. In

20 accordance with an embodiment, a fabric manager (not shown) can be implemented on a host which is a member of the network fabric and can be employed within the fabric to manage both physical and logical resources that are part of the fabric. For instance, management tasks such as discovering fabric resources, controlling connectivity between physical servers, collecting and viewing real-time network statistics, disaster recovery, and 25 setting quality of service (QoS) settings, among others, may be performed by a user through the fabric manager. In accordance with an embodiment, the fabric manager may span all subnets defined in the fabric. That is, the fabric manager can manage physical and logical resources that are members of, or associated with, the fabric at large, regardless of which subnet the resources are a member of.

[0106] In accordance with an embodiment, the fabric manager can include a graphical user 30 interface (GUI) through which a user can perform administration functions. The fabric manager GUI can incorporate visualization tools that allow a user to monitor and control fabric resources. For example, in an embodiment, a user can view server connections, configuration settings and performance statistics for servers across the fabric through the fabric interface. Other examples of fabric functionality that can be monitored and/or managed through the fabric manager GUI include discovering inter-subnet fabric topology, viewing 35 visual representations of these topologies, creating fabric profiles (e.g., virtual machine fabric

profiles), and building and management of a fabric manager database that can store fabric profiles, metadata, configuration settings and other data required by, and related to, the network fabric. In accordance with an embodiment, the fabric manager database is a fabric-level database.

5 [0107] In addition, the fabric manager can define legal inter-subnet connectivity in terms of which subnets are allowed to communicate via which router ports using which partition numbers. In accordance with an embodiment, the fabric manager is a centralized fabric management utility. The above examples are not meant to be limiting.

10 [0108] In accordance with an embodiment, some of the fabric manager's functionality can be initiated by a user, and other functionality can be abstracted from the user, or be automated (e.g., some functionality may be performed by the fabric manager upon startup, or at other predetermined events).

[0109] In an exemplary embodiment of a management event, a user may initiate, at the fabric manager interface, a configuration change directed towards a network fabric device.

15 After receiving the configuration change request, the fabric manager may, in turn, ensure that the configuration change request is properly carried out. For example, the fabric manager may communicate the request to the device and ensure that the configuration change is written to the device's configuration. In one embodiment, the physical device acknowledges to the fabric manager that the configuration change has successfully completed. In 20 accordance with an embodiment, the fabric manager may then update the interface to give a visual confirmation that the request has been carried out. Further, the fabric manager may persist the configuration of the device to the fabric manager database, e.g., for disaster recovery or other purposes.

[0110] In accordance with an embodiment, the fabric manager can have other interfaces, 25 such as a command line interface, that includes some, all, or more functionality than the GUI.

Fabric-Level Resource Domains

[0111] As discussed above, a fabric manager can allow users to perform administrative tasks throughout the network fabric through an interface of the fabric manager. In 30 accordance with an embodiment, an additional function of the fabric manager is facilitation of hierarchical role-based access control. In an embodiment, role-based access control is achieved through fabric-level resource domains.

[0112] In accordance with an embodiment, role-based access control is based on the concept of fabric users. Access from both human administrators and external management 35 applications can represent an authenticated context that defines legal operations on all or a subset of the fabric infrastructure or fabric resources. For example, a user can be represented in the fabric by a user profile. That is, within the fabric a user can be defined by

creating a profile of the user and assigning attributes to the profile. A user profile can be assigned a username attribute, and a password attribute, where the username is unique within the fabric, thereby uniquely identifying the user. Further, the user profile may be associated with certain roles defined in the fabric that assign certain access levels to different resources within the fabric. In accordance with an embodiment, setting up user profiles can be accomplished through the fabric manager interface. All or part of the user profile can be stored in the fabric manager database. Moreover, in an embodiment, the fabric manager can integrate with well-known user directories, such as Microsoft's® Active Directory or LDAP directories, or with, e.g., the RADIUS networking protocol for remote authentication.

5 [0113] In accordance with an embodiment, the fabric manager can manage fabric resources that it discovers through fabric-level resource domains (also referred to as "resource domains", or simply "domains" herein). A resource domain is a logical grouping of fabric resources defined at the fabric level. Fabric resources include both physical and logical resources. Some examples of resources include fabric devices (such as HCAs, physical 10 nodes, and switches), fabric profiles (such as virtual machine fabric profiles, and user profiles), virtual machines, clouds, and I/O modules, among others.

15 [0114] In accordance with an embodiment, all fabric resources discovered and managed by the fabric manager reside in the default domain, which exists by default (i.e., without the need to setup or configure it) in the fabric, and can be accessed through the fabric manager 20 interface. The default domain is the highest level domain – that is, it is the parent domain to all other resource domains, and all other resource domains exist within the default domain. The default domain is associated with a fabric-level administrator, which also exists by default, and is configured with administrative privileges in the default domain by default.

25 [0115] In accordance with an embodiment, resource domains represent a hierarchical form of resource management. For example, the process of configuring and managing the default domain is available only to the fabric-level administrator. However, child domains can be created within the default domain by the fabric-level administrator. For instance, the fabric-level administrator can create a child domain and can add domain resources to the child domain. Additionally the fabric-level administrator can create domain-level "domain admin" 30 users and add (i.e., associate) the domain admin users to the child domain. Making the domain admin user a member of the resource domain allows the domain admin user to manage the child domain and its contained subset of the fabric resources. In accordance with an embodiment, the domain admin user cannot manage resources outside of the child domain (i.e., resources at a parallel or a higher level than the domain admin is associated 35 with). However the domain admin can manage resources contained in resource domains that have been created as child domains of the resource domain. In accordance with an

embodiment, the fabric manager is responsible for providing the security that ensures that resource domain boundaries are strictly enforced.

[0116] Figure 11 shows the hierarchical structure of resource domains. As shown, default domain 1102 exists within network fabric 1100. Fabric-level administrator 1110 has access rights to manage fabric-level resources 1112, 1124, and 1134. Fabric-level administrator 1110 can also create and administer new resource domains within default domain 1102. Fabric-level administrator 1110 has created resource domain 1120 and 1130, and corresponding domain-level domain admin users 1122 and 1132. Domain admin user 1122 has access rights to manage fabric resources 1124 (assigned to resource domain 1120 by fabric-level administrator 1110), but has no access rights to manage fabric resources 1112 (at a higher level) or domain resources 1134 (at a parallel level). Likewise, Domain admin user 1132 has access rights to manage fabric resources 1134 (assigned to resource domain 1130 by Fabric-level administrator 1110), but has no access rights to manage fabric resources 1112 (at a higher level) or domain resources 1124 (at a parallel level).

15

Admin Partitions

[0117] In accordance with an embodiment, a resource domain can be represented at the subnet level by an administration, or “admin” partition (as they are referred to herein). An admin partition represents a group membership which grants access rights at the subnet level to subnet resources. Members of an admin partition are considered privileged, in that the members have access rights to any subnet resources that are associated with the admin partition, in accordance with an embodiment. At the fabric manager level, an admin partition is associated with a resource domain and a corresponding domain admin user. Thus, user-role separation can be ensured in multi-tenant environments at the subnet level. Further, resource domain membership can be correlated with admin partition membership, so that resources that are members of an admin partition that is associated with a particular resource domain are also members of the resource domain.

[0118] In accordance with an embodiment, an admin partition can be defined at the subnet level in the same way that a data partition is defined, but with an additional attribute specifying that the partition being created is an admin partition. Like data partitions (discussed in detail, above), admin partitions can be created by an administrator through the fabric manager interface, in accordance with an embodiment. In an embodiment, the fabric manager can support an “admin partition” flag as an optional parameter during the creation of a partition. If selected by the creating administrator, the fabric manager will include the additional attribute specifying that the newly created partition is an admin partition, and will be treated as an admin partition by the fabric manager and the local master subnet manager.

[0119] In accordance with an embodiment, the fabric manager can be configured to automatically create a corresponding admin partition for each resource domain that is created, and associate the automatically created partition with the corresponding resource domain. In such an embodiment, when fabric-level resources are added to the resource domain, the fabric manager also associates them with the admin partition that was automatically created and associated with the resource domain. Thus, resources added to the resource domain will have subnet-level access rights to each other upon being added to the resource domain, with no further action being taken by the administrator (e.g., the fabric-level administrator or the domain admin).

[0120] Moreover, in accordance with an embodiment, entire subnets of the network can represent a special resource domain in a domain hierarchy that has a top-level domain (e.g., the default domain). For instance, in a domain hierarchy, where the default domain represents the top-level domain, each subnet of the network fabric can then be recognized by the fabric manager as a child domain of the default domain. Recognition of entire subnets as child domains of a top-level domain can be configured as default behavior of the fabric manager, or these default domains can be manually defined by an administrator. Here again, in order to have user role separation and enforcement of domain boundaries and resource associations at the subnet level, admin partitions corresponding to entire-subnet resource domains can be defined. In accordance with an embodiment, an admin partitions that is defined in a subnet and includes each resource in that subnet (as either a member, or associated with the admin partition) can be termed a "domain global" admin partition, since in this configuration, every resource in the subnet would have access rights to every other resource.

[0121] In accordance with an embodiment, an admin partition can be transparent to a domain admin. As noted above, a domain global admin partition can be created automatically for a resource domain at the fabric manager level, and then all resources assigned to or created within the scope of this domain can automatically be associated with the corresponding admin partition. In another embodiment, however, the domain admin can explicitly create different admin partitions within the relevant resource domain, and then resources within the domain can be explicitly associated with the explicitly created admin partition instead of with the admin partition that was created by default for the resource domain.

[0122] In accordance with an embodiment, the fabric manager can support the creation of both shared and private admin partitions. Admin partitions created by a fabric-level administrator in the default domain can be shared partitions that can be made available to individual resource domains. Admin partitions created by a domain admin (i.e., a user with credentials associated with a specific resource domain) in the domain in which the domain

admin is a member can be private partitions associated with and available only to the specific resource domain in whose context the admin partitions were created.

[0123] In accordance with an embodiment, end-ports of HCAs and vHCAs can be members of an admin partition, just as they can be members of a data partition. Admin

5 partitions are differentiated from data partitions, however, in that admin partitions can be associated with other subnet resources, in accordance with an embodiment. For example, a data partition can be associated with an admin partition. Further, an admin partition can be associated with another admin partition, as a child or as a parent, thus making admin partitions a hierarchical concept and able to correspond with the hierarchy of the resource

10 domains they are associated with, in accordance with an embodiment.

[0124] As a technical matter, end-ports of HCAs (and vHCAs) can be referred to as “members” of partitions, in traditional terminology, and other resources can be “associated with” admin partitions, in accordance with an embodiment. The technical differences in these

15 two concepts are explained below. For convenience and readability, however, this document may occasionally, in reference to admin partitions, use the terms “member” and “associated with” interchangeably. Notwithstanding the use of these terms interchangeably, it is to be understood that the technical differences between end-port/HCA membership in, and resource association with, admin partitions is meant to be consistently applied by the reader.

[0125] In accordance with an embodiment, an admin partition is defined by a P_Key, just

20 as a data partition is defined. However, while an end-port is aware of the data partitions that it is a member of, it is not necessary that end-ports be aware of what admin partitions they are members of. Thus, in one embodiment, a P_Key defining an admin partition is not entered in the P_Key table of member end-ports. In this way the creation of an admin partition does not waste P_Key table entries – which are a limited resource – if an admin partition is not used for IB packet traffic. In another embodiment, however, an admin partition may function as both an admin partition and a data partition. In such an embodiment, all P_Key tables of end-ports that are members of the admin partition can have a P_Key entry for the admin partition in their respective P_Key tables. In accordance with an embodiment, a data partition may be defined as any partition that is not also an admin partition.

25 [0126] In accordance with an embodiment, a data partition can be associated with one or more admin partitions. For example, a data partition, being defined by a P_Key value, can be associated with an admin partition that is defined by its own distinct P_Key value. Moreover, the data partition can be associated with a second admin partition defined by yet another distinct P_Key value. In accordance with an embodiment, the association of a data partition with a specific admin partition can define the maximum membership level for end-ports that are members of the specific admin partition.

[0127] As noted above, an admin partition represents a group membership which grants access rights to subnet resources. In accordance with an embodiment, any end-port member of an admin partition has access rights to any subnet resource that is associated with the same admin partition based solely on the end-port's membership in the admin partition.

5 Thus, any end-port that is a member of an admin partition has access rights to any data partition that is associated with that same admin partition. Notably, this does not necessarily mean that the member end-port is a member of the associated data partition, but that it has access rights to the associated data partition, and therefore could be a member of the data partition.

10 **[0128]** Such a scheme obviates the need for administrators to grant end-ports access to, e.g., data partitions by manually including the data partition's P_Key in the P_Key table of the end-port. In an embodiment, when an end-port is initialized in the subnet, the master subnet manager can query a data store (e.g., an admin partition registry, as discussed below) that holds admin partition definitions (e.g., P_Keys), and relationships that define membership in 15 the defined admin partitions and that define associations with the defined admin partitions, to determine which admin partitions the end-port is a member of. The subnet manager can then further check to see if there are any data partitions associated with the admin partitions of which the end-port is a member. If the SM finds that 1) the end-port is a member of an admin partition, and 2) that that admin partition is associated with a data partition, then the SM can 20 automatically place the P_Key of the associated data partition in the P_Key table of the end-port, thereby automatically granting the end-port access to the data partition. Thus, the admin partition represents a simpler, more scalable solution than manual partition mapping by administrators.

25 **[0129]** **Figure 12** shows an exemplary network fabric having both admin partitions and data partitions. As shown in Figure 12, admin partitions 1230, 1240, and 1250 have been defined within the fabric. Nodes A-E 1201-1205, are physically connected to the fabric by their respective HCAs 1211-1215. Additionally, each HCA is a member of at least one admin partition. HCA 1211 and HCA 1214 are members of admin partition 1230. HCA 1211 is also 30 a member of admin partition 1240, along with HCAs 1212 and 1213. HCA 1213 is, additionally, a member of admin partition 1250, along with HCA 1215.

[0130] With further reference to Figure 12, and in accordance with an embodiment, data partitions 1232, 1242, and 1252 have been defined within the fabric. Data partition 1232 is associated with admin partition 1230, data partition 1242 is associated with admin partition 1240, and data partition 1252 is associated with admin partition 1250. In accordance with an 35 embodiment, HCA 1211 and HCA 1214 have access rights to membership in data partition 1232 based on their membership in admin partition 1230. Likewise, HCAs 1211-1213 have access rights to membership in data partition 1242 based on their membership in admin

partition 1240. Moreover, HCAs 1213 and 1215 have access rights to membership in data partition 1252 based on their membership in admin partition 1250.

[0131] In accordance with an embodiment, admin partitions can also be used to determine whether a vHCA can be registered with the virtual function of a physical HCA. A vHCA describes a host channel adapter which is planned and configured for a specific virtual machine (VM), in accordance with an embodiment. A vHCA differs from a virtual function (VF) in that a vHCA migrates with a VM, while a VF stays with the physical adapter. As discussed above, however, both physical HCAs and vHCAs (and, at a lower level, the end-ports of these (v)HCAs) can be members of admin partitions. Thus, in accordance with an embodiment, admin partition membership can be used by the SM to determine whether a request from a physical HCA to register a vHCA with a virtual function of the requesting physical HCA is permissible.

[0132] Figure 13 shows an exemplary network fabric having HCAs and vHCAs as members of admin partitions. As shown in Figure 13, subnet 1302 is part of network fabric 1300. HCA 1310, 1324, 1332, and 1344 represent physical HCAs physically connected through their respective end-ports to network fabric 1300 in subnet 1302. HCA 1310 is associated with physical function (PF) 1312 and with virtual functions (VFs) 1314 and 1316. HCA 1324 is associated with PF 1326 and with VFs 1328 and 1329. HCA 1332 is associated with PF 1334 and with VFs 1336 and 1338. HCA 1344 is associated with PF 1346 and with VFs 1348 and 1349. Further, vHCA 1320 is depicted as registered with VF 1314, and associated with Virtual machine (VM) 1318 (i.e., VM 1318 obtains access to network fabric 1300 through vHCA 1320, and ultimately through physical HCA 1310). vHCA 1340 is registered VF 1337, and associated with VM 1338.

[0133] With continued reference to Figure 13, as shown, HCAs 1310 and 1324, and vHCA 1320 are members of admin partition 1350. Additionally, HCA 1332 and 1344, and vHCA 1340 are members of admin partition 1360. Consequently, vHCA 1320 can be legally registered with VF 1314 or 1316 of HCA 1310, or with VF 1328 or 1329 of HCA 1324, by virtue of the fact that HCA 1310 and 1324, and vHCA 1320 are each members of admin partition 1350. Similarly, vHCA 1340 can be legally registered with VF 1336 or 1338 of HCA 1330, or with VF 1348 or 1349 of HCA 1344, by virtue of the fact that HCA 1332 and 1324, and vHCA 1340 are each members of admin partition 1360.

[0134] As noted above, the fabric-level fabric database holds information related to the fabric and fabric resources, and is managed by the fabric manager. In accordance with an embodiment, the fabric database can have “complete knowledge” of the fabric resource inventory (i.e., every resource that is a part of the network fabric is represented, at least, by a record held in the fabric database). Further, the access rights and name spaces associated

with each resource in the fabric can be either stored in the fabric database, or derived from information and relationships contained in the fabric database.

[0135] For example, in accordance with an embodiment, information pertaining to admin partition membership and/or resource association with an admin partition can be stored in the fabric database. The tables holding this information and the relationships that link these tables together can be a subset of the fabric database, and can be referred to as the admin partition registry. In accordance with an embodiment, the admin partition registry is a collection of admin partition group resources. For example, an admin partition group within the admin partition registry can be a collection of HCA members (including vHCAs) and associated resources of a particular admin partition, where the group is looked up by the P_Key that defines the particular admin partition. Moreover, admin partition group members and associated resources can be looked up in the registry using keys such as GUID or vGUID for member HCAs or vHCAs, respectively, or P_Keys for associated data partitions. Relationships between the P_Key of an admin partition and the unique identifier of members or associated resources define membership or association, respectively, in the admin partition, and are maintained by the admin partition registry, and by the fabric database, at a higher level.

[0136] In accordance with an embodiment, all or part of the admin partition registry may be held as records in a cache of the SM. For instance, records of the admin partition registry that correspond to resources of a particular subnet can be duplicated in a cache in a resident memory of a subnet manager (e.g., the master subnet manager) of the particular subnet. The admin partition registry records can either be retrieved (i.e., copied) from the fabric database by the SM (e.g., when the SM boots), or be placed in the cache before it is persisted to the fabric database. The cache can be a volatile or non-volatile memory. Regardless of when the registry records are placed in the cache, synchronization can then occur between the cached copy of the admin partition registry and the copy of the admin partition registry found in a fabric-level data base.

[0137] By holding all, or a subnet-relevant part, of the admin partition registry in a high-speed cache on the SM, rather than retrieving admin partition information from a persisted state (i.e., from the fabric database) every time a query is received, the lookup of admin partition information can impose minimal overhead on the SM. This can be especially important during subnet initialization, when access rights among subnet resources are being automatically assigned.

In accordance with an embodiment, logical names or identifiers can be assigned to resources within a resource domain (by, e.g., the fabric-level or domain-level admin user). These logical names can be private to the resource domain. The fabric manager, through the

fabric database, can create relationships that map unique identifiers used within the fabric (e.g., vGUIDs and P_Keys) to logical or symbolic names given to resources within the fabric.

[0138] For instance, the fabric database, in accordance with an embodiment, can store records of resources, and domain membership and/or admin partition membership of

5 resources. Logical names can be assigned to the resources upon discovery of the resources by the fabric manager. These names can be linked to the unique identifiers of the fabric

resources in the fabric database. Moreover, the fabric manager can keep track of each resource's membership in resource domains and admin partitions through a relationship in the fabric manager database. With these records and relationships, the fabric manager can

10 allow like logical names across disparate resource domains and admin partitions. In accordance with an embodiment, the logical domain name scheme can reflect the hierarchy of the resource domain or domains that a particular domain resource is a member of. In such

an embodiment, logical resource names can be unique to the highest level resource domain that the resource is a member of.

15 [0139] In accordance with an embodiment, the identifier of a resource in the fabric –

regardless of what the identifier is – can be unique within the scope of the admin partition. Then, global uniqueness (i.e., at the fabric level) can be achieved by prefixing the resource

name with the corresponding admin partition.

[0140] **Figure 14** shows an exemplary network fabric having both resource domains and

20 admin partitions. As shown in Figure 14, fabric manager 1402 is executing on network fabric 1400. In accordance with an embodiment, fabric manager 1402 can execute from a node

(not shown) of network fabric 1400. Fabric manager 1402 is administered by fabric-level

administrator 1404, and includes fabric manager database 1414. Admin partition registry

1416 is part of fabric manager database 1414, as is a logical name table 1418.

25 [0141] With continued reference to Figure 14, subnet 1420 is defined within network fabric

1400. Subnet manager 1422 is associated with subnet 1420, and, in accordance with an

embodiment, performs the semantic runtime operations required by subnet 1420 for

operation in network fabric 1400. Setup and administrative tasks required by subnet 1420

can be performed by fabric-level administrator 1404 and fabric manager 1402.

30 [0142] Node 1444, 1454, 1474 and 1484 are part of subnet 1420. HCA 1446 is associated

with node 1444, and includes PF 1448 and VFs 1450 and 1452. Similarly, HCA 1456 is

virtual function 1452 of HCA 1446. vHCA 1472 has been planned and configured for VM 1470, is associated with VM 1470, and is registered with virtual function 1482 of HCA 1476.

[0143] In accordance with an embodiment, HCAs 1446, 1456, 1476, and 1486 are considered domain resources, and a record of each is stored in fabric manager database

1414. The record can include an identifier, such as a GUID, which is used to identify the HCA resource in the fabric. Further, vHCAs 1442 and 1472 are also considered domain resources, and a record of each is stored in fabric manager database 1414. The record can include an identifier, such as a GUID, which is used to identify the vHCA.

[0144] With further reference to Figure 14, and in accordance with an embodiment,

10 resource domain 1410 and resource domain 1412 have been created within fabric manager 1402. In accordance with an embodiment, fabric-level administrator 1404 is responsible for the creation of resource domain 1410 and resource domain 1412. Additionally, domain admin 1406 is a domain-level administrator associated with resource domain 1410. Likewise, domain admin 1408 is a domain-level administrator associated with resource domain 1412.

15 In accordance with an embodiment, fabric-level administrator 1404 can create domain admins 1406 and 1408, as admins of their respective resource domains, adhering to the hierarchical nature of resource domains.

[0145] In accordance with an embodiment, admin partition 1424 and admin partition 1426

20 have been defined in subnet 1420. Admin partition 1424 is associated with resource domain 1410, and admin partition 1426 is associated with resource domain 1412.

[0146] As shown in Figure 14, vHCA 1442 and HCAs 1446 and 1456 are members of

resource domain 1410. In accordance with an embodiment, because admin partition 1424 is associated with resource domain 1410, when vHCA 1442 and HCAs 1446 and 1456 are added as members of resource domain 1410, they also become members of admin partition 1424, and a relationship is created in admin partition registry 1416 between the P_Key defining admin partition 1424 and the identifiers of HCAs 1446 and 1456, and vHCA 1442. In accordance with an embodiment, this relationship defines HCAs 1446 and 1456, and vHCA 1442 as members of admin partition 1424.

[0147] Likewise, vHCA 1472 and HCAs 1476 and 1486 are members of resource domain

30 1412. In accordance with an embodiment, because admin partition 1426 is associated with resource domain 1410, when vHCA 1472 and HCAs 1466 and 1486 are added as members of resource domain 1412, they also become members of admin partition 1426, and a relationship is created in admin partition registry 1416 between the P_Key defining admin partition 1426 and the identifiers of HCAs 1476 and 1486, and vHCA 1472. In accordance 35 with an embodiment, this relationship defines HCAs 1476 and 1486, and vHCA 1472 as members of admin partition 1426.

[0148] As noted above, VM 1440 (including vHCA 1442), node 1444 (including HCA 1446) and node 1454 (including HCA 1456) are members of resource domain 1410, in accordance with an embodiment. In an embodiment of the invention, fabric-level administrator 1404 is responsible for adding node 1444 and node 1454 to resource domain 1410. For example, 5 fabric-level administrator 1404 can, through the interface of fabric manager 1402, add nodes 1444 and 1454 to resource domain 1410. Once fabric-level administrator 1404 has added nodes 1444 and 1454 to resource domain 1410, domain admin 1406 can perform administrative tasks on nodes 1444 and 1454. In keeping with the hierarchical scheme of 10 resource domains, however, domain admin 1406 could not perform administrative tasks on nodes 1444 and 1454 before they were added to resource domain 1410 (i.e., while they were a member of the higher-level default domain (not shown)). Further, in accordance with an embodiment, domain admin 1408 cannot perform administrative tasks on nodes 1444 and 1454, because nodes 1444 and 1454 are members of a parallel-level resource domain that domain admin 1408 is not associated with.

15 [0149] With continued reference to Figure 14, and in accordance with an embodiment, admin partitions 1424 and 1426 have been defined within subnet 1420. In keeping with the hierarchical scheme of resource domains, in one embodiment admin partitions 1424 and 1426 were defined by fabric-level administrator 1404. In another embodiment, domain admin 1406 defined admin partition 1424, and domain admin 1408 defined admin partition 1426. In 20 accordance with an embodiment, admin partition 1424 is associated with resource domain 1410, and admin partition 1426 is associated with resource domain 1412. As discussed above, admin partitions 1424 and 1426 represent resource domains 1410 and 1412, respectively, at the subnet level, in accordance with an embodiment. In addition to being associated with their respective resource domains, admin partitions 1424 and 1426 are 25 associated with domain admins 1406 and 1408, respectively (i.e., the corresponding admin user of the resource domains each of the admin partitions is associated with). As noted above, this association between admin partitions and domain-level admins can ensure user-role separation in multi-tenant environments at the subnet level, in accordance with an embodiment.

30 [0150] Data partitions 1428 and 1430 have been defined in subnet 1420, in accordance with an embodiment. In keeping with the hierarchical scheme of resource domains, in one embodiment data partitions 1428 and 1430 were defined by fabric-level administrator 1404. In another embodiment, domain admin 1406 defined data partition 1428, and domain admin 1408 defined data partition 1430. As shown in figure 14, data partition 1428 is associated 35 with admin partition 1424, and data partition 1430 is associated with admin partition 1426. Moreover, as noted above and shown in Figure 14, HCAs 1446 and 1456 and vHCA 1442 are members of admin partition 1424. Consequently, in accordance with an embodiment,

HCA 1446 and 1456 and vHCA 1442 have access permissions to data partition 1428 because they are members of an admin partition (i.e., admin partition 1424) that data partition 1428 is associated with.

[0151] In accordance with an embodiment, when data partition 1428 is associated with admin partition 1424, a relationship between the identifier of data partition 1428 (e.g., the P_Key of data partition 1428) and the P_Key of admin partition 1424 is created in the admin partition registry 1416. This relationship defines data partition 1428 as associated with admin partition 1424. Likewise, when data partition 1430 is associated with admin partition 1426 a relationship between the identifier of data partition 1430 (e.g., the P_Key of data partition 1430) and the P_Key of admin partition 1426 is created in the admin partition registry 1416. This relationship defines data partition 1430 as associated with admin partition 1426.

[0152] In accordance with an embodiment, if a request were received from either of HCAs 1446 and 1456 or vHCA 1442 to join data partition 1428, SM 1422 could check with admin partition registry 1416, and find that HCAs 1446 and 1456 and vHCA 1442 are members of admin partition 1424, and that data partition 1428 is associated with admin partition 1424. Then, SM 1422 could allow HCAs 1446 and 1456 and vHCA 1442 to become members of data partition 1428 based on HCAs 1446 and 1456 and vHCA 1442 being members of admin partition 1424 and data partition 1428 being associated with admin partition 1424. No manual mapping from either fabric-level administrator 1404 or domain-level administrator 1406 would be necessary to allow HCAs 1446 and 1456 and vHCA 1442 to join data partition 1428.

[0153] Moreover, vHCA 1442 can be registered with either of VF 1452 or 1450 of HCA 1446, or either of VF 1462 or 1460 of HCA 1456, because HCAs 1446 and 1456 and vHCA 1442 are members of admin partition 1424 (vHCA 1442 is depicted as registered with VF 1452). Here again, SM 1422 could check with admin partition registry 1416, and find that HCAs 1446 and 1456 and vHCA 1442 are members of admin partition 1424. Upon finding that HCAs 1446 and 1456 and vHCA 1442 are members of admin partition 1424, SM 1422 could allow registration of vHCA 1442 with any of VFs 1452, 1450 1462, and 1460 without intervention from any fabric user.

30 Virtual Machine Fabric Profiles

[0154] As discussed above, virtual machines (VMs) can be employed in a fabric, such as an IB fabric in order to improve efficient hardware resource utilization and scalability. Yet, live migration of virtual machines (VMs), still remains an issue due to the addressing and routing schemes used in these solutions. In accordance with an embodiment, methods and systems provide for facilitating pre-defined, highly available, and physical-location independent virtual machine fabric profiles that can support addressing schemes aimed at overcoming such VM migration issues. In accordance with an embodiment, VM fabric profiles enable centralized

setup and configuration administration for VMs using fabric connectivity, and support optimized VM migration for VMs based on SR-IOV. In accordance with an embodiment, a VM fabric profile represents a single, centralized repository of detailed fabric configuration information for a virtual machine. A database associated with the fabric manager (e.g., the fabric database) can persist the information that makes up a VM fabric profile.

5 [0155] In accordance with an embodiment, a VM fabric profile can be identified in a network fabric, such as an IB fabric, through a virtual machine identifier (VM-id). In one embodiment, the VM-id is a unique 128-bit number that is a universally unique identifier (UUID), that can be unique across the entire fabric. However, uniqueness of the VM-id is only necessary across differently administered VM manager domains (e.g., a VM-id can be unique within an admin partition). Therefore, in other embodiments, the VM-id can be some other appropriate type of ID that is at least unique across such domains. In accordance with an embodiment, all management entities, at either the fabric or the subnet level, look up information about a VM fabric profile by referencing the VM-id of the profile.

10 15 [0156] **Figure 15** shows an exemplary database structure for storing VM fabric profile information. Figure 15 depicts several tables in a traditional relational database design. However, any suitable data structure can be used to store VM fabric profile data (e.g., a flat-file table, or a delimited structure, etc.). In Figure 15, an asterisk (*) denotes a key value. Figure 15 depicts VM fabric profile data as being part of a larger fabric database 1500, but in other embodiments, VM profile data may be contained in its own database, or may be a 20 separate database with access to fabric database 1500.

25 [0157] As shown in Figure 15, the contents of a VM fabric profile may include, but are not limited to: a virtual machine identifier (VM-id) 1502 used as a lookup key; a logical name 1504 for ease of use and improved quality in administration of the fabric; a profile type 1506 used to distinguish between, e.g., a VM fabric profile and other profiles that have been defined for the fabric; a profile ID 1508 – a unique id within the set of all profiles defined for the fabric; and a content update enumerator (CUE) 1510, which can be a sequence number for a profile where the highest sequence represents the most recent update. As shown in Figure 15, this VM fabric profile content can be stored in a VM profile table 1512, where the 30 VM-id acts as a unique key for identifying each VM fabric profile.

35 [0158] As discussed above, virtual HCAs (vHCAs) may be used in conjunction with the virtual functions of a physical HCA to provide network access to VMs. In accordance with an embodiment, each VM fabric profile is also associated with at least one vHCA. A vHCA can be planned and configured for a specific VM, and this configuration can also be included in and stored with the fabric profile of the VM for which the vHCA is configured. The configured vHCA can then migrate with the virtual machine – in contrast to virtual functions, which can be defined by, and can stay with, the physical HCA. With further reference to Figure 15, a

vHCA can be represented in the fabric manager database as a unique combination of a VM-id 1502 and a vHCA Instance ID 1514. Moreover, this combination can be stored in a vHCA table related to the VM Profile table 1512 through the VM-id key 1502.

[0159] Like a physical HCA, a vHCA can have a plurality of (virtual) ports. In accordance with an embodiment, these virtual ports can act as end-ports in the network environment, just as physical ports do. As an example, all end-ports, including vHCA ports, can be assigned a GUID (e.g., a 64-bit GUID as used in an IB network). This GUID can be used to request a LID destination address from the routing tables of a SM. In accordance with an embodiment, a virtual GUID (vGUID) can represent the current fabric address of each vHCA port. In one embodiment, vGUIDs can be assigned to the vHCA ports from a list of GUIDs allocated to, and stored with, the VM fabric profile, as discussed above. vGUIDs can be assigned to a fabric profile from a dedicated pool of GUIDs owned and controlled by the fabric manager in accordance with fabric manager GUID policy. For instance, a free and fabric-wide unique vGUID can be allocated to a vHCA configured for a VM when the fabric profile is created in the fabric manager.

[0160] With continued reference to Figure 15, and in accordance with an embodiment, a vHCA port can be represented in the fabric manager database as the unique combination of a vHCA Instance ID 1514, and a vGUID 1522. A vHCA configuration can also include a vHCA port number 1520. This configuration can be stored in the vHCA port table 1518, which can be related the vHCA table through the vHCA Instance ID* 1514 key, and ultimately to the VM profile table 1512 through the VM-id key 1502.

[0161] In accordance with an embodiment, a vHCA can be a member of both data partitions and admin partitions. Partition membership of a vHCA can be represented in a VM fabric profile by relationships in the fabric database linking the vHCA record to partitions (both admin and data partitions) that the vHCA is a member of. In one embodiment, for example, the vGUID key 1522 can be related to tables (not shown) containing data and admin partitions P_Keys that have been defined in the network fabric. In an embodiment, the vGUID key is linked, through a relationship, to the admin partition registry (discussed above). In an embodiment, there is alternatively, or also, a relationship linking the admin partition registry to a vHCA Instance ID 1514. These relationships allow fabric components to identify which data and admin partitions the vHCA is a member of.

[0162] Figure 15 has been provided for exemplary purposes only, and one skilled in the art will appreciate that there are numerous ways to design and manage the storage of the data that make up a VM fabric profile. Further, the foregoing list of VM fabric profile contents is meant to be exemplary, not limiting, and other embodiments of virtual machine fabric profiles may include more, less, or other contents. In accordance with an embodiment, the database components depicted in Figure 15 can be a part of a much larger fabric manager database

that holds other relevant information about the fabric and can be interrelated with other tables to enhance the functionality of the fabric manager and other fabric components.

[0163] In one embodiment, a user interacts with virtual machine fabric profiles through an interface of the fabric manager. For instance, a user may create, edit, and delete VM fabric

5 profiles through the fabric manager. In accordance with an embodiment, some of the fabric profile information is supplied by a user creating or editing the VM fabric profile (e.g., the logical name of the VM fabric profile), while other of the information is generated or supplied by the fabric manager or the local SM of the subnet in which the fabric manager is being created (e.g., the VM-id, the vHCA instance ID, or the vGUID).

10 **[0164]** In accordance with an embodiment, the creation of the virtual machine fabric profile can take place in a management context that represents administrative privileges of the fabric resources the VM fabric profile is associated with. For example, a domain admin creates a VM fabric profile for use by nodes having HCAs that are a member of the same resource domain(s) (and the same admin partition(s)) as the vHCAs being configured for the

15 VM fabric profile. The created VM fabric profile is considered a (logical) resource, and a member of the resource domain in which it is created. Thus, by virtue of being a member of the same admin partition, the vHCA of the VM fabric profile has permission to be registered with any of the VF's of the HCAs that are also members of the resource domain, thereby easing administrative overhead.

20 **[0165]** In accordance with an embodiment, a fabric-level or domain-level administrator user can use a component of the fabric manager termed the "Virtual Machine Manager" (VMM) to set up and configure VM fabric profiles. VMM can use a fabric REST API in the creation of a VM fabric profile. A user can access VMM, for example, through a GUI of the fabric manager.

25 In accordance with an embodiment, the administrative user can supply certain parameters related to the VM fabric profile (such as logical name, profile type, and the number of vHCAs that will be associated with the profile), and other parameters can be automatically generated and assigned by the VMM (such as VM-id, and vGUIDs and vHCA instance IDs of each vHCA associated with the profile). Other CRUD actions, such as updating and deleting VM fabric profiles can also be available to the administrator user through the fabric manager, and

30 VMM, specifically.

[0166] In accordance with an embodiment, once all of the parameters necessary to build a VM fabric profile have been supplied via the VMM, the fabric manager can create an instance of the VM fabric profile object having the attributes specified by the administrator user and VMM, and persist the VM fabric profile object to the fabric level database.

35 **[0167]** In accordance with an embodiment, in an operational network fabric, a VM fabric profile can be held as one or multiple records in a cache of a SM. The VM fabric profile can either be retrieved (i.e., copied) from the fabric database by the SM (e.g., when the SM

boots), or be placed in the cache before it is persisted to the fabric database. The cache can be a volatile or non-volatile memory. The VM-id can be used as a key for querying the cache to retrieve attributes of a specific VM fabric profile.

[0168] By holding the VM fabric profiles in a high-speed cache on the SM rather than retrieving them from a persisted state (i.e., from the fabric database) every time a query is received, the lookup of VM fabric profile attributes can impose minimal overhead on the SM. This is especially important during VM and host boot-up, when fabric profile data will be needed to establish which HCAs VMs can and will be paired with.

[0169] **Figure 16** is a flow chart for making a VM fabric profile available to subnet resources.

[0170] At step 1610, a VM fabric profile including setup parameters and configuration of the VM is defined.

[0171] At step 1620, the VM fabric profile is stored in a fabric-level database.

[0172] At step 1630, the VM fabric profile is made available through a high-speed memory cache on the subnet manager.

[0173] At step 1640, VM fabric profile data is returned from the subnet manager based on VM-id lookup requests directed to the high-speed memory cache.

[0174] **Figure 17** is a flow chart for creating a virtual machine fabric profile for a virtual machine.

[0175] At step 1710, a virtual machine identifier (VM-id) is generated.

[0176] At step 1720, a virtual host channel adapter instance ID that identifies a virtual host channel adapter (vHCA) is generated.

[0177] At step 1730, a virtual globally unique identifier (vGUID) is assigned as a current fabric address of a virtual port of the vHCA, from a pool of globally unique identifiers.

[0178] At step 1740, a first relationship between a P_Key that defines an admin partition and the vGUID is created, where the relationship between the P_Key and the vGUID defines the virtual port for which the vGUID is assigned as the current fabric address as a member of the admin partition defined by the P_Key.

[0179] At step 1750, a second relationship between the VM-id and the vHCA instance ID is created, where the second relationship allows the vHCA instance ID to be retrieved through access to the VM-id.

[0180] At step 1760, a third relationship between the VM-id and the vGUID is created, where the third relationship allows the vGUID to be retrieved through access to the VM-id.

[0181] At step 1770, the VM-id, the virtual host channel adapter instance ID, the vGUID, the first relationship, the second relationship, and the third relationship is persisted in a format that preserves the VM-id, the virtual host channel adapter, the vGUID, the first relationship, the second relationship, and the third relationship.

[0182] Many features of the present invention can be performed in, using, or with the assistance of hardware, software, firmware, or combinations thereof. Consequently, features of the present invention may be implemented using a processing system (e.g., including one or more processors).

5 **[0183]** Features of the present invention can be implemented in, using, or with the assistance of a computer program product which is a storage medium (media) or computer readable medium (media) having instructions stored thereon/in which can be used to program a processing system to perform any of the features presented herein. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or any type of media or device suitable for storing instructions and/or data.

10 **[0184]** Stored on any one of the machine readable medium (media), features of the present invention can be incorporated in software and/or firmware for controlling the hardware of a processing system, and for enabling a processing system to interact with other mechanism utilizing the results of the present invention. Such software or firmware may include, but is not limited to, application code, device drivers, operating systems and execution environments/containers.

15 **[0185]** Features of the invention may also be implemented in hardware using, for example, hardware components such as application specific integrated circuits (ASICs). Implementation of the hardware state machine so as to perform the functions described herein will be apparent to persons skilled in the relevant art.

20 **[0186]** Additionally, the present invention may be conveniently implemented using one or more conventional general purpose or specialized digital computer, computing device, machine, or microprocessor, including one or more processors, memory and/or computer readable storage media programmed according to the teachings of the present disclosure. Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art.

25 **[0187]** While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example, and not limitation. It will be apparent to persons skilled in the relevant art that various changes in form and detail can be made therein without departing from the spirit and scope of the invention.

30 **[0188]** The present invention has been described above with the aid of functional building blocks illustrating the performance of specified functions and relationships thereof. The boundaries of these functional building blocks have often been arbitrarily defined herein for the convenience of the description. Alternate boundaries can be defined so long as the

specified functions and relationships thereof are appropriately performed. Any such alternate boundaries are thus within the scope and spirit of the invention.

[0189] The foregoing description of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the

5 invention to the precise forms disclosed. The breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments. Many modifications and variations will be apparent to the practitioner skilled in the art. The

modifications and variations include any relevant combination of the disclosed features. The embodiments were chosen and described in order to best explain the principles of the

10 invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalents.

Claims

What is claimed is:

1. A method of creating a virtual machine (VM) fabric profile for a virtual machine comprising the steps of:
 - 5 generating a virtual machine identifier (VM-id);
 - generating a virtual host channel adapter instance ID that identifies a virtual host channel adapter (vHCA);
 - 10 assigning, from a pool of globally unique identifiers, a virtual globally unique identifier (vGUID) as a current fabric address of a virtual port of the vHCA;
 - creating a first relationship between a P_Key that defines an admin partition and the vGUID, wherein the relationship between the P_Key and the vGUID defines the virtual port for which the vGUID is assigned as the current fabric address as a member of the admin partition defined by the P_Key;
 - 15 creating a second relationship between the VM-id and the vHCA instance ID, wherein the second relationship allows the vHCA instance ID to be retrieved through access to the VM-id;
 - creating a third relationship between the VM-id and the vGUID, wherein the third relationship allows the vGUID to be retrieved through access to the VM-id; and
 - 20 persisting the VM-id, the virtual host channel adapter instance ID, the vGUID, the first relationship, the second relationship, and the third relationship in a format that preserves the VM-id, the virtual host channel adapter, the vGUID, the first relationship, the second relationship, and the third relationship.
- 25 2. The method of Claim 1, further comprising the steps of:
 - placing the VM-id, the vHCA instance ID, the vGUID, the first relationship, the second relationship, the third relationship, and the P_Key in a cache of a subnet manager of a subnet, wherein the admin partition defined by the P_Key is defined within the subnet.
- 30 3. The method of Claim 2, wherein the VM-id, the virtual host channel adapter instance ID, the vGUID, the first relationship, the second relationship, the third relationship, and the P_Key placed in the cache are copies.
4. The method of Claim 2 or 3 wherein the cache is a volatile memory.

5. The method of any preceding Claim, further comprising the steps of:
using the VM-id as a lookup key to access the vHCA instance ID, the vGUID, and the
P_Key.

5 6. The method of Claim 5, wherein the using the VM-id as a lookup key is initiated by a
physical host channel adapter attempting to register the vHCA instance ID with a virtual
function of the physical host channel adapter.

10 7. The method of any of Claims 2 to 6, further comprising the steps of:
retrieving from the cache, by the subnet manager, the vHCA instance ID and the
vGUID; and
sending the vHCA instance ID and the vGUID to the physical host channel adapter.

15 8. A system for creating and accessing a virtual machine (VM) fabric profile for a virtual
machine comprising:

a node, including a processor;
a memory accessible to the node; and
wherein the node operates to:

generate a virtual machine identifier (VM-id);

20 generate a virtual host channel adapter instance ID that identifies a virtual
host channel adapter (vHCA);

assign, from a pool of globally unique identifiers, a virtual globally unique
identifier (vGUID) as a current fabric address of a virtual port of the vHCA;

25 create a first relationship between a P_Key that defines an admin partition and
the vGUID, wherein the relationship between the P_Key and the vGUID defines the
virtual port for which the vGUID is assigned as the current fabric address as a
member of the admin partition defined by the P_Key;

30 create a second relationship between the VM-id and the vHCA instance ID,
wherein the second relationship allows the vHCA instance ID to be retrieved through
access to the VM-id;

create a third relationship between the VM-id and the vGUID, wherein the third
relationship allows the vGUID to be retrieved through access to the VM-id; and

35 persist the VM-id, the virtual host channel adapter instance ID, the vGUID, the
first relationship, the second relationship, and the third relationship to the memory in a
format that preserves the VM-id, the virtual host channel adapter, the vGUID, the first
relationship, the second relationship, and the third relationship.

9. The system of Claim 8, further comprising a subnet manager of the subnet, wherein the VM-id, the vHCA instance ID, the vGUID, the first relationship, the second relationship, the third relationship, and the P_Key are placed in a cache of the subnet manager, and wherein the admin partition defined by the P_Key is defined within the subnet.

5

10. The system of Claim 9, wherein the VM-id, the virtual host channel adapter instance ID, the vGUID, the first relationship, the second relationship, the third relationship, and the P_Key placed in the cache are copies.

10 11. The system of Claim 9 or 10 wherein the cache is a volatile memory.

12. The system of any of Claims 8 to 11, wherein the VM-id is used as a lookup key to access the vHCA instance ID, the vGUID, and the P_Key.

15 13. The system of Claim 12, wherein a physical host channel adapter initiates the lookup in an attempt to register the vHCA instance ID with a virtual function of the physical host channel adapter.

20 14. The system of any of Claims 9 to 13, wherein the subnet manager retrieves the vHCA instance ID and the vGUID from the cache; and

wherein the subnet manager sends the vHCA instance ID and the vGUID to the physical host channel adapter.

25 15. A non-transitory computer readable storage medium, including instructions stored thereon for creating a virtual machine (VM) fabric profile for a virtual machine, which when read and executed by one or more computers cause the one or more computers to perform steps comprising:

generating a virtual machine identifier (VM-id);

generating a virtual host channel adapter instance ID that identifies a virtual host

30 channel adapter (vHCA);

assigning, from a pool of globally unique identifiers, a virtual globally unique identifier (vGUID) as a current fabric address of a virtual port of the vHCA;

35 creating a first relationship between a P_Key that defines an admin partition and the vGUID, wherein the relationship between the P_Key and the vGUID defines the virtual port for which the vGUID is assigned as the current fabric address as a member of the admin partition defined by the P_Key;

creating a second relationship between the VM-id and the vHCA instance ID, wherein

the second relationship allows the vHCA instance ID to be retrieved through access to the VM-id;

creating a third relationship between the VM-id and the vGUID, wherein the third relationship allows the vGUID to be retrieved through access to the VM-id; and

5 persisting the VM-id, the virtual host channel adapter instance ID, the vGUID, the first relationship, the second relationship, and the third relationship in a format that preserves the VM-id, the virtual host channel adapter, the vGUID, the first relationship, the second relationship, and the third relationship.

10 16. The non-transitory computer readable storage medium of Claim 15, the steps further comprising:

placing the VM-id, the vHCA instance ID, the vGUID, the first relationship, the second relationship, the third relationship, and the P_Key in a cache of a subnet manager of a subnet, wherein the admin partition defined by the P_Key is defined within the subnet.

15

17. The non-transitory computer readable storage medium of Claim 16, wherein the VM-id, the virtual host channel adapter instance ID, the vGUID, the first relationship, the second relationship, the third relationship, and the P_Key placed in the cache are copies.

20 18. The non-transitory computer readable storage medium of Claim 16 or 17, wherein the cache is a volatile memory.

19. The non-transitory computer readable storage medium of any of Claims 15 to 18, the steps further comprising:

25 using the VM-id as a lookup key to access the vHCA instance ID, the vGUID, and the P_Key, wherein the using the VM-id as a lookup key is initiated by a physical host channel adapter attempting to register the vHCA instance ID with a virtual function of the physical host channel adapter.

30 20. The non-transitory computer readable storage medium of any of Claims 16 to 19, the steps further comprising:

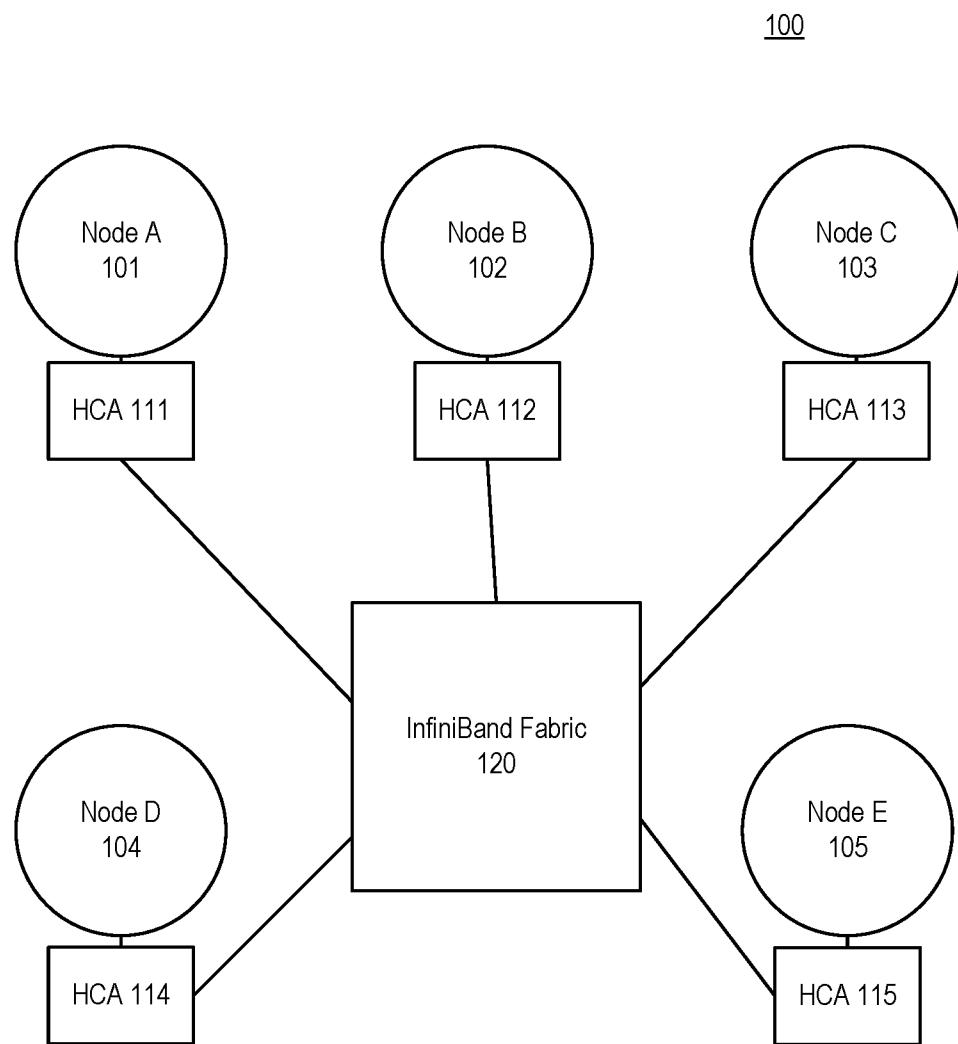
retrieving from the cache, by the subnet manager, the vHCA instance ID and the vGUID; and

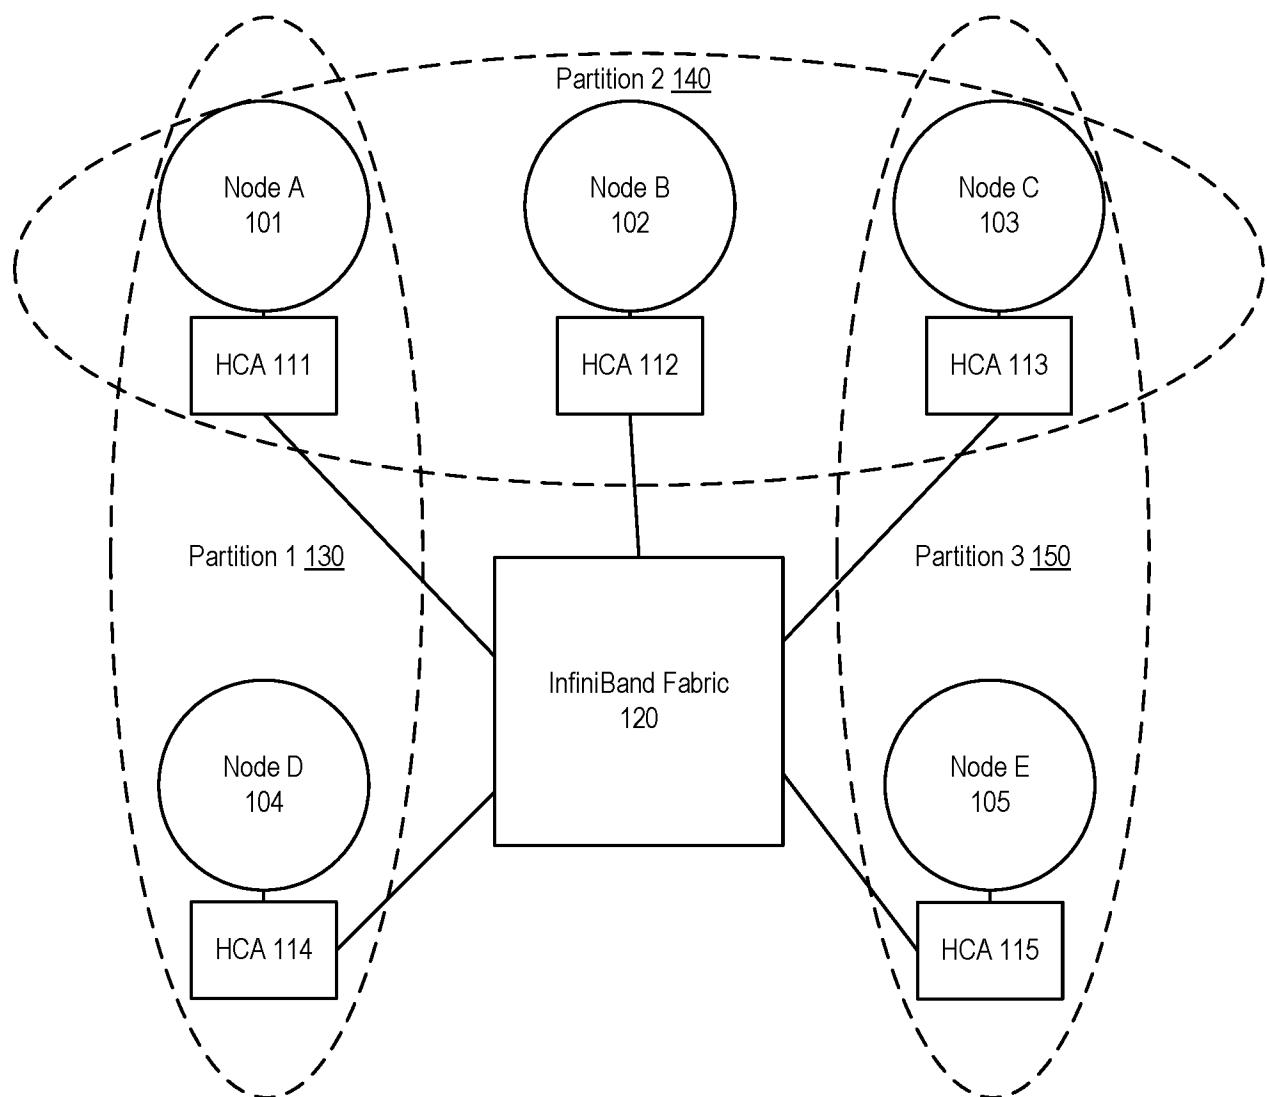
sending the vHCA instance ID and the vGUID to the physical host channel adapter.

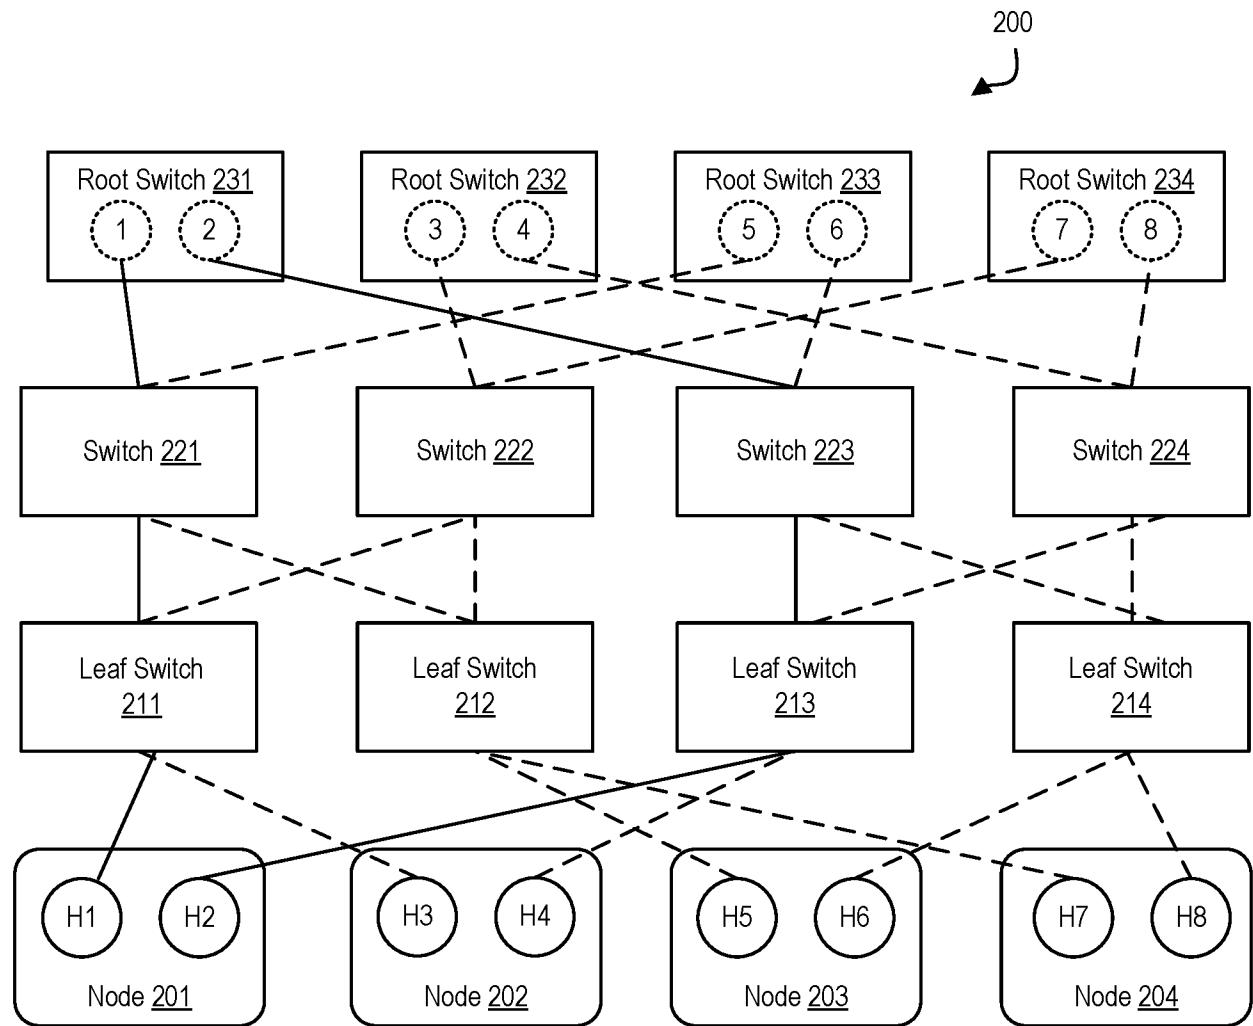
35

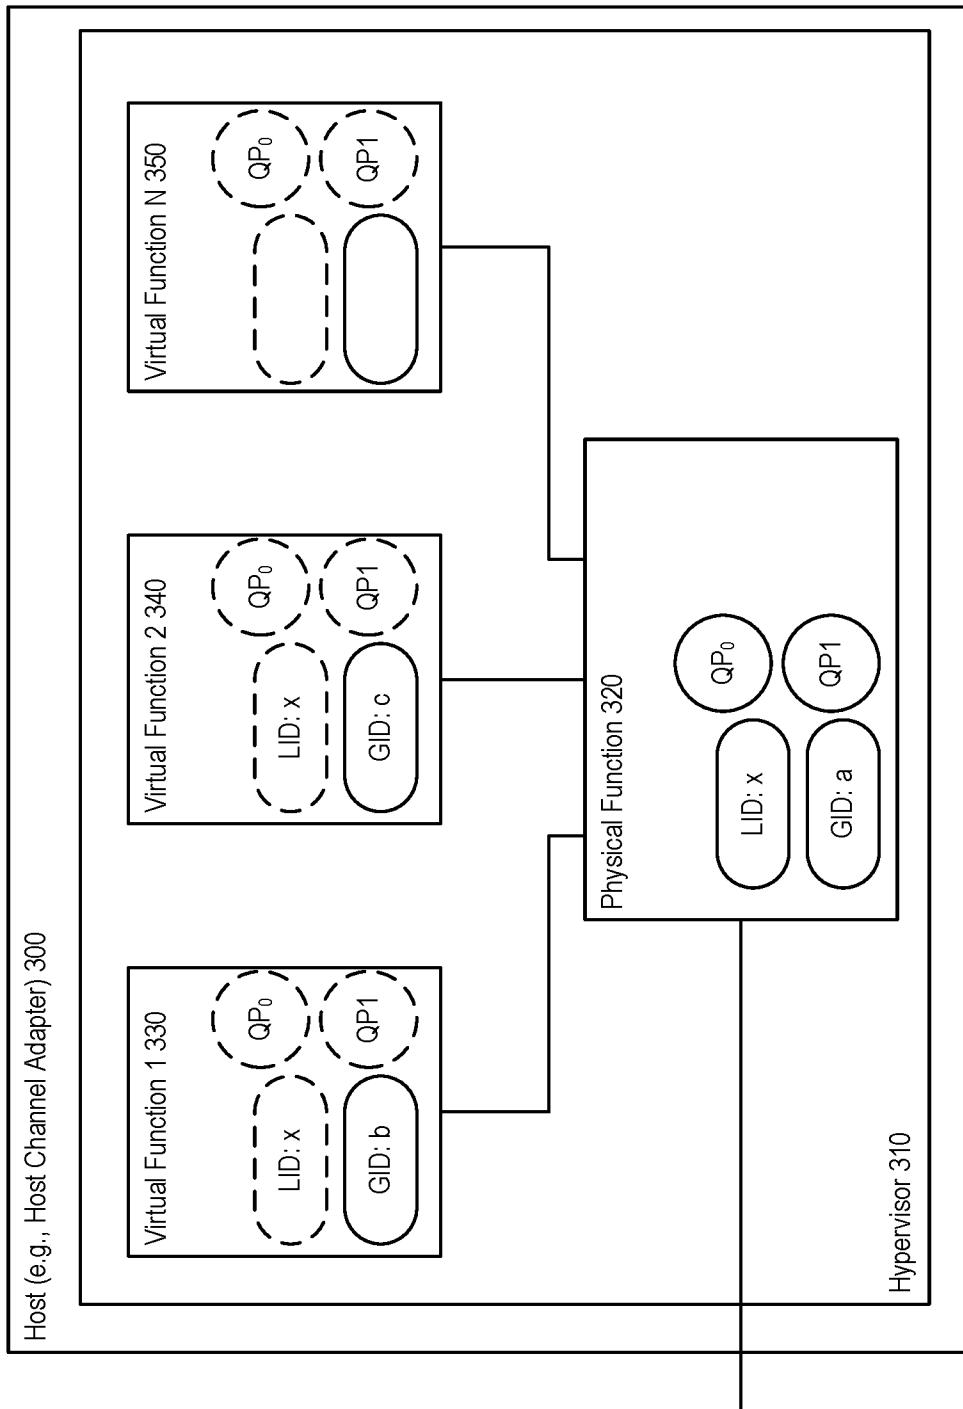
21. A computer program comprising program instructions in machine-readable format that when executed by a computer system cause the computer system to perform the method of any of Claims 1 to 7.

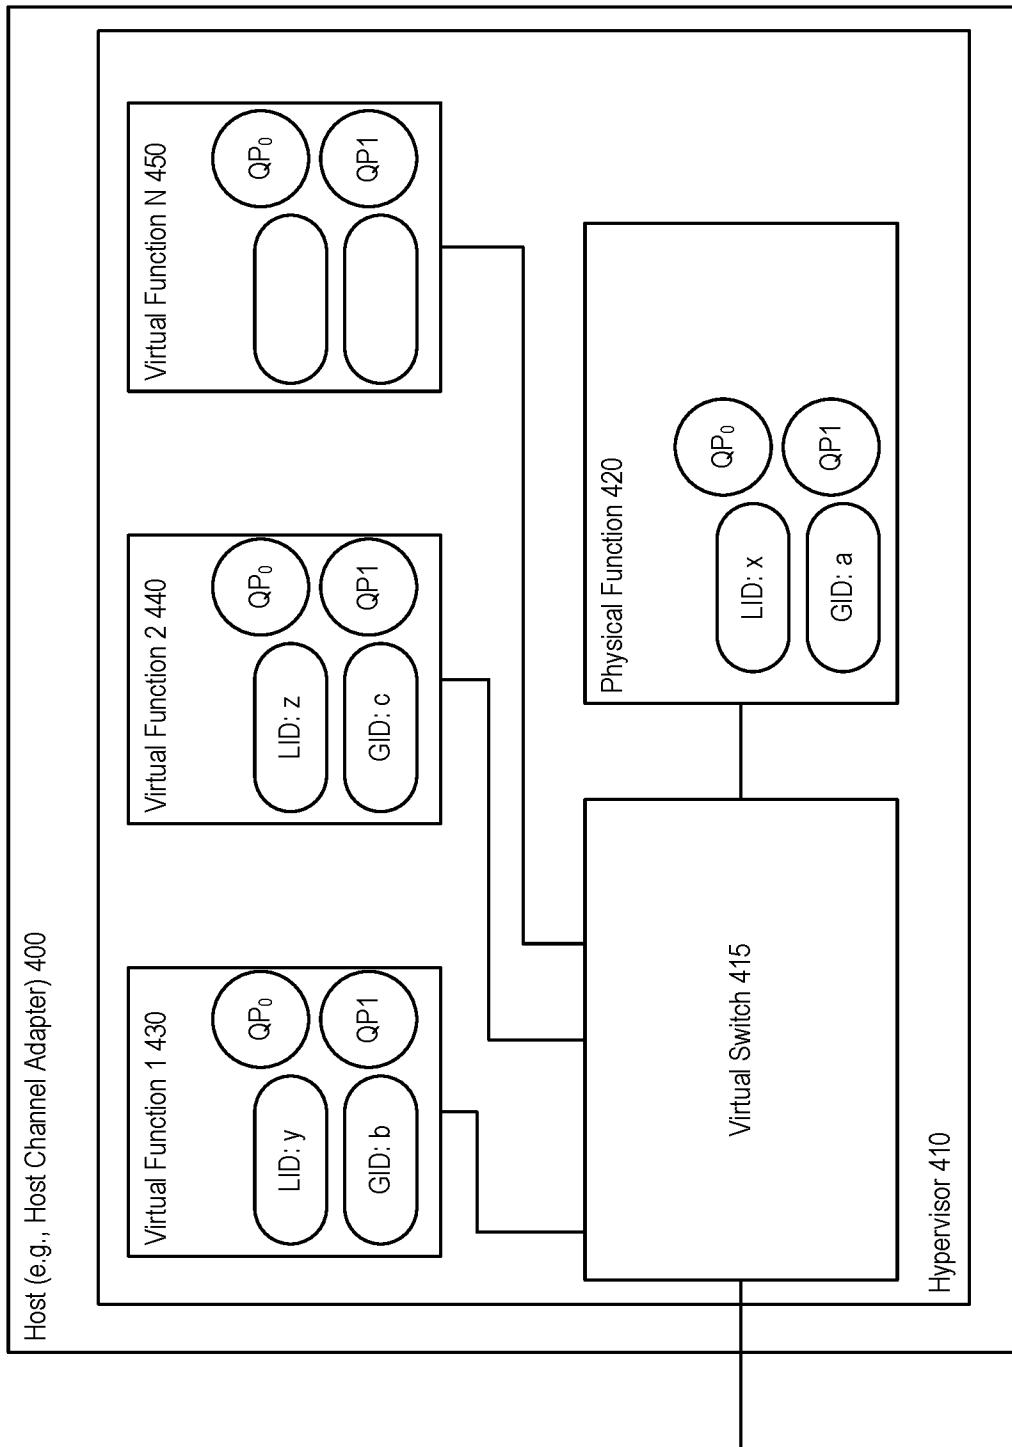
5 22. A computer program product comprising the computer program of Claim 21 stored in a non-transitory machine readable data storage medium.

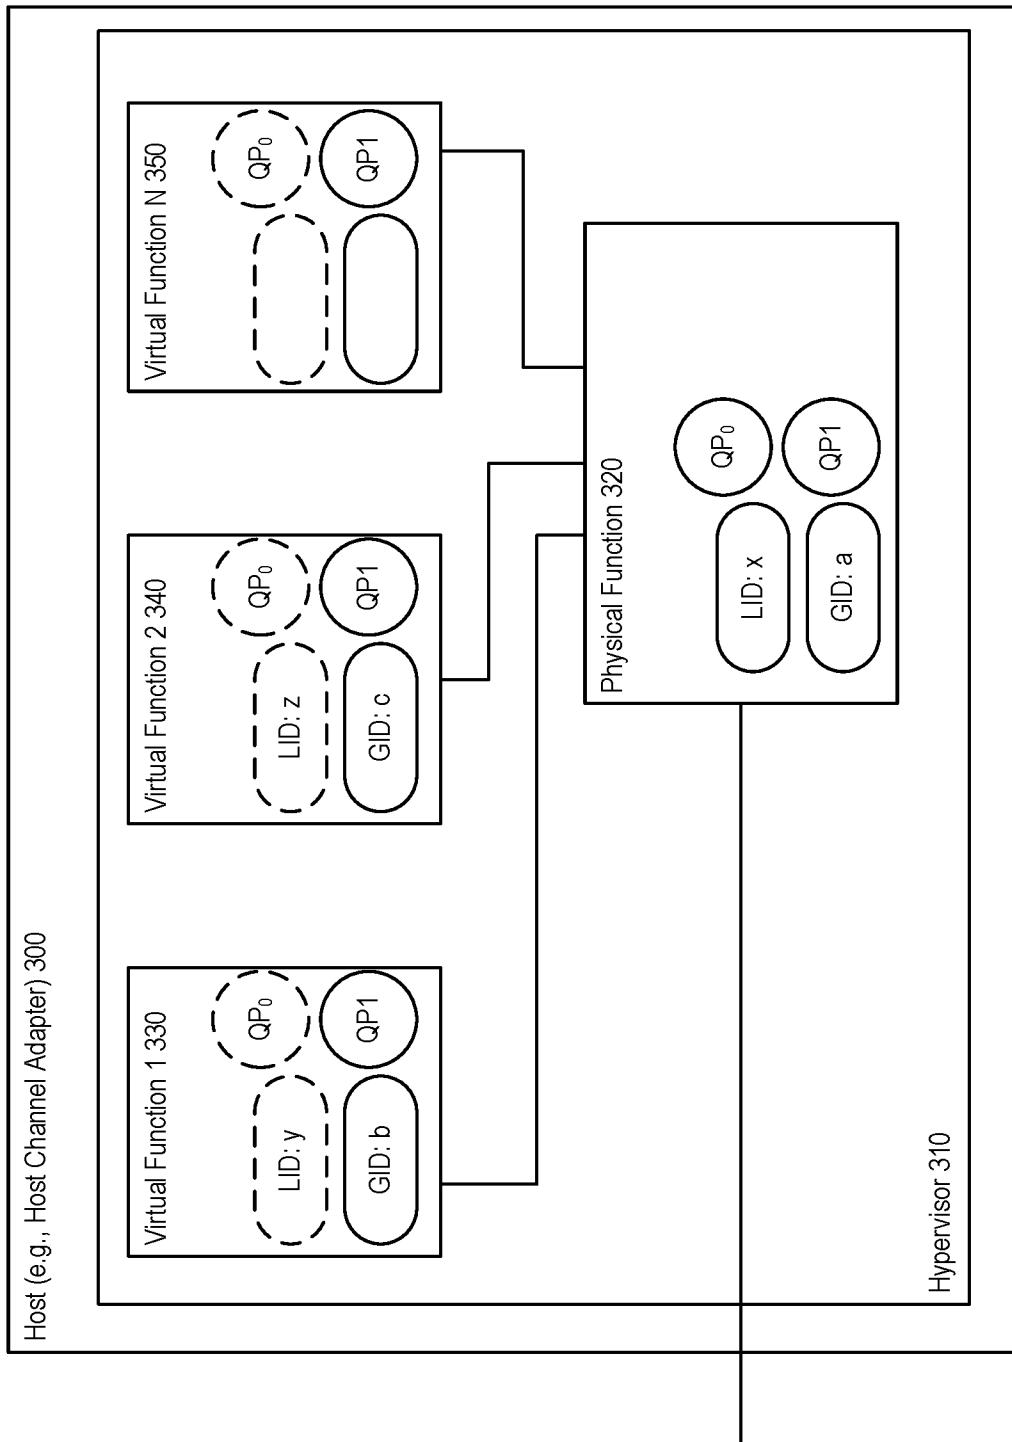
**FIGURE 1**

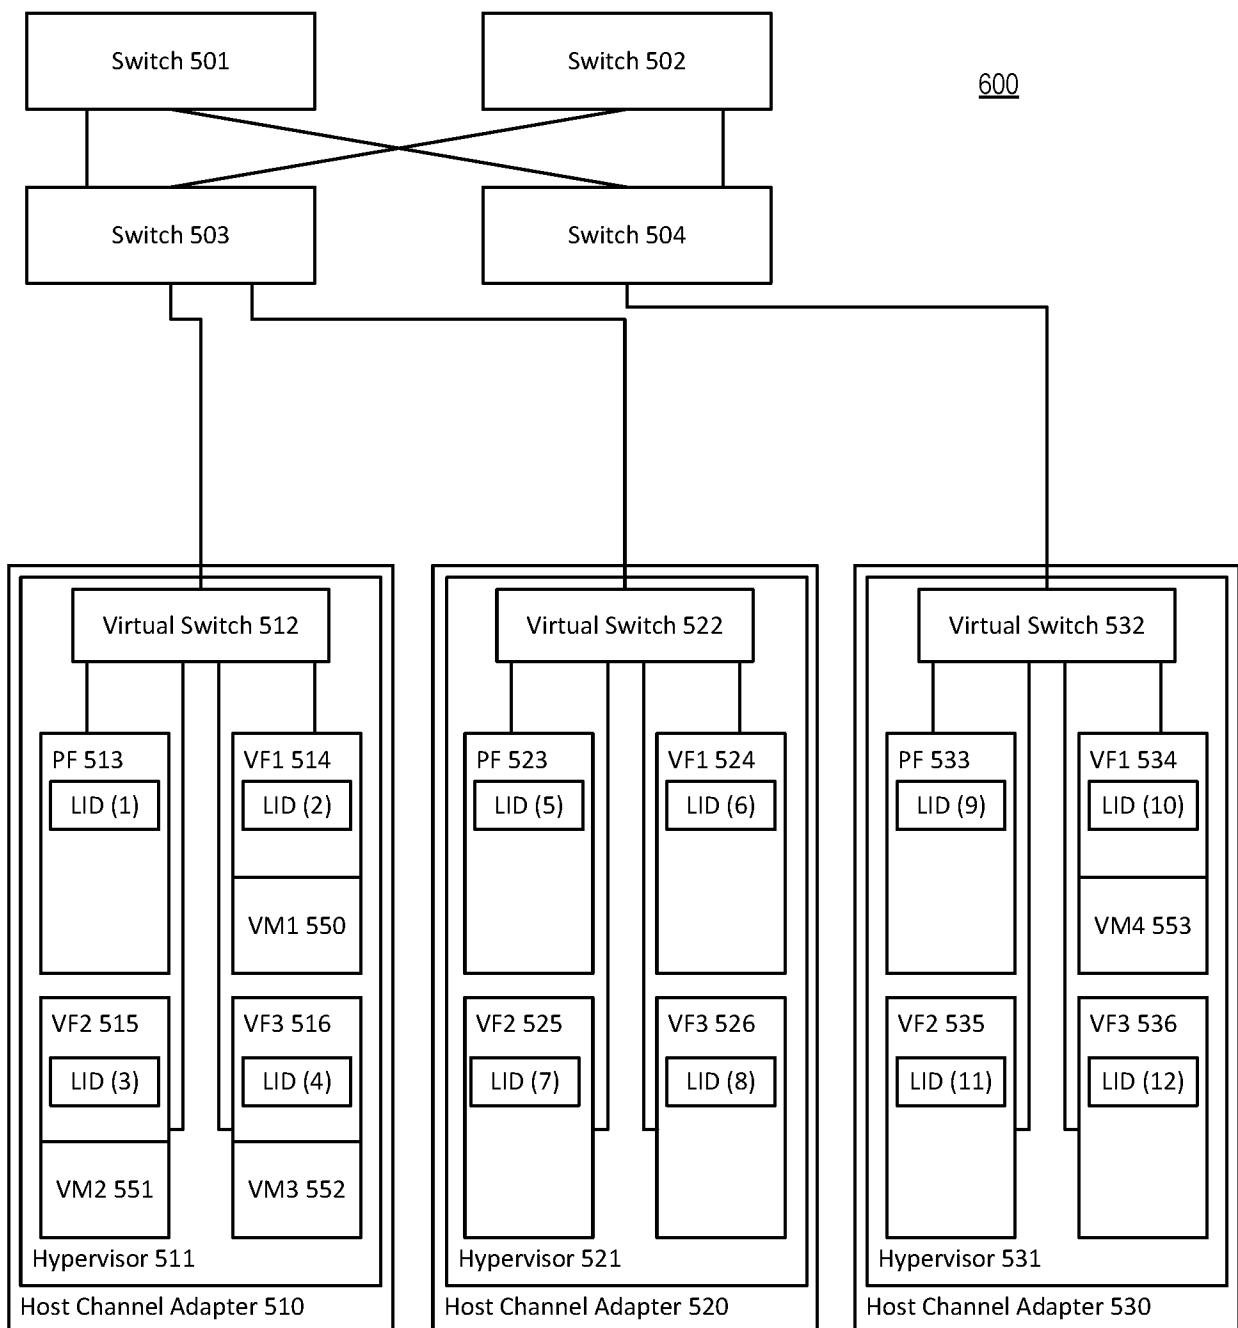
**FIGURE 2**

**FIGURE 3**

**FIGURE 4**

**FIGURE 5**

**FIGURE 6**

**FIGURE 7**

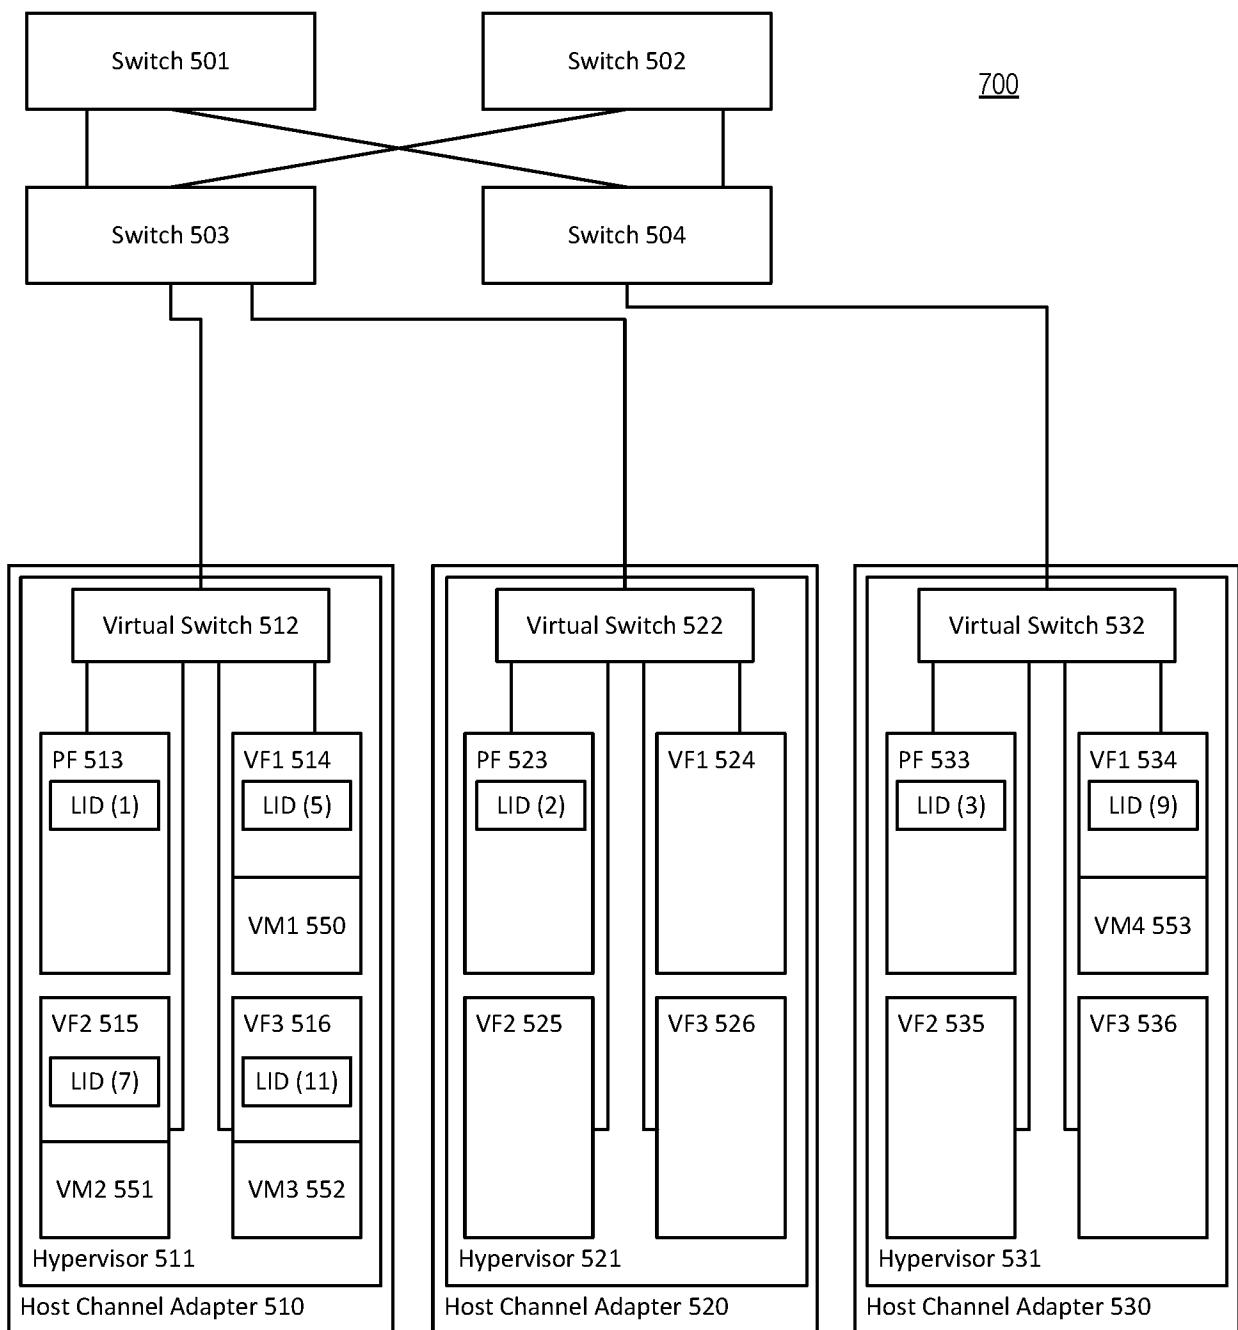
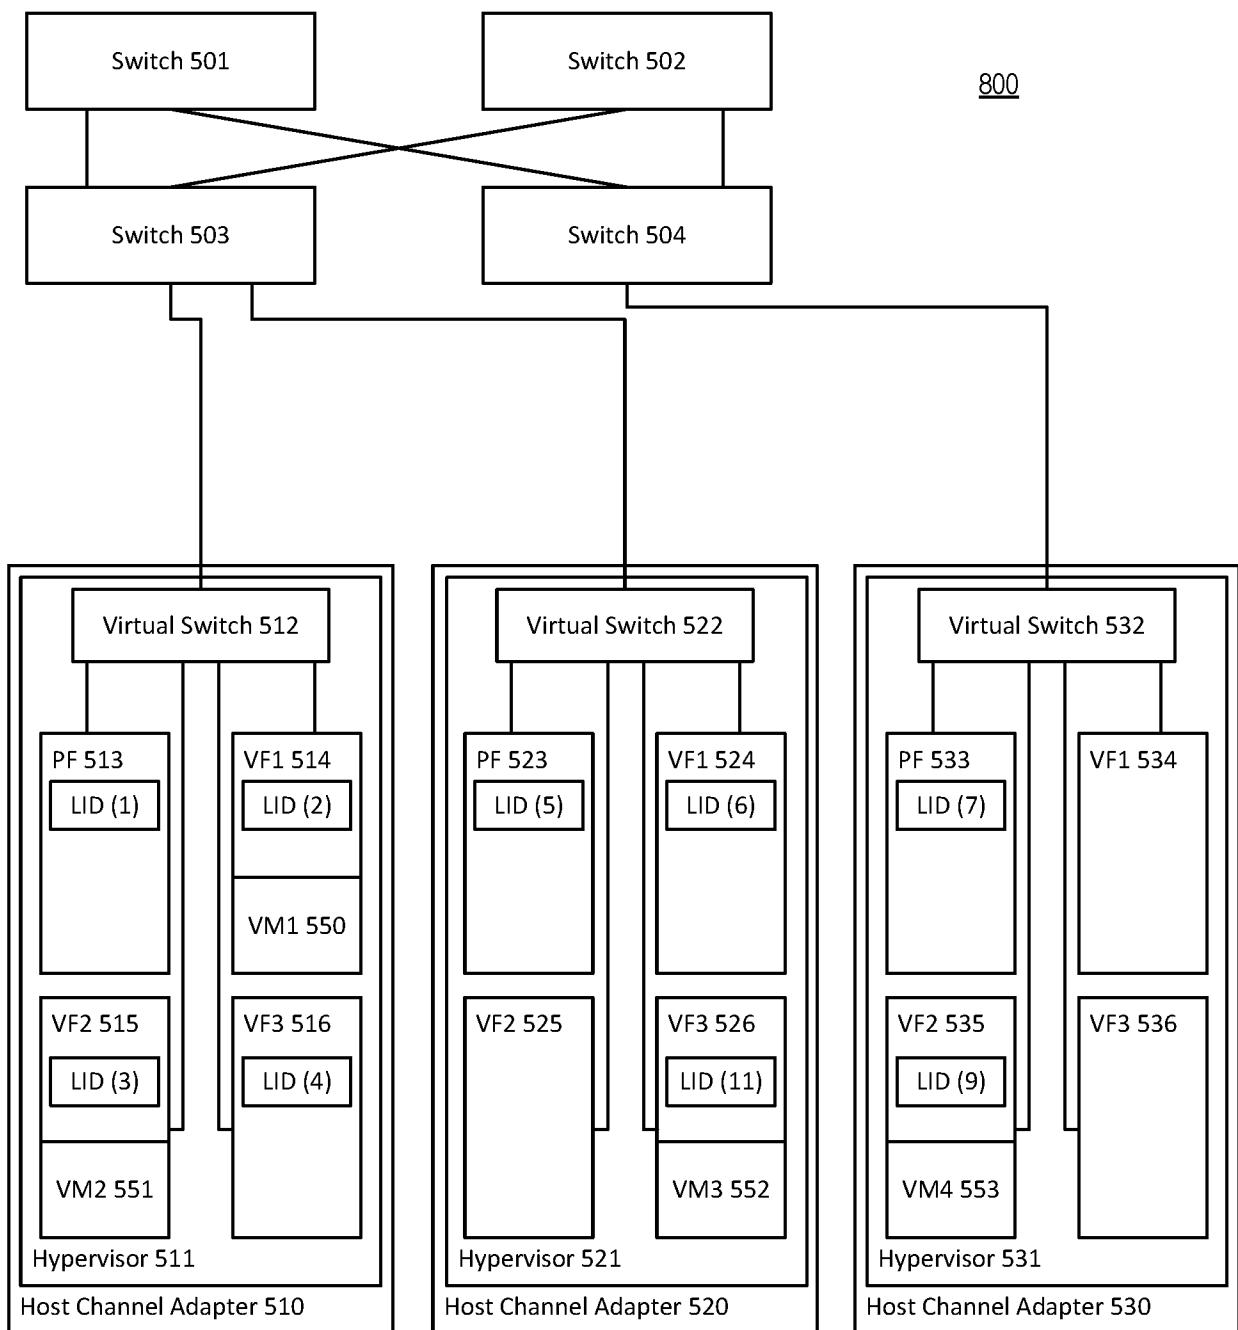


FIGURE 8

**FIGURE 9**

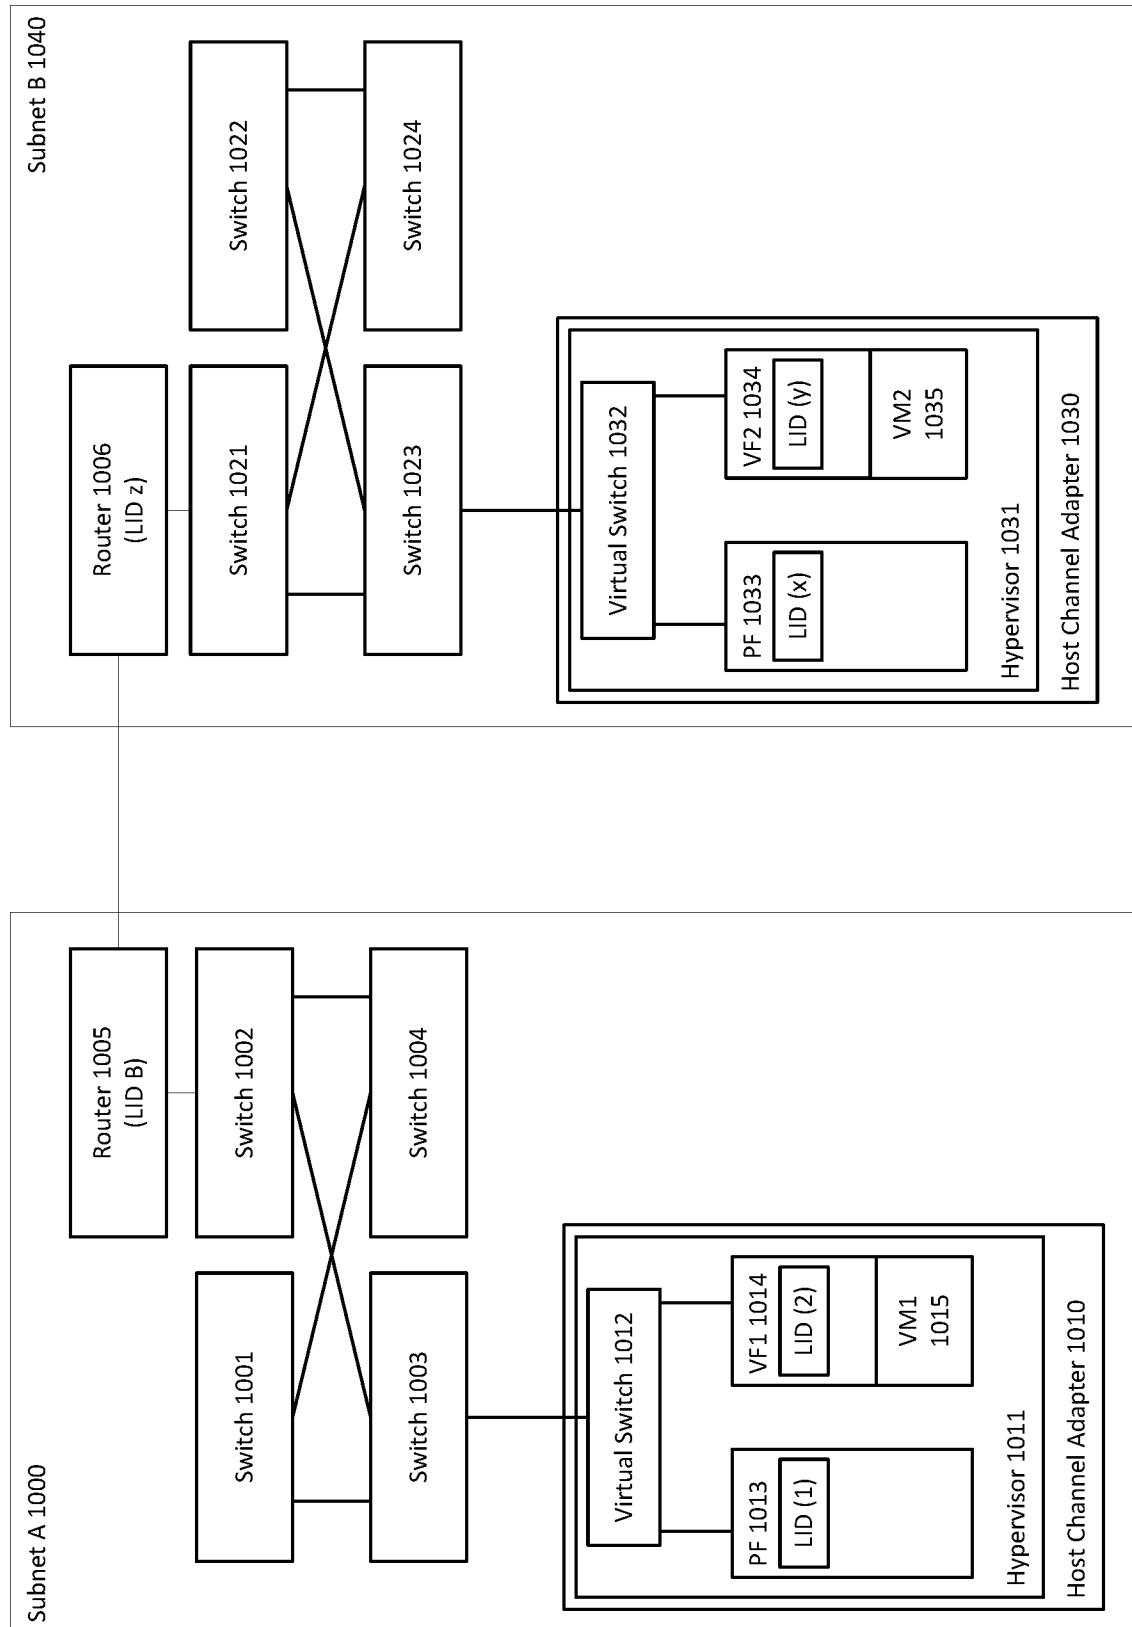


FIGURE 10

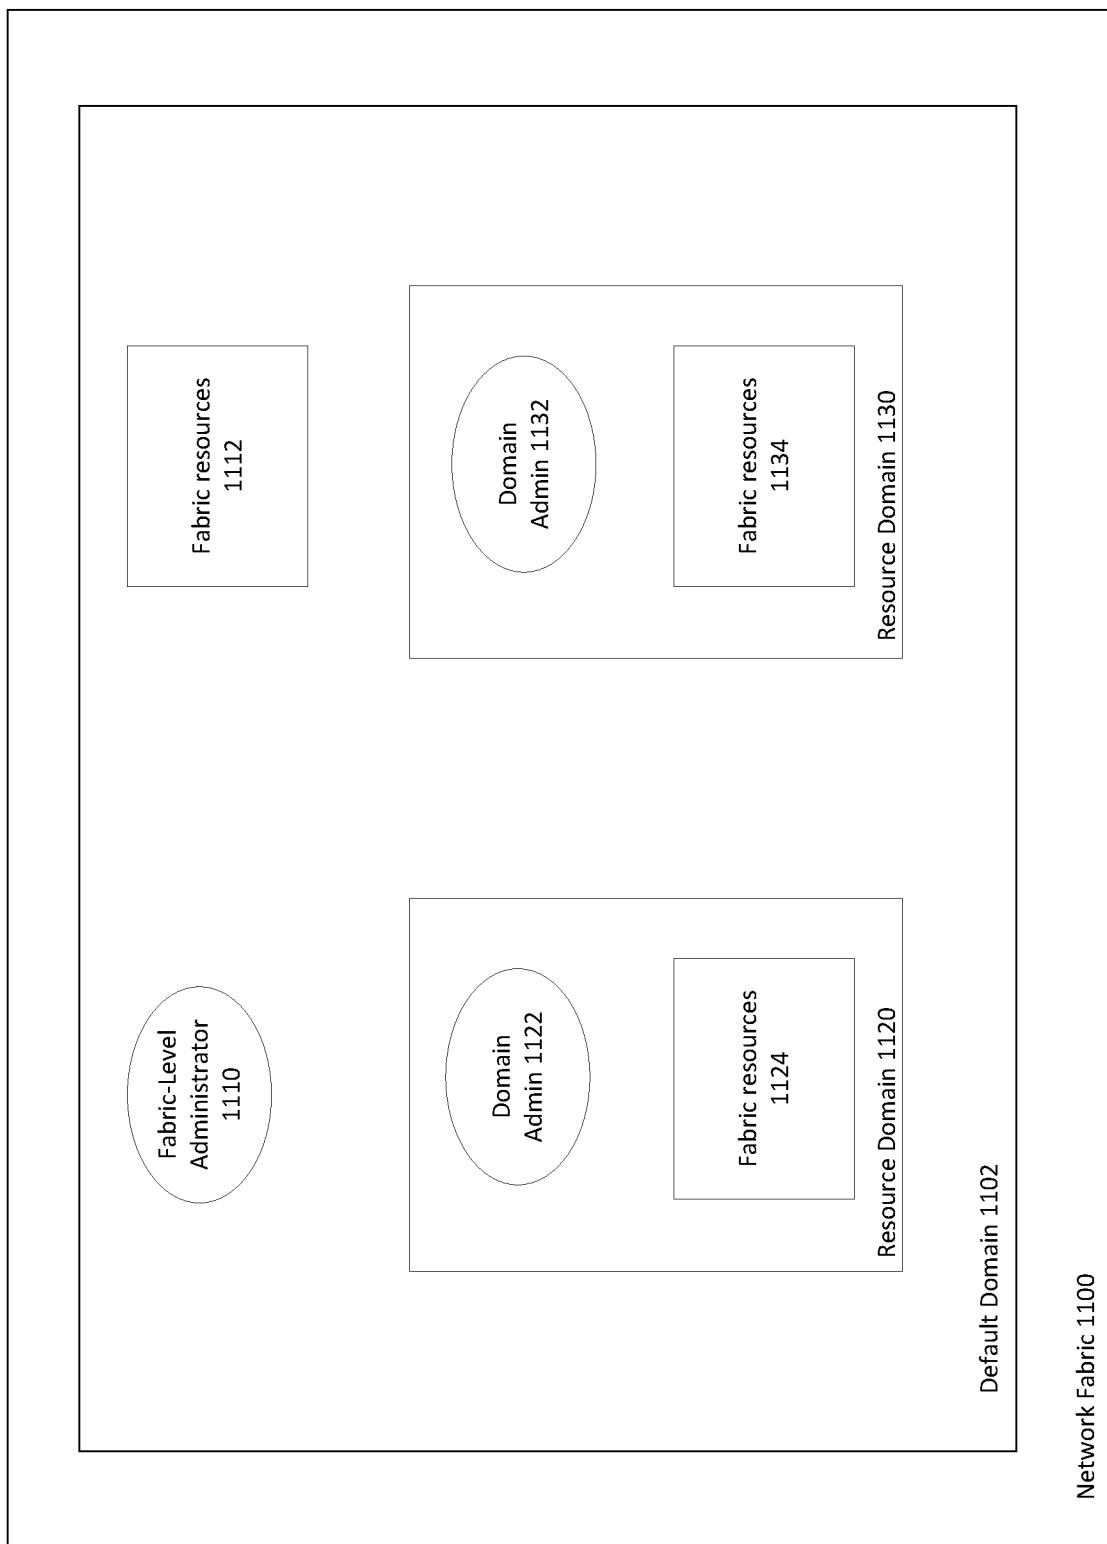
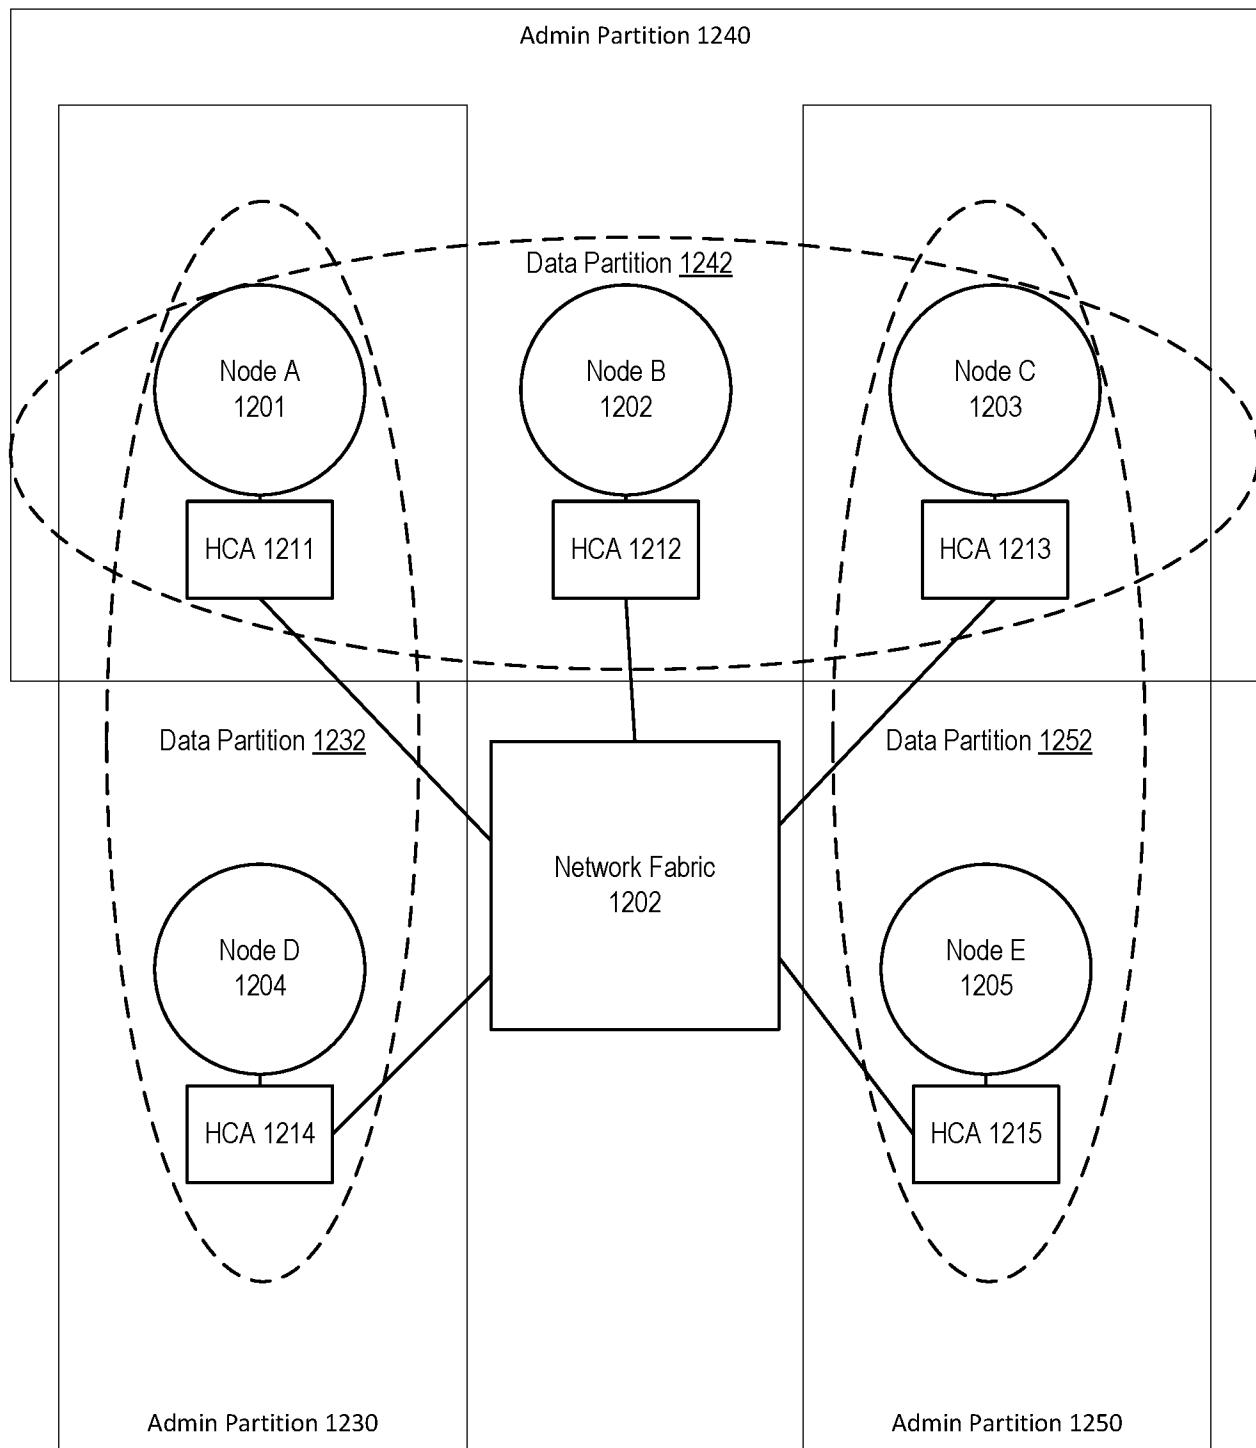


FIGURE 11

**FIGURE 12**

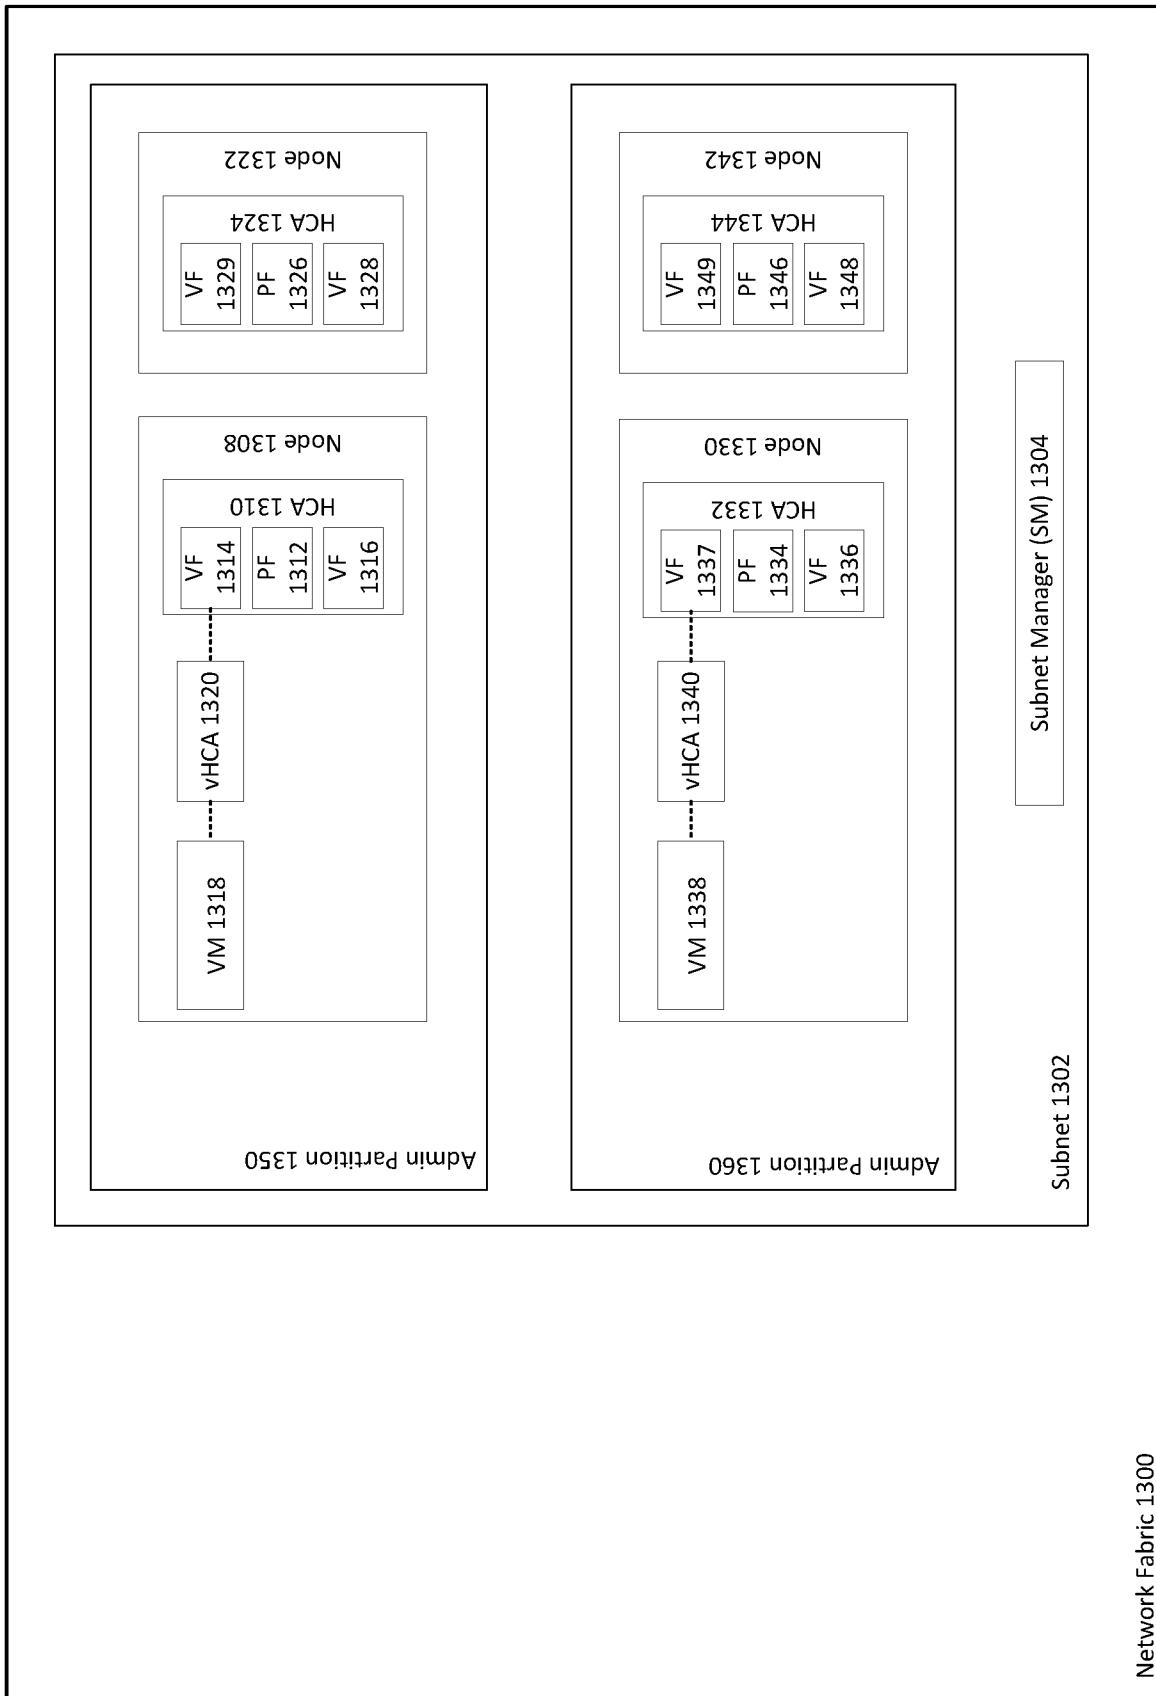


FIGURE 13

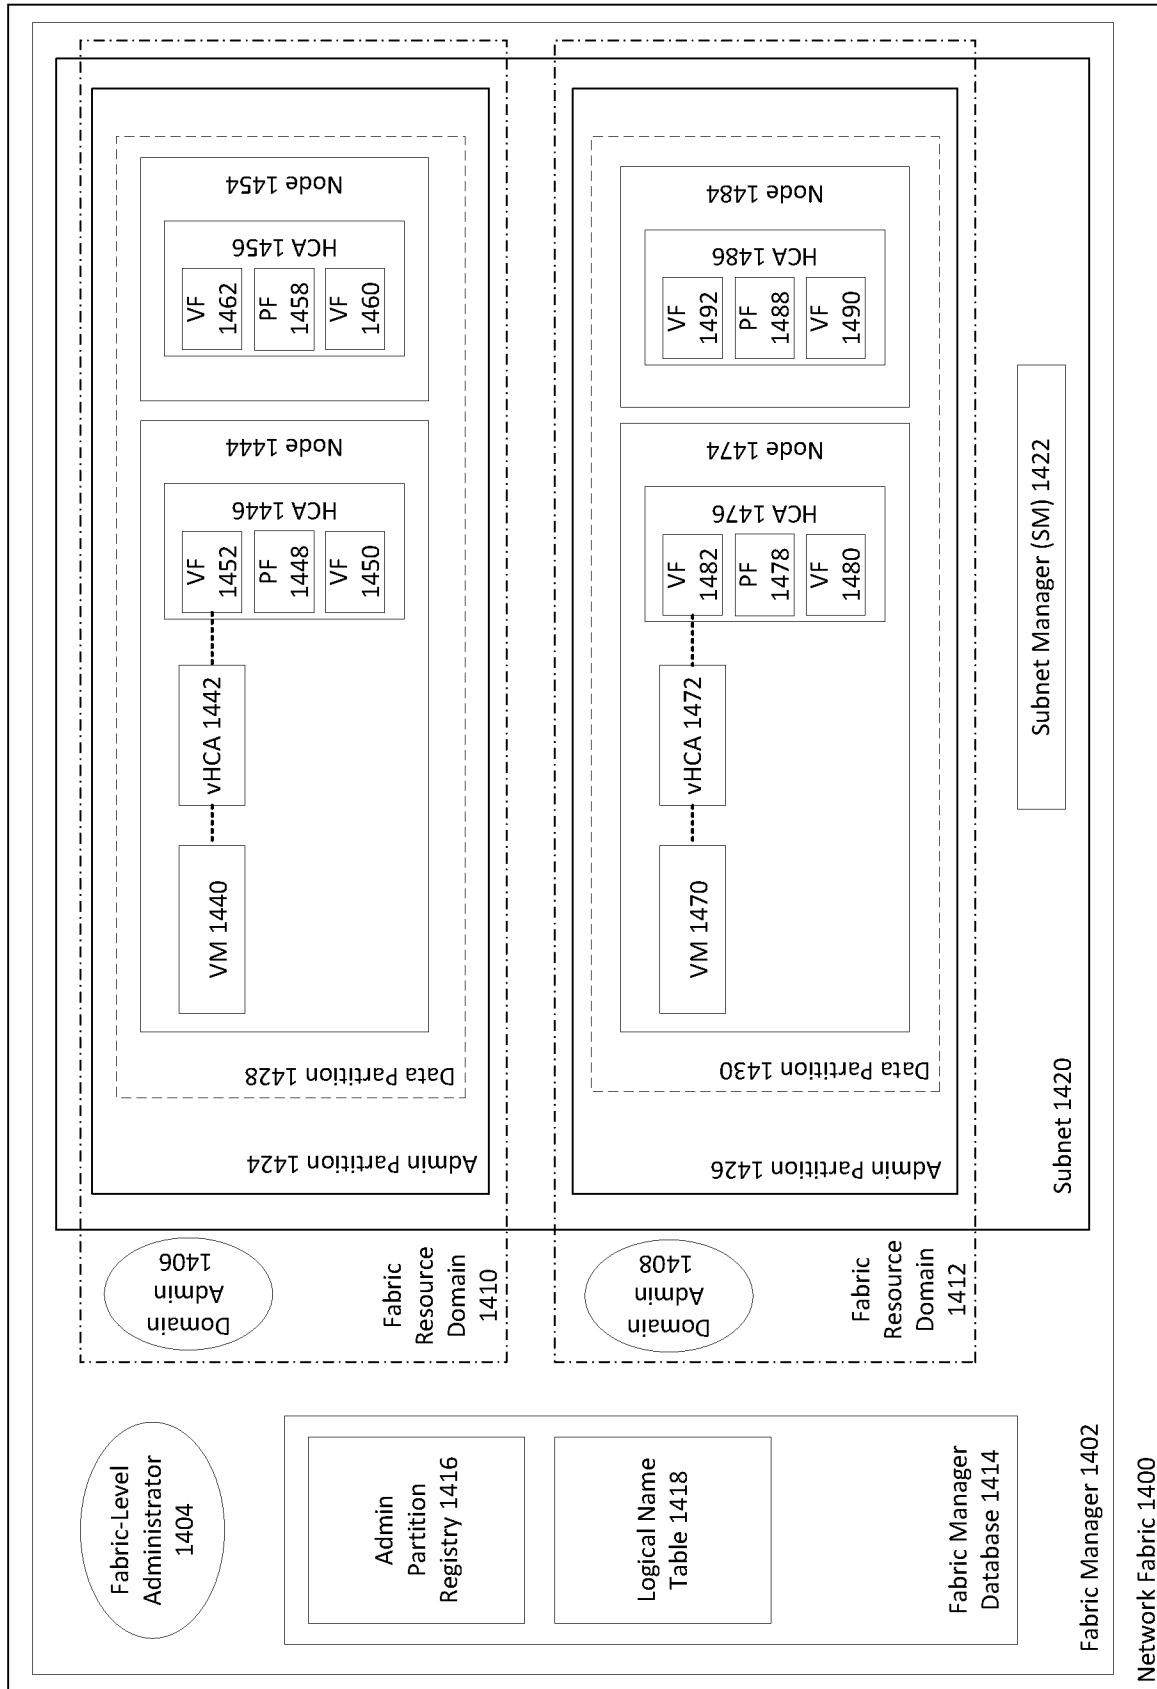
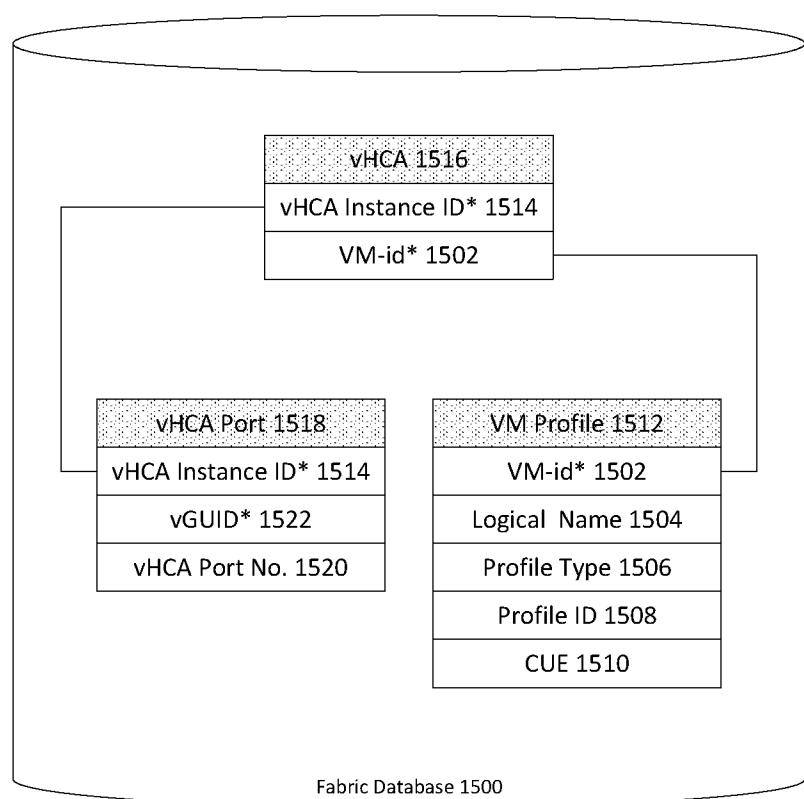
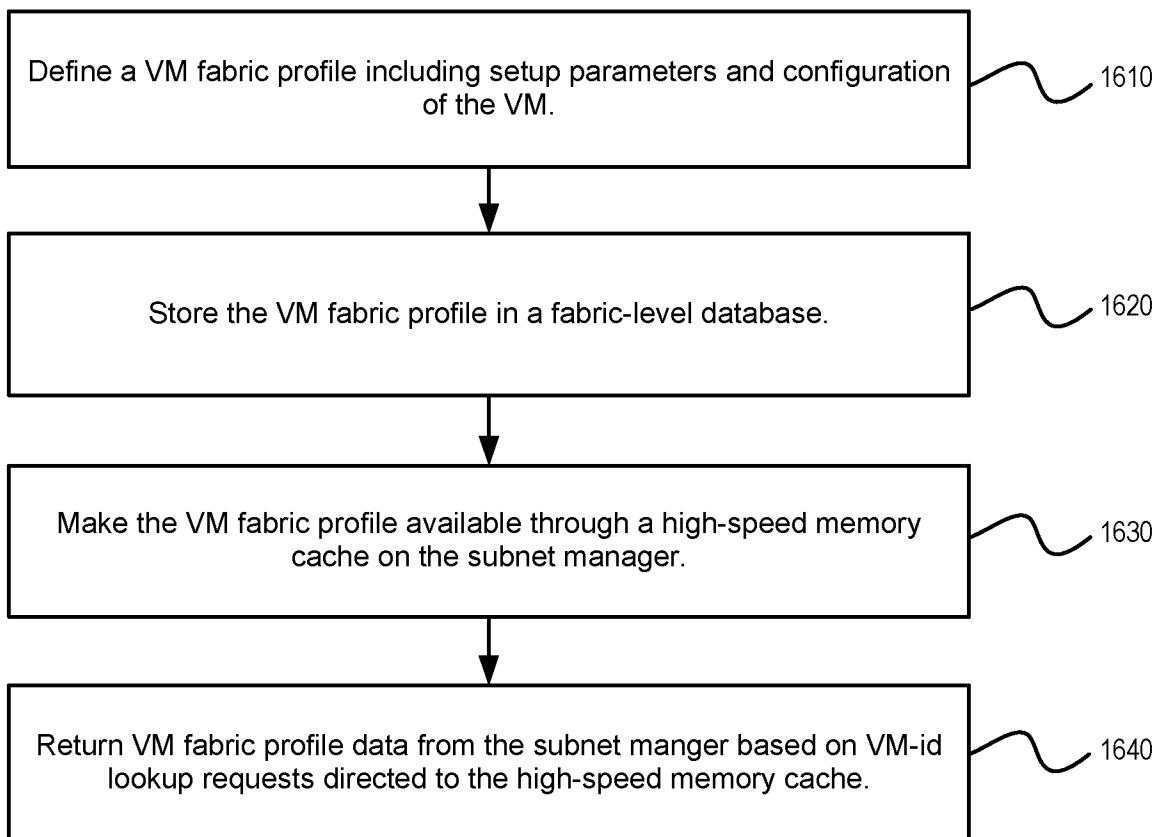


FIGURE 14

**FIGURE 15**

**FIGURE 16**

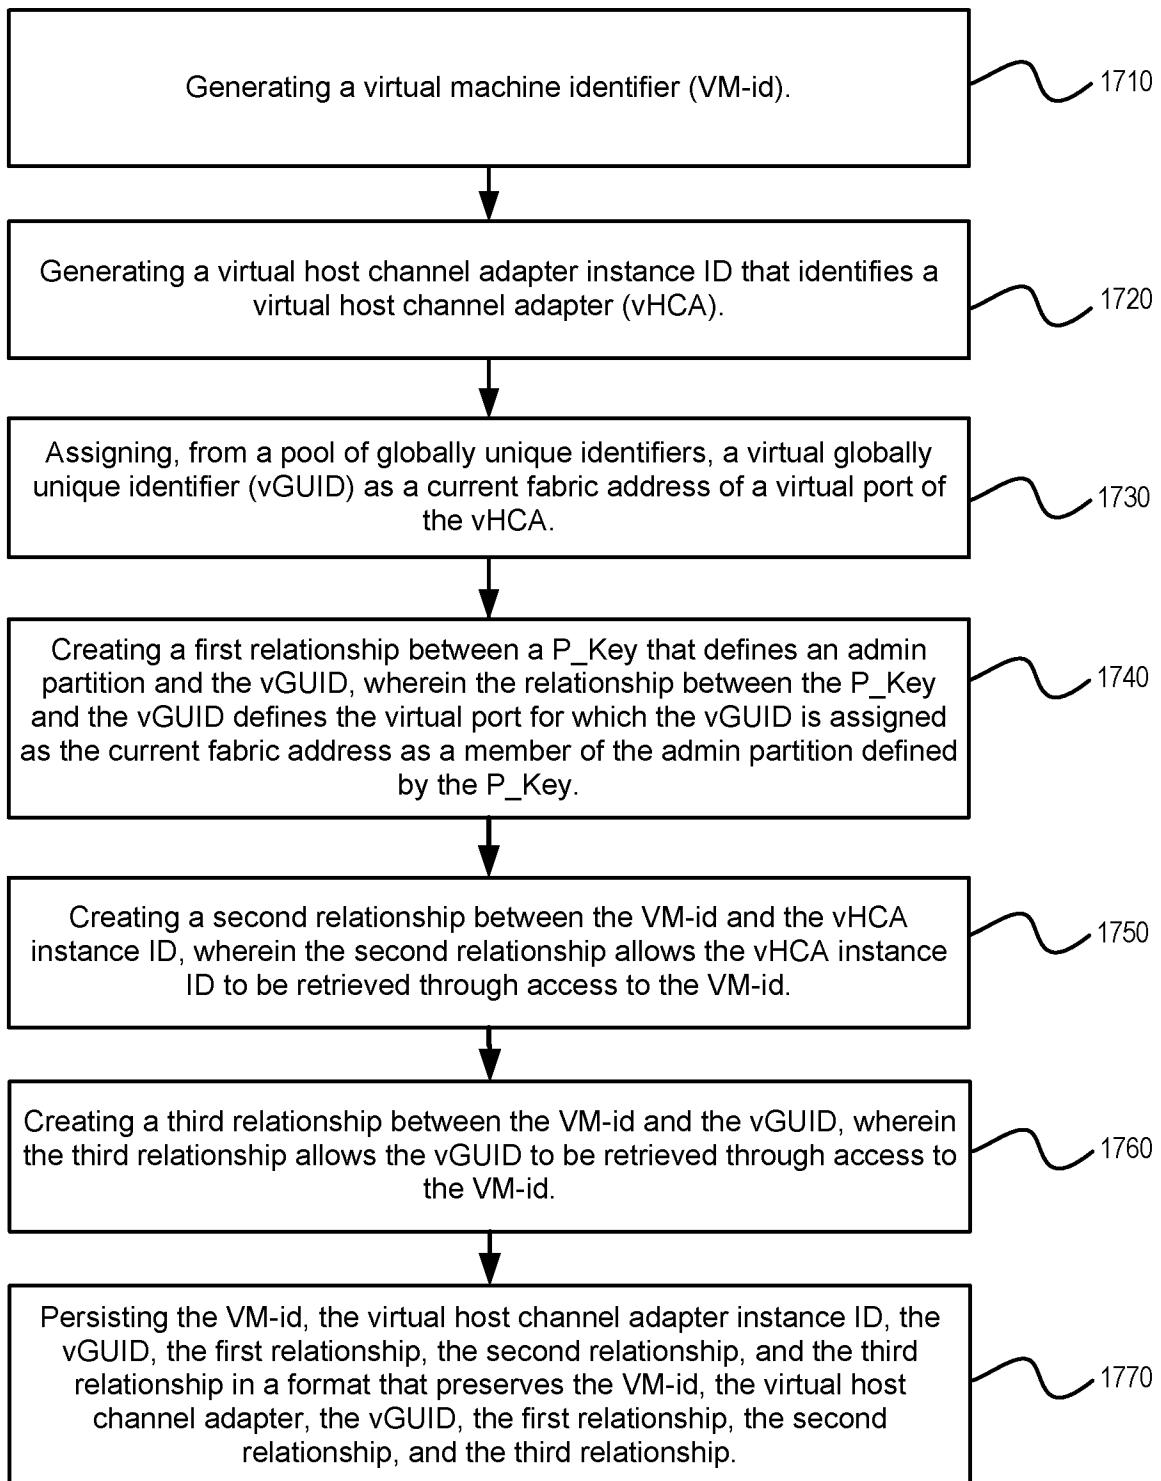


FIGURE 17

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/015169

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/455 H04L12/931 H04L29/12
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)
G06F H04L

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

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

EPO-Internal, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2013/051232 A1 (GAVRILOV CONSTANTINE [IL] ET AL) 28 February 2013 (2013-02-28) paragraphs [0001] - [0011] paragraphs [0026] - [0079] figures 1-11 ----- US 2015/063356 A1 (MAKHHERVAKS VADIM [US] ET AL) 5 March 2015 (2015-03-05) paragraphs [0007] - [0010] paragraphs [0029] - [0155] figures 1-17 ----- US 6 438 130 B1 (KAGAN MICHAEL [IL] ET AL) 20 August 2002 (2002-08-20) column 1, line 10 - column 3, line 47 column 4, line 3 - column 6, line 12 figures 1-4B -----	1-22 1-22 1-22



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	Date of mailing of the international search report
4 April 2017	11/04/2017
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 Noll, Joachim

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2017/015169

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013051232	A1 28-02-2013	US 2013051232 A1 US 2013229919 A1	28-02-2013 05-09-2013
US 2015063356	A1 05-03-2015	NONE	
US 6438130	B1 20-08-2002	NONE	