

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

AMAZON TECHNOLOGIES, INC.,)
)
Plaintiff,)
)
v.) C.A. No. _____
)
NOKIA CORPORATION, NOKIA) **JURY TRIAL DEMANDED**
SOLUTIONS AND NETWORKS OY, and)
NOKIA OF AMERICA CORPORATION,)
)
Defendants.)

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Amazon Technologies, Inc. (“ATI,” “Amazon,” or “Plaintiff”) files this Complaint against Nokia Corporation, Nokia Solutions and Networks Oy, and Nokia of America Corporation (collectively, “NOKIA,” “Nokia,” or “Defendants”), and alleges as follows:

NATURE OF THE ACTION

1. This is a civil action arising under 35 U.S.C. § 271 for NOKIA’s infringement of Amazon’s U.S. Patent Nos. 11,516,080 (“the ’080 patent”); 11,425,194 (“the ’194 patent”); 9,329,909 (“the ’909 patent”); 8,296,419 (“the ’419 patent”); 9,253,211 (“the ’211 patent”); 9,621,593 (“the ’593 patent”); 9,106,540 (“the ’540 patent”); 8,117,289 (“the ’289 patent”); 9,766,912 (“the ’912 patent”); 9,756,018 (“the ’018 patent”); 11,336,529 (“the ’529 patent”); and 11,909,586 (“the ’586 patent”), together, the “Asserted Patents,” based on NOKIA’s unauthorized commercial manufacture, use, importation, offer for sale, and sale of products covered by the Asserted Patents.

PARTIES

2. Amazon Technologies, Inc. is a Nevada corporation with a principal place of business at 410 Terry Avenue North, Seattle, Washington 98109.

3. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the Asserted Patents, including the right to assert all causes of action arising under the Asserted Patents and the right to sue and obtain remedies for past, present, and future infringement.

4. On information and belief, defendant Nokia Corporation is a foreign corporation organized under the laws of Finland, with offices located at Karakaari 7, FIN-02610, Espoo, Finland.

5. On information and belief, defendant Nokia Solutions and Networks Oy is a foreign corporation organized under the laws of Finland, with offices located at Karaportti 3, FIN-02610, Espoo, Finland.

6. On information and belief, defendant Nokia of America Corporation is a Delaware corporation, with offices located at 600 Mountain Avenue, Murray Hill, New Jersey, 07974.

7. NOKIA makes, uses, sells, offers for sale and/or imports the following products or services in or into the United States and the State of Delaware: Nokia Airframe Data Center, Nokia CloudBand, Nokia CloudBand Application Manager, Nokia CloudBand Infrastructure Software, Nokia Container Services, Nokia Cloud Operations Manager, Nokia Nuage Networks Virtualized Cloud Services, Nokia Nuage Networks Virtualized Services Platform, and Nokia Nuage Software Defined Network (SDN) (the “Accused Products”).

JURISDICTION AND VENUE

8. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

9. This Court has subject matter jurisdiction over the patent infringement claims under 28 U.S.C. §§ 1331 and 1338(a).

10. This Court has personal jurisdiction over NOKIA because NOKIA conducts business in Delaware and because infringement by NOKIA has occurred and continues to occur in Delaware. This Court has personal jurisdiction over defendants Nokia Corporation and Nokia Solutions and Networks OY because each has operated, and continues to operate, websites and mobile applications within this jurisdiction. Through Nokia Corporation and Nokia Solutions and Networks OY websites and mobile applications, each has conducted, and continues to conduct, business by making, using, selling, offering for sale, and/or importing the Accused Products that infringe ATI's Asserted Patents. Upon information and belief, each of Nokia Corporation and Nokia Solutions and Networks OY has derived substantial revenue from the making, using, selling, offering for sale, and/or importing of the Accused Products within this jurisdiction, and have caused foreseeable harm and injury to ATI. This Court has personal jurisdiction over defendant Nokia of America Corporation because it is incorporated in the state of Delaware.

11. Venue is proper in this District pursuant to 28 U.S.C. §§ 1391 and 1400(b). Defendant Nokia of America Corporation is incorporated in the state of Delaware. NOKIA has transacted business in this District and has committed acts of infringement in this District.

BACKGROUND

12. Amazon Web Services (AWS) is the world's leading provider of cloud computing services. AWS offers customers access to an expansive suite of on-demand cloud-based products, including compute, storage, databases, analytics, networking, mobile, developer tools, management tools, IoT, security, and enterprise applications.¹

¹<https://docs.aws.amazon.com/whitepapers/latest/aws-overview/introduction.html>.

13. Since its inception in 2006, AWS has fundamentally transformed large-scale computing and internet communications. Before AWS launched its pioneering technology, large-scale computing relied predominantly on on-premise physical servers and software installations. This model required high up-front costs, required continuous maintenance, and lacked scalability. AWS's technology changed this by democratizing access to computing infrastructure and software through its cloud-based, on-demand services model. The breadth of offerings and flexible pricing structure provides enterprises, start-ups, small and medium-sized businesses, and customers in the public sector with tools that allow them to swiftly adapt to evolving business conditions and requirements.² This transformative model replaced bulky up-front capital infrastructure expenses with low variable costs, scaling up and down to meet business demands, eliminating the need for long-term infrastructure planning.³

14. Today, AWS's cloud computing technology is ubiquitous, used by companies large and small, as well as individual users, in 190 countries around the world.⁴

15. Over 90% of Fortune 100 companies and most Fortune 500 companies use AWS.⁵ AWS is also widely used by federal and local governments. Indeed, the NSA, Pentagon, U.S. Navy, and CIA are some of the top AWS customers.⁶

16. The concepts of cloud computing and virtualized computing resources date back to the 1960s.⁷ But it was Amazon⁸ that went on to become the acknowledged frontrunner in

²*Id.*

³*Id.* at Introduction.

⁴*Id.*

⁵<https://www.thomsondata.com/blog/fortune-500-companies-that-use-aws/>.

⁶<https://www.govconwire.com/articles/the-top-government-contracts-won-by-amazon-web-services/>.

⁷<https://www.dataversity.net/brief-history-cloud-computing/>.

⁸“Amazon” in the Background and Asserted Patents sections of this Complaint generally refers to Amazon.com, Inc., rather than ATI.

integrating and transforming these technologies into a service that is available to the general public, spearheaded by its introduction of Amazon Web Services (AWS) in 2006.⁹ Amazon developed its groundbreaking cloud computing technology as part of its rapid expansion in the early 2000s, when Amazon had begun onboarding thousands of software engineers. To avoid duplicating effort building fundamental components such as storage systems, compute resources, or databases, Amazon created a communal layer of infrastructure services.¹⁰

17. Amazon realized that other companies could benefit from the infrastructure that Amazon developed for its own purposes. In 2003, Amazon formed a team focused on converting Amazon's internal computing infrastructure into a business-centric cloud computing service available to others.¹¹ Amazon's vision for this service was akin to an operating system for the internet—allowing organizations, companies, or developers to run their applications using Amazon's technology infrastructure as the base.¹²

18. After three years of intensive internal development, Amazon unveiled its web infrastructure services to the public.¹³ Amazon launched the Simple Storage Service (S3) first, in March 2006, followed shortly by the Elastic Compute Cloud (EC2) five months later.¹⁴ Both services use distributed computing—S3 offers robust, scalable cloud-based storage, and EC2 provides secure, adjustable virtual computing capacities.¹⁵ Both products experienced immediate

⁹*Id.*

¹⁰<https://fortune.com/longform/amazon-web-services-ceo-adam-selipsky-cloud-computing/>.

¹¹*Id.*

¹²<https://thinkproduct.org/2022/10/09/aws-story-internets-os/>.

¹³<https://fortune.com/longform/amazon-web-services-ceo-adam-selipsky-cloud-computing/>.

¹⁴*Id.*

¹⁵<https://www.whizlabs.com/blog/amazon-ec2-vs-amazon-s3-comparison-guide/>.

success: over 12,000 developers signed up for S3 on the first day of its release.¹⁶ In 2007, Amazon premiered SimpleDB, bringing to life Amazon’s vision for creating core components of a web OS: computational capabilities, storage, and database.¹⁷

19. AWS has been growing ever since, the result of Amazon’s considerable investment and the hard work of thousands of engineers. Amazon has also focused its development efforts on technology that meets its customers’ business needs, consistent with one of the company’s key leadership principles—a “customer obsession.”¹⁸

20. Over the last decade, Amazon has continued to innovate, developing scores of new and successful cloud computing products, including CloudFront, CloudWatch, Auto Scaling, and Virtual Private Cloud. As a result, Amazon is the recipient of thousands of cloud computing patents awarded by the United States Patent & Trademark Office (“USPTO”). Amazon Technologies, Inc. owns and maintains Amazon’s intellectual property, including patents.

21. Amazon’s two decades of groundbreaking technological development in cloud computing stand in stark contrast to Nokia’s recent efforts. Nokia’s lengthy corporate history spans a wide range of businesses, from paper mills and rubber goods to electronics.¹⁹ Nokia is best known for its mass production of cellular telephones, first introduced in 1992.²⁰ However, with the advent of smartphones developed by Apple and Samsung, among others, Nokia’s

¹⁶<https://fortune.com/longform/amazon-web-services-ceo-adam-selipsky-cloud-computing/>.

¹⁷<https://www.aakashg.com/aws/>.

¹⁸<https://www.amazon.jobs/content/en/our-workplace/leadership-principles>.

¹⁹<https://www.sec.gov/Archives/edgar/data/924613/000155837021002363/nok-20201231x20f.htm>.

²⁰<https://history-computer.com/the-real-reason-nokia-failed-spectacularly/>.

prominence in the mobile phone market has plummeted.²¹ Nokia’s failure to anticipate the importance of smartphone technology led it to the verge of bankruptcy in 2013.²²

22. To save the company, Nokia exited the mobile device business in 2014—an act its board chairman referred to as a “moment of reinvention”²³—and pivoted to the sale of 5G network infrastructure and associated services that it acquired from Alcatel-Lucent in 2016.²⁴

23. In late 2020—nearly 15 years after AWS launched—Nokia announced a “new company strategy” focused on cloud computing, establishing a “Cloud and Network Services” business division.²⁵ Nokia attributes this latest transition to the “profound changes” within the industry, marked by trends favoring “open interfaces, virtualization, and cloud native software.”²⁶ In July 2020, Nokia entered the data center and switching business,²⁷ and in November 2021 entered the Software-as-a-Service (SaaS) market.²⁸

24. However, Nokia’s “new company strategy” involved leveraging *Amazon’s* innovative solutions, including Amazon’s patented technology, to address issues faced by cloud service providers. For instance, as discussed in detail below, Nokia CloudBand infringes Amazon patents related to configuring virtual machines and managing distributed application execution, in addition to autoscaling resources used during program execution. Similarly, Nokia Nuage

²¹*Id.*

²²*Id.*

²³<https://news.microsoft.com/2013/09/03/microsoft-to-acquire-nokias-devices-services-business-license-nokias-patents-and-mapping-services/>.

²⁴<https://www.sec.gov/Archives/edgar/data/924613/000155837021002363/nok-20201231x20f.htm>.

²⁵<https://www.nokia.com/about-us/news/releases/2020/10/29/nokia-announces-first-phase-of-its-new-strategy-changes-to-operating-model-and-group-leadership-team/>.

²⁶*Id.*

²⁷<https://www.fiercetelecom.com/telecom/make-way-cisco-arista-and-juniper-nokia-enters-data-center-switching-fray>.

²⁸<https://www.nokia.com/about-us/news/releases/2021/11/17/nokia-announces-entry-into-software-as-a-service-for-csps-with-multiple-services/>.

Networks infringes Amazon's patents related to managing communications in virtual networks and emulating physical network devices.

25. Amazon brings this suit against Nokia because it was Amazon that pioneered in the cloud, and now Nokia is using Amazon's patented cloud innovations without permission.

THE ASSERTED PATENTS

26. Amazon's pioneering work in the field of cloud computing is reflected in its extensive patent portfolio that includes the Asserted Patents. The Asserted Patents address a variety of challenges faced by cloud service providers that aim to build a highly configurable and dynamic architecture to meet the computing, networking, and storage needs of various organizations. The patented technologies make possible a highly configurable cloud platform that supports multiple virtual computing systems and virtual networks for multiple organizations using a shared substrate network based on physical computing systems, applications, and resources. The Asserted Patents provide many improvements to cloud computing techniques that can be classified broadly into different categories, although the techniques disclosed in each Asserted Patent may span multiple categories and may include categories not described here.

Virtual Networking Infrastructure: managing communications in virtual networks and emulating physical devices to create virtual devices that can be used for virtual networking.

Virtual Networking Security: managing security across virtual networks as well as providing secure access to virtual networks from remote locations.

Virtual Networking Performance: techniques for effectively supporting and exploiting virtual networking technology, for example, using network scaling.

Distributed Program Execution & Management: serverless computing for allowing cloud platforms to dynamically allocate, scale, and manage resources used for executing distributed programs in both virtual and physical environments.

VIRTUAL NETWORKING INFRASTRUCTURE

27. The following Asserted Patents provide and improve techniques for managing communications in virtual networks, for example, core communication constructs for implementing virtual computer networks such as forwarding or dropping communications, spoofing addresses of devices to manage virtual network communications, and more. This category of Asserted Patents further provides and improves techniques for emulating traditional networking hardware including routers and load balancers with virtual components.

A. The '540 patent

28. On August 11, 2015, the USPTO issued the '540 patent, titled "Providing logical networking functionality for managed computer networks." The '540 patent issued from U.S. patent application number 12/414,260, filed on March 30, 2009. A true and correct copy of the '540 patent is attached as Exhibit A.

29. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '540 patent, including the right to assert all causes of action arising under the '540 patent and the right to sue and obtain remedies for past, present, and future infringement.

30. The '540 patent concerns "providing logical networking functionality for managed computer networks, such as for virtual computer networks." ('540 patent at 2:6-8.)

31. The '540 patent improves virtualization technologies by "providing logical networking functionality for managed computer networks, such as for virtual computer networks." (*Id.* at 2:6-8.) The techniques disclosed in the '540 patent provide "logical networking functionality ... without physically implementing the network topology for the virtual computer network." (*Id.* at 2:18-21.) Accordingly, the techniques disclosed in the '540 patent provide

“logical networking functionality” for a virtual computer network “corresponding to [a] network topology ... without physically implementing the network topology.” (*Id.*)

32. The '540 patent specification explains that virtualization techniques, such as VMWare, XEN, or User-Model Linux, are beneficial for managing large-scale computing resources. (*Id.* at 1:34-44.) For example, these techniques may allow a single computing machine to be shared among multiple users by providing each user with a virtual machine hosted on a single computing machine. (*Id.*) Additionally, some virtualization techniques have the benefit of providing a single virtual machine with multiple virtual processes that span multiple distinct physical computing systems. (*Id.* at 1:44-48.) However, as explained by the '540 patent, managing physical computing resources has become increasingly complicated as the scale and scope of data centers and computer networks has increased. (*Id.* at 1:24-27.) The '540 patent addresses the need for an improved computing resource management technique and addresses this by providing a logical overlay, which allows the user to configure resources independently of a fixed physical network topology. The '540 patent thus improves upon non-configurable fixed network topologies.

33. The '540 patent provides a technical solution that improves upon existing virtualization technologies by allowing users to configure or specify a network topology, thereby allowing users to create virtual computer networks that may have network topologies distinct from the network topology of the underlying physical network. (*Id.* at Abstract, 12:3-11.) Thus, communications between multiple nodes of the virtual network “are managed so as to emulate functionality that would be provided by specified logical networking devices if they were physically present and/or to otherwise emulate functionality corresponding to a specified network topology if it was physically implemented.” (*Id.* at 2:21-28.) Such a technical solution provides

many technical advantages, such as: enabling an overlay of the virtual computer network on the physical substrate network (*id.* at 6:35-42), modifying with ease the number of computing nodes in a virtual computer network (*id.* at 6:62-67), enabling multiple computer networks to share parts of a physical network while maintaining network isolation for computing nodes of a particular virtual network (*id.* at 6:58-62), and scaling the underlying substrate network to include additional computing nodes regardless of their geographic locations (*id.* at 7:2-15).

34. Thus, the '540 patent's claimed techniques for providing logical networking functionality for computer networks are a concrete technical contribution and not simply the embodiment of an abstract idea. The '540 patent involves a specific system for implementing a virtual computer network without physically implementing the network topology that has concrete and valuable technical advantages in the field of virtual computer networks. (*Id.*)

35. The claimed elements of the '540 patent also provide an inventive concept individually and as an ordered combination. As the scale of data centers increased, the number of interconnected computing resources of the data centers increased. These data centers support complex distributed applications that require large numbers of computing resources and complex communications between the computing resources. However, traditional computer networks use physical routers, with a fixed topology, to implement networking functionality. The elements in the claims of the '540 patent provide a virtual overlay that allows users to configure different network topologies that are suited to their specific distributed applications. A user first provides "configuration information" for a "virtual computer network" that specifies a "virtual router device" that connects logical groups of computing nodes. (*Id.* at claim 4.) The system performs further actions to implement the virtual network topology, including by "emulating functionality of the specified virtual router device," for example to perform actions such as responding

to requests for information and other actions handled by router devices. (*Id.*) Further, the system forwards communications between groups of computing nodes by “emulating further functionality of the specified virtual router device for modifying the communication.” (*Id.*) By performing these actions, the system implements a logical network topology overlayed on a physical network with a fixed network topology. These actions and the ordered combination of elements together provide a technological improvement over the conventional technology at the time resulting in benefits of a user configurable logical network with greater flexibility to address complex resource requirements in a large-scale network.

B. The '289 patent

36. On February 14, 2012, the USPTO issued the '289 patent, titled “Using virtual networking devices to manage substrate devices.” The '289 patent issued from U.S. patent application number 13/091,986, filed on April 21, 2011. A true and correct copy of the '289 patent is attached as Exhibit B.

37. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '289 patent, including the right to assert all causes of action arising under the '289 patent and the right to sue and obtain remedies for past, present, and future infringement.

38. The '289 patent concerns “providing managed virtual computer networks [with] configured logical network topology [with] virtual networking devices.” ('289 patent at Abstract.)

39. The '289 patent improves virtualization technologies by “providing virtual networking functionality for managed computer networks” with “configured logical network topology.” (*Id.* at 2:7-18.) The techniques disclosed in the '289 patent “handle communications between computing nodes of a managed computer network in accordance with its specified

network topology by emulating at least some types of functionality that would be provided by virtual networking devices.” (*Id.* at 2:35-42.)

40. The '289 patent specification explains that “virtualization technologies for commodity hardware” provided “benefits with respect to managing large-scale computing resources.” (*Id.* at 1:36-40.) Prior to the '289 patent, providing, administering, and managing large-scale computing resources associated with data centers and computer networks had become increasingly complicated. (*Id.* at 1:31-35.) Computer networks operated by companies and other organizations interconnected numerous computing systems, for which data centers housing the systems provided different levels of service depending on their geographic location and operating entity. (*Id.* at 1:15-31.) Existing virtualization technologies for commodity hardware provided some benefits but allowed only a single physical computing machine to be shared among multiple users. (*Id.* at 1:36-51.) Thus, there was no well-understood, routine, and conventional way to manage large-scale computer networks effectively and securely. (*See id.* at 1:15-52.)

41. The '289 patent improves upon the prior art by inventing techniques performed by a Network Routing Manager module that enable a user to configure or otherwise specify a network topology for a managed computer network provided by a configurable network service and subsequently provide the user virtual networking functionality corresponding to the specified topology. (*Id.* at 2:19-56.) As the '289 patent explains, the managed computer network may be a virtual computer network overlaid on an underlying substrate network that includes or otherwise provides access to particular devices that are of use to the networking functionality to be provided. (*Id.* at 3:35-48.) If a particular managed computer network is configured to include one or more devices providing a particular type of networking-related functionality, and there are network-accessible devices available via the substrate network that provide that particular type of

networking-related functionality, modules of the configurable network service may operate to route appropriate communications via the substrate network through those network-accessible devices. (*Id.* at 3:48-58.) These concrete techniques enable a user or entity to configure or otherwise specify one or more networking devices for a managed computer network that are supported by various automated actions. (*Id.* at 3:62-4:8.) The technological solutions offered by the '289 patent thus achieve the goal of enabling users or entities to effectively and securely manage large-scale computer networks provided by a configurable network service.

42. The claimed elements of the '289 patent also provide an inventive concept individually and as an ordered combination. These elements encompass technical solutions and improvements including, but not limited to, the claimed elements “providing the configured virtual computer network for the client in accordance with the configuring by overlaying the virtual computer network on a distinct substrate network” and “forwarding multiple communications between the multiple computing nodes in accordance with the configuring, the forwarding including routing at least one of the multiple communications to at least one of the selected devices to enable the at least one selected device to provide the indicated type of functionality for the at least one communication.” (*Id.* at claim 20.) The ordered combination of these steps is also inventive because it was not well-understood, routine, or conventional at the time of invention to receive information for use in virtual computer network configuration by specifying interconnections between multiple computing nodes and providing a functionality for handling communications between the nodes. Moreover, the '289 patent provides the concrete, technical improvement of providing the configured network by selecting specific network devices and subsequently routing communications between the nodes to the devices.

C. The '586 patent

43. On February 20, 2024, the USPTO issued the '586 patent, titled “Managing communications in a virtual network of virtual machines using telecommunications infrastructure systems.” The '586 patent issued from U.S. patent application number 18/047,239, filed on October 17, 2022. A true and correct copy of the '586 patent is attached as Exhibit C.

44. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '586 patent, including the right to assert all causes of action arising under the '586 patent and the right to sue and obtain remedies for past, present, and future infringement.

45. The '586 patent concerns “logical networking functionality for managed computer networks, such as for virtual computer networks provided on behalf of users or other entities.” ('586 patent at Abstract.)

46. The '586 patent improves virtual networking technologies by “embedding virtual network address information in substrate network addresses for an underlying physical substrate network,” and providing benefits such as “limiting communications to and/or from computing nodes of a particular virtual computer network to other computing nodes that belong to that virtual computer network” and “allow[ing] computing nodes to easily be added to and/or removed from a virtual computer network, such as to allow a user to dynamically modify the size of a virtual computer network.” (*Id.* at 6:51-7:18.) The system “supports changes to an underlying substrate network—for example, if the underlying substrate network is expanded to include additional computing nodes at additional geographical locations, existing or new virtual computer networks being provided may seamlessly use those additional computing nodes.” (*Id.* at 7:18-32.)

47. The '586 patent specification explains that virtualization techniques, such as VMWare, XEN, or User-Model Linux, are beneficial for managing large-scale computing

resources. (*Id.* at 1:45-55.) For example, these techniques may allow a single computing machine to be shared among multiple users by providing each user with a virtual machine hosted on a single computing machine. (*Id.* at 1:33-44.) Additionally, some virtualization techniques have the benefit of providing a single virtual machine with multiple virtual processes that span multiple distinct physical computing systems. (*Id.* at 1:44-48.) However, as explained by the '586 patent, managing physical computing resources has become increasingly complicated as the scale and scope of data centers and computer networks has increased. (*Id.* at 1:35-39.) The '586 patent addresses the need for an improved computing resource management technique and addresses this by providing a logical overlay, which allows the user to configure resources independently of a fixed physical network topology. The '586 patent thus improves upon non-configurable fixed network topologies.

48. The '586 patent provides a technical solution that improves upon existing virtualization technologies by allowing users to configure or specify a network topology, thereby allowing users to create virtual computer networks that may have network topologies distinct from the network topology of the underlying physical network. (*Id.* at Abstract, 12:3-11.) Thus, communications between multiple nodes of the virtual network “are managed so as to emulate functionality that would be provided by specified logical networking devices if they were physically present and/or to otherwise emulate functionality corresponding to a specified network topology if it was physically implemented.” (*Id.* at 2:32-40.) One such technique to implement this solution is by spoofing a response to an ARP request that includes a virtual hardware address for a computing node instead of a hardware address (*e.g.*, a 48-bit MAC address) for the computing node. (*Id.* at 14:21-26.) Such a technical solution provides many benefits, such as: enabling an overlay of the virtual computer network on the physical substrate network (*id.* at 6:35-42),

modifying with ease the number of computing nodes in a virtual computer network (*id.* at 6:62-67), enabling multiple computer networks to share parts of a physical network while maintaining network isolation for computing nodes of a particular virtual network (*id.* at 6:58-62), and scaling the underlying substrate network to include additional computing nodes regardless of their geographic locations (*id.* at 7:2-15).

49. The claimed elements of the '586 patent also provide an inventive concept individually and as an ordered combination. As the scale of data centers increased the number of interconnected computing resources of the data centers increased. These data centers support complex distributed applications that require large numbers of computing resources and complex communications between the computing resources. However, traditional computer networks use physical routers to implement networking functionality, making network topology fixed and non-configurable. The following claimed elements allow users to configure different network topologies that are suited to their specific distributed applications: “based at least in part on the configuration information, modifying the first communication and forwarding the first communication to the computing node,” and “based at least in part on the configuration information, dropping the second communication without forwarding the second communication to the computing node,” and “in response to an address resolution protocol (ARP) communication from the computing node regarding a second computing node in the virtual computer network, sending a spoofed response to the ARP communication indicating a virtual hardware address of the second computing node.” (*Id.* at claim 1.) These actions and the ordered combination of elements together provide a technological improvement over the conventional technology, and the technical advantages of a user-configured logical network.

D. The '529 patent

50. On May 17, 2022, the USPTO issued the '529 patent, titled “Providing virtual networking device functionality for managed computer networks.” The '529 patent issued from U.S. patent application number 16/798,070, filed on February 21, 2020. A true and correct copy of the '529 patent is attached as Exhibit D.

51. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '529 patent, including the right to assert all causes of action arising under the '529 patent and the right to sue and obtain remedies for past, present, and future infringement.

52. The '529 patent concerns “providing virtual networking functionality for managed computer networks.” ('529 patent at Abstract.)

53. The '529 patent improves virtual networking technologies by “embedding virtual network address information in substrate network addresses for an underlying physical substrate network,” and providing benefits such as “limiting communications to and/or from computing nodes of a particular virtual computer network to other computing nodes that belong to that virtual computer network” and “allow[ing] computing nodes to easily be added to and/or removed from a virtual computer network, such as to allow a user to dynamically modify the size of a virtual computer network.” (*Id.* at 11:24-59.) The system “supports changes to an underlying substrate network—for example, if the underlying substrate network is expanded to include additional computing nodes at additional geographical locations, existing or new virtual computer networks being provided may seamlessly use those additional computing nodes.” (*Id.* at 11:24-12:2.)

54. The '529 patent describes techniques for “providing virtual networking functionality.” (*Id.* at Abstract.) As described in the '529 patent, “logical sub-networks” are specified for a “virtual computer network” with a “substrate network functionality used to emulate

various functionality corresponding to the specified logical subnets.” (*Id.* at 12:7-11.) The techniques disclosed in the ’529 patent “configure, authorize and otherwise manage communications sent to and from associated computing nodes” and “support providing various virtual networking functionality for ... virtual computer networks.” (*Id.* at 14:29-33.)

55. The techniques disclosed “provide various ... benefits in various situations, such as limiting communications to and/or from computing nodes of a particular virtual computer network to other computing nodes that belong to that virtual computer network.” (*Id.* at 11:42-48.) The techniques disclosed improve upon the technology of virtual computer networks, for example, “[b]y not delivering unauthorized communications to computing nodes, network isolation and security of entities’ virtual computer networks is enhanced.” (*Id.* at 15:67-16:2.) The described techniques improve upon the technology of virtual computer networks by “provid[ing] functional decomposition and/or isolation for the various component types” and “virtualiz[ing] physical networks to reflect almost any situation that would conventionally necessitate physical partitioning of distinct computing systems and/or networks.” (*Id.* at 66:39-67:6.)

56. As disclosed in the ’529 patent, “[n]etwork access constraint information” may be “configured for a provided computer network in various manners.” (*Id.* at 9:44-46.) For example, “a client may specify information about whether and how some or all of the computing nodes of a provided computer network are allowed to communicate with other computing nodes of the provided computer network and/or with other external computing systems” based on various types of information such as “types of communication protocols used, such as to allow HTTP requests for text but not images and to not allow FTP requests.” (*Id.* at 9:46-57.) As described in the ’529 patent the “Communication Manager module” receives “outgoing communications” and “determines whether to authorize the sending of the outgoing communication” based on

“information about the sending virtual machine computing node” or “the destination virtual machine computing node.” (*Id.* at 15:54-67.)

57. The claimed elements of the '529 patent also provide an inventive concept individually and as an ordered combination of elements, in that they provide the technical solution and advantages of a virtual computer network. In a virtual computer network, a computing node may be “assigned one or more virtual network addresses for the virtual computer network that are unrelated to those computing nodes’ substrate network addresses.” (*Id.* at 3:37-50.) As a result, traditional networking protocols such as Address Resolution Protocol (ARP) and any techniques based on such protocols may not work in a virtual computer network. The claims of the '529 patent improve upon computer networking technology to modify the ARP protocol and further use it to implement access control policies. As recited in the claims of the '529 patent, a configurable network service creates a virtual computer network based on client requests and creates a logical sub-network. (*Id.* at claim 1.) After creating the virtual computer network, the configurable network service manages communications of the virtual computer network by intercepting an ARP request sent by a virtual machine and responding to the ARP request by providing a MAC address. (*Id.*) The configurable network service further receives frames comprising the MAC address from the virtual machine and implements an access control policy by allowing or denying communications based on information such as source, destination, direction, or protocol used for the communication. (*Id.*) These actions performed in the ordered combination provide several technical advantages: limiting communications to and/or from computing nodes of a virtual computer network to other computing nodes; sharing parts of the network while still providing network isolation; allowing computing nodes to be added to or removed from a virtual network;

and allowing users to perform technically-beneficial modifications, such as modify the underlying substrate network while the logical network presents a consistent appearance. (*Id.* at 11:42-65.)

VIRTUAL NETWORKING SECURITY

58. The following Asserted Patents provide and improve techniques for managing security in virtual networks, including security of communications between multiple virtual networks managed by a cloud platform as well as providing secure access to virtual networks from client devices operating in remote locations.

E. The '018 patent

59. On September 5, 2017, the USPTO issued the '018 patent, titled “Establishing secure remote access to private computer networks.” The '018 patent issued from U.S. patent application number 15/179,700, filed on June 10, 2016. A true and correct copy of the '018 patent is attached as Exhibit E.

60. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '018 patent, including the right to assert all causes of action arising under the '018 patent and the right to sue and obtain remedies for past, present, and future infringement.

61. The '018 patent concerns “access to private computer networks” for allowing “users to interact with a remote configurable network service to create and configure computer networks.” ('018 patent at Abstract.)

62. The '018 patent improves security in virtual computer networks by “providing users with access to computer networks, such as under the control of a configurable network service available to remote users.” (*Id.* at 2:21-23.) Furthermore, “[a]fter configuring such a computer network, the user may interact from one or more remote locations with the computer network being

provided to the user by the configurable network service, such as to execute programs on the computing nodes of the provided computer network.” (*Id.* at 2:30-35.)

63. The '018 patent describes techniques for “[e]stablishing secure remote access” to “computer networks.” (*Id.* at Title.) The '018 patent specification describes that “virtualization technologies” allow management of “large-scale computing resources” thereby “allowing various computing resources to be efficiently and securely shared by multiple customers.” (*Id.* at 1:44-47.) However, these approaches do not provide a secure solution for remote users to interact with virtual resources.

64. The '018 patent provides a technological solution that allows “a remote user [to] interact with a configurable network service over public networks in order to create and configure a computer network.” (*Id.* at 2:24-30.) The computer network of the network service is configured according to “network topology information” received from the remote user. (*Id.*) The “configuration information” received from the remote user may “include routing information or other interconnectivity information between networking devices and/or groups of computing devices.” (*Id.* at 9:53-57.) The remote user configures a computer network and “interact[s] from one or more remote locations with the computer network being provided to the user.” (*Id.* at 2:30-35.)

65. The '018 patent techniques provide the technical advantage of enhanced security for remote users of the computer networks. For example, the '018 patent techniques “manage communications between the computing nodes” and “external computing systems” in order to “enforce specified network access constraints, as well as to manage configured access mechanisms for remote resource services and secure connections to remote client private computer networks.” (*Id.* at 10:36-43.)

66. The claimed elements of the '018 patent also provide an inventive concept individually and as an ordered combination. Organizations use private networks for communication between computing devices. (*Id.* at 1:24-30.) However, access may have to be provided to such computing devices from remote locations, thereby increasing the likelihood of malicious attacks on the private network. (*Id.* at 18:53-63.) These problems are further exacerbated in virtual networks since conventional networking techniques often do not work on virtual networks as computing devices are mapped to virtual network addresses that are distinct from their physical network addresses. The claims of the '018 patent provide secure access to virtual computer networks. (*Id.* at claim 18.) The system provides a computer network to a client by performing an ordered sequence of actions including selecting multiple computing nodes, provisioning the selected computing nodes for use in a computer network, and configuring hardware devices to route communications according to a network topology specified by a user. (*Id.*) After providing the computer network, the system receives a request for a secure connection to the computer network from a remote location. (*Id.*) After receiving the request, the system responds by providing the configuration information to allow devices in the remote location to participate in the secure connection. (*Id.*) These actions and the configuration interface of the ordered combination together provided an advancement over the conventional networking technology at the time resulting in multiple benefits based on a user's configuration specification including beneficially "provid[ing] private or other specialized access to one or more remote resource services." (*Id.* at 13:7-10.)

F. The '080 patent

67. On November 29, 2022, the USPTO issued the '080 patent, titled "Using virtual networking devices and routing information to associate network addresses with computing

nodes.” The ’080 patent issued from U.S. patent application No. 17/119,944, filed on December 11, 2020. A true and correct copy of the ’080 patent is attached as Exhibit F.

68. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the ’080 patent, including the right to assert all causes of action arising under the ’080 patent and the right to sue and obtain remedies for past, present, and future infringement.

69. The ’080 patent “provid[es] virtual networking functionality for managed computer networks” that may be accessed by clients “from remote locations over one or more intervening networks.” (’080 patent at 2:13-23.)

70. The ’080 patent provides improvements to virtualization technologies by “providing virtual networking functionality for managed computer networks.” (*Id.* at 2:13-16.) The techniques disclosed in the ’080 patent “enable a user to configure ... a network topology for a managed computer network” by “separat[ing] multiple computing nodes of the managed computer network into multiple logical sub-networks” and providing “networking functionality corresponding to the ... network topology.” (*Id.* at 2:23-40.) For example, a physical computer network is “emulate[ed]” using “virtual networking devices” without “physically implementing the ... network topology” and “without physically providing those networking devices.” (*Id.* at 2:40-48.)

71. Prior to the ’080 patent, providing, administering, and managing large-scale computing resources associated with data centers and computer networks had become increasingly complicated. (*Id.* at 1:36-40.) Computer networks operated by companies and other organizations interconnected numerous computing systems, for which data centers housing the systems provided different levels of service depending on their geographic location and operating entity. (*Id.* at 1:19-36.) Existing virtualization technologies for commodity hardware provided some benefits

but allowed only a single physical computing machine to be shared among multiple users. (*Id.* at 1:41-57.) Thus, there was no well-understood, routine, and conventional way to manage large-scale computer networks provided by a configurable network service effectively and securely. (*See id.* at 1:19-57.)

72. The '080 patent improves upon the prior art by inventing techniques that enable a user to configure or otherwise specify a network topology for a managed computer network and subsequently provide the user virtual networking functionality corresponding to the specified topology. (*Id.* at 2:23-48.) Such functionality is provided by allowing at least some computing nodes of the managed computer network to dynamically signal particular types of uses of one or more indicated target network addresses based on routing information in routing communications received from the computing nodes. (*Id.* at 4:32-38.) The configurable network service can then use this information to dynamically configure the network topology of the managed computer network. (*Id.* at 4:39-5:9.) These techniques provide benefits including enabling an overlay of the virtual computer network on the substrate network without encapsulating communications or configuring physical networking devices of the substrate network; limiting communications to and from computing nodes of a particular virtual computer network to other computing nodes that belong to that virtual computer network; and allowing a user to dynamically modify the size of a virtual computer network. (*Id.* at 14:41-15:23.) The technological solutions offered by the '080 patent thus overcome the shortcomings of the prior art by disclosing specific and concrete techniques for users or entities to effectively and securely manage large-scale computer networks provided by a configurable network service.

73. The claimed elements of the '080 patent also provide an inventive concept individually and as an ordered combination. These elements encompass new technical solutions

and improvements including “providing the first virtual computer network to the first client according to the first configuration information, wherein the first virtual computer network is overlaid on a substrate network of the configurable network service,” “assigning one of the network addresses in the first range to one of the computing nodes in the first virtual computer network,” and “routing, using the virtual peering router, network traffic from the computing nodes of the first virtual computer network to the computing nodes of the second virtual computer network over the substrate network.” (*Id.* at claim 1.) The ordered combination of these steps is inventive because it was not well-understood, routine, or conventional at the time of invention to associate providing different virtual computer networks to different clients according to different sets of received configuration information together with assigning network addresses from the configuration information to and subsequently routing network traffic between the computing nodes of the different networks. The ordered combination of these steps thus provides a technical solution allowing computing nodes of the network to dynamically signal particular types of uses of indicated target network addresses based on routing information, which can further be used to dynamically configure the topology of the network. (*Id.* at 4:39-56.)

VIRTUAL NETWORKING PERFORMANCE

74. The following Asserted Patent provides and improves techniques for effectively supporting and exploiting virtual networking technology, for example, by improving the efficiency of virtual networks in cloud platforms.

G. The '912 patent

75. On September 19, 2017, the USPTO issued the '912 patent, titled “Virtual machine configuration.” The '912 patent issued from U.S. patent application number 13/686,683, filed on November 27, 2012. A true and correct copy of the '912 patent is attached as Exhibit G.

76. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '912 patent, including the right to assert all causes of action arising under the '912 patent and the right to sue and obtain remedies for past, present, and future infringement.

77. The '912 patent concerns “launching a virtual machine and then configuring the virtual machine after launch.” ('912 patent at Abstract.)

78. The '912 patent improves virtualization technologies by allowing “launching [of] a virtual machine” using a virtual machine image “and then configuring the virtual machine after launch” with “metadata configuration data.” (*Id.* at Abstract.) Prior to the '912 patent, snapshots were used “to create a new virtual machine image,” but “updates [were] often ignored.” (*Id.* at 1:23-29.) The techniques disclosed in the '912 patent provide the “advantages of using a snapshot, but with updates applied to the VM image.” (*Id.* at 5:21-23.) A “virtual machine image configuration” is “used to launch a virtual machine” and “metadata configuration information” is “used to further configure the virtual machine after launch.” (*Id.* at Abstract.) This allows “third-party vendor[s]” to apply “updates [to] the VM image.” (*Id.* at 5:21-23.)

79. The '912 patent describes techniques for “[v]irtual machine configuration.” (*Id.* at Title.) As described in the '912 patent, virtual machines are configured using a “virtual machine image” that contains “data needed to launch a virtual machine in a virtual environment.” (*Id.* at 1:5-7.) Traditional techniques for configuring virtual machines use “[s]napshots ... to create a new virtual machine image.” (*Id.* at 1:23-24.) However, these techniques use “virtual machine images made from snapshots” that ignore “updates.” (*Id.* at 1:28-29.) Such an approach has drawbacks, “[f]or example, if Linux has a new security patch release, a cloud provider may update the source virtual machine image, but any already-created snapshot[s]” are not updated. (*Id.* at 1:29-35.) Alternate approaches to solve these problems also have drawbacks, for example,

significant overhead is required if a “third party” “takes the updated source virtual machine image, adds the same components to it and again re-packages it as a new virtual machine image.” (*Id.*)

80. The ’912 patent provides a technological solution to these problems by storing “[m]etadate configuration information” “in conjunction with virtual machine image configuration data.” (*Id.* at Abstract.) The “virtual machine image configuration data” is used to “launch a virtual machine” and “metadate configuration information” is used to “further configure the virtual machine after launch.” (*Id.*)

81. The techniques disclosed in the ’912 patent provide several technical advantages. For example, the VM image is improved upon in that it is customizable and more easily controlled by the user. Specifically, the VM image can be configured differently for different users by allowing a “third-party vendor” to inject “metadate configuration information ... into the VM registration record to change how a VM is configured upon launch by” different users. (*Id.* at 5:3-6.) The VM image can be changed in a manner that is “transparent to a user, as the user continues to use the same VM image identifier” and a “third-party vendor” obtains “the advantages of using a snapshot, but with updates applied to the VM image.” (*Id.* at 5:16-23.) Accordingly, the ’912 patent provides a technical improvement to the process of installing VM images by allowing changes to the VM images in a manner transparent to the user.

82. Thus, the ’912 patent’s claimed methods and systems for providing logical networking functionality for computer networks are a concrete technical contribution and not simply the embodiment of an abstract idea. They involve a specific system for launching a VM image and using metadate configuration information after a VM image is launched to customize the virtual machine.

83. The claimed elements of the '912 patent also provide an inventive concept individually and as an ordered combination. Prior art techniques to launch VMs used a snapshot of the VM that created a copy of the volume of the VM at a given point in time. The snapshot preserved the state of the virtual machine and therefore ignored subsequent updates to the VM image unless the VM image was repackaged with the recent updates. (*Id.* at 1:28-35.) The ordered combination of claimed elements in the claims of the '912 patent established techniques to allow the launched VM to be updated without requiring the VM image to be repackaged. As a first action, the system determines a virtual machine image, a virtual machine image configuration, and metadata configuration information, together associated with launching the virtual machine. (*Id.* at claim 10.) Both the virtual machine image configuration and the metadata configuration information are identified by a single Application Programming Interface request. As a second action, the system next launches the virtual machine using the virtual machine image and the virtual machine image configuration. (*Id.*) As a third action, the system supplies the metadata configuration information to the launched virtual machine as part of the launching process so that the virtual machine can use the metadata configuration information to customize itself after launching. (*Id.*) The claimed elements beneficially overcome the problem of the prior art techniques that used snapshots, for example, by allowing a third-party vendor to inject metadata configuration information into a VM registration record while a service center updates a pointer to a virtual machine image. (*Id.* at 4:57-62.) These actions and their ordered combination together improve upon the conventional technology and create the technical advantage of allowing the VM image to be updated transparently for the user, since the user used the same VM Image identifier. (*Id.* at 5:16-23.)

DISTRIBUTED PROGRAM EXECUTION & MANAGEMENT

84. The following Asserted Patents provide and improve techniques for implementing serverless computing by allowing cloud platforms to dynamically allocate and manage resources used to execute distributed programs on the cloud platform. These techniques include autoscaling of virtual machines used to execute distributed programs, automatically terminating distributed programs as necessary, effectively placing and distributing distributed programs, load balancing, and other techniques.

H. The '194 patent

85. On August 23, 2022, the USPTO issued the '194 patent, titled "Dynamically modifying a cluster of computing nodes used for distributed execution of a program." The '194 patent issued from U.S. patent application number 17/128,746, filed on December 21, 2020. A true and correct copy of the '194 patent is attached as Exhibit H.

86. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '194 patent, including the right to assert all causes of action arising under the '194 patent and the right to sue and obtain remedies for past, present, and future infringement.

87. The '194 patent concerns "managing distributed execution of programs" by "dynamically modifying the distributed program execution in various manners," such as, for example, "adding and/or removing computing nodes from a cluster that is executing the program, modifying the amount of computing resources that are available for the distributed program execution, terminating or temporarily suspending execution of the program (*e.g.*, if an insufficient quantity of computing nodes of the cluster are available to perform execution), etc." ('194 patent at Abstract.)

88. The '194 patent provides improvements to virtualization technologies by “dynamically monitoring the ongoing distributed execution of a program on a cluster of multiple computing nodes and dynamically modifying the distributed program execution in various manners based on the monitoring.” (*Id.* at 2:30-36.) The dynamic monitoring includes “determining the status of execution of the program on each of the multiple computing nodes and/or determining the aggregate usage of one or more types of computing resources across the cluster of multiple computing nodes by the distributed program execution.” (*Id.* at 2:36-41.) The distributed program execution is dynamically modified by “adding and/or removing computing nodes from the cluster that is executing the program, modifying the amount of computing resources that are available for the distributed program execution, temporarily throttling usage of computing resources by the distributed program execution,” or “terminating or temporarily suspending execution of the program.” (*Id.* at 2:41-54.)

89. Prior to the '194 patent, there had been rapid growth in distributed computing and “data centers housing significant numbers of interconnected computing systems [had] become commonplace.” (*Id.* at 1:35-41.) The '194 patent explains that “as the scale and scope of typical data centers has increased, the task of provisioning, administering, and managing the physical computing resources has become increasingly complicated.” (*Id.* at 1:46-49.) Without proper management of the physical computing resources, programs could use more resources than they should be allocated, becoming bottlenecks for the entire system. (*Id.* at 2:41-55, 30:41-50.) In other scenarios, programs could fail to execute if “insufficient quantity of computing nodes of the cluster are available to perform execution” or to complete execution in the allocated timeframe. (*Id.* at 2:41-54, 34:66-35:7.) These problems were compounded by virtualization technologies that allowed “a single physical computing machine to be shared among multiple users” or “a single

virtual machine with multiple virtual processors that actually spans multiple distinct physical computing systems.” (*Id.* at 1:54-2:2.) There was no well-understood, routine, and conventional way to effectively provision, administer, and manage the physical computing resources for virtual machines in large-scale distributed systems to address these growing issues.

90. The ’194 patent provides a technical solution to communicate program configuration information, monitor the status of program execution, and dynamically allocate resources for virtual machines based on that monitoring information in large-scale distributed computing systems. (*Id.* at 2:29-55.)

91. The ’194 patent describes the Distributed Program Execution Service System Manager (“DPESSM”) that provides a service for distributed program execution by “executing multiple programs on behalf of multiple customers.” (*Id.* at 2:56-3:4.) The DPESSM manages the computing nodes allocated to customers, allows customers to communicate program information and execution requirements, and configures and initiates execution of the programs in the distributed system. (*Id.* at 8:20-8:49, 8:50-9:4.) For example, the ’194 patent’s claim 1 recites “receiving, by one or more configured computing systems of a program execution service that provides computing resources available to multiple users of the program execution service, instructions from a first user to execute a program using specified configuration information, the specified configuration information including a first number of virtual machines (VMs) to use to execute the program and a set of instructions specifying how to modify a number of VMs used during execution of the program based on resource utilization metrics.” The ability to receive specific program execution and configuration information related to computing resources for virtual machines on a distributed system was a concrete, technological improvement over the prior art. This type of personalized configuration allows users to specify the computing resources

needed to execute a program, but also allows users to limit the computing resources consumed to limit costs. This technical advantage over the prior art was recognized in the industry. For example, by using the AWS cloud, companies were able to efficiently use “many GPU cores, CPU cores, and AWS instances” to train “multiple models either from different datasets or configurations.” (Alex Chen, et al., Distributed Neural Networks with GPUs in the AWS Cloud, Medium (Feb. 10, 2014), <https://netflixtechblog.com/distributed-neural-networks-with-gpus-in-the-aws-cloud-ccf71e82056b>.)

92. The DPESSM also solved the technological problem of programs failing to execute due to insufficient resources by monitoring all computing nodes, storing information related to physical resources of each computing node, and dynamically allocating additional computing nodes to allow the program to complete execution or temporarily throttling usage of computing resources. (’194 patent at Figs. 2A-2C, 2:41-54, 13:41-14:16, 14:46-15:6, 17:17-44, 18:59-19:5, 22:1-46.) For example, claim 1 recites “monitoring, by the one or more configured computing systems during the execution of the program, a resource utilization of the group of VMs, wherein the resource utilization is based on a measured amount of a resource used by the group of VMs.” (*Id.* at claim 1.) The ’194 patent explains that “[s]uch monitoring of the execution of execution jobs may provide various benefits, such as to determine when to later initiate execution of other execution jobs.” (*Id.* at 19:60-20:9.) This monitoring and tracking of performance of the distributed program execution also “provide[s] various benefits, such as to enable the ongoing intermediate execution and data state from the partial execution of the execution job to be tracked and used” in re-allocating node clusters to increase or decrease computing resources within the distributed system. (*Id.* at 4:14-38, 20:9-23.)

93. The claimed elements of the '194 patent also provide an inventive concept individually and as an ordered combination. The technological solutions provided by the '194 patent solve problems specific to large-scale distributed systems that effectively provision, administer, and manage distributed program execution by dynamically modifying clusters of computing nodes running virtual machines. For example, claim 1 recites “modifying, by the one or more configured computing systems, a quantity of VMs in the group of VMs for use in further execution of the program, wherein the modifying includes adding one or more additional VMs to the group of VMs while the execution of the program is ongoing and using the one or more additional VMs for further execution of the program, wherein adding the one or more additional VMs to the group of VMs includes allocating computing resources of one or more physical computing systems to the one or more additional VMs.” (*Id.* at claim 1.) It was an improvement upon the prior art to modify a quantity of VMs in a group of VMs for use in further execution of a program by adding one or more additional VMs while the execution of the program is ongoing and furthermore, using the additional VMs for further execution of the program, and allocating computing resources of physical computing systems to the additional VMs. This technical advantage of dynamically adjusting computing resources based on configuration or demand was, and is, widely recognized as one of the main benefits of using the AWS Cloud over other solutions. The industry widely recognized this benefit in “us[ing] a reasonable number of machines to implement a powerful machine learning solution” through program customization and by “leverag[ing] the full, on-demand computing power we can obtain from AWS.” (Alex Chen, et al., Distributed Neural Networks with GPUs in the AWS Cloud, Medium (Feb. 10, 2014), <https://netflixtechblog.com/distributed-neural-networks-with-gpus-in-the-aws-cloud-ccf71e82056b>.) The claim elements, and the ordered combination of each of the elements of the

claim, provide an inventive technological solution for the particular technological environment of virtual machines in a distributed computing, that solved critical shortcomings in the prior art.

I. The '909 patent

94. On May 3, 2016, the USPTO issued the '909 patent, titled “Dynamically modifying a cluster of computing nodes used for distributed execution of a program.” The '909 patent issued from U.S. patent application number 13/620,805, filed on September 15, 2012. A true and correct copy of the '909 patent is attached as Exhibit I.

95. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '909 patent, including the right to assert all causes of action arising under the '909 patent and the right to sue and obtain remedies for past, present, and future infringement.

96. The '909 patent concerns “distributed execution of programs,” for example, by “dynamically modifying a cluster of computing nodes used for distributed execution of a program.” ('909 patent at Abstract, Title.)

97. The '909 patent improves virtualization technologies by allowing users to “manag[e] distributed execution of programs” by “dynamically monitoring the ongoing distributed execution of a program on a cluster of multiple computing nodes, and dynamically modifying the distributed program execution.” (*Id.* at 2:21-27.) The “dynamic modifying of the distributed program execution” provides several advantages, for example “[c]luster expansion” may be performed “to enable program execution to complete sooner,” for example, “if execution on one or more cluster computing nodes is taking longer than expected” “additional computing nodes” provide access to “additional computing resources that were lacking.” (*Id.* at 2:33-34, 4:5-35.) Similarly, “[c]luster shrinking” may be performed to “efficiently use resources,” for example, “if the distributed program execution is progressing faster than expected.” (*Id.* at 4:29-35.)

98. Prior to the '909 patent, there had been rapid growth in distributed computing and “data centers housing significant numbers of interconnected computing systems [had] become commonplace.” (*Id.* at 1:27-29.) The '909 patent explains that “as the scale and scope of typical data centers has increased, the task of provisioning, administering, and managing the physical computing resources has become increasingly complicated.” (*Id.* at 1:38-41.) Without proper management of the physical computing resources, programs could use more resources than they should be allocated, becoming bottlenecks for the entire system. (*Id.* at 2:37-43, 30:45-53.) In other scenarios, programs could fail to execute if “insufficient quantity of computing nodes of the cluster are available to perform execution” or complete execution in the allocated timeframe. (*Id.* at 2:44-46, 34:1-9.) These problems were compounded by virtualization technologies that allowed “a single physical computing machine to be shared among multiple users” or “a single virtual machine with multiple virtual processors that actually spans multiple distinct physical computing systems.” (*Id.* at 1:46-61.) There was no well-understood, routine, and conventional way to effectively provision, administer, and manage the physical computing resources in large-scale distributed systems to address these growing issues.

99. The '909 patent provides an inventive technological solution to communicate program configuration information, monitor the status of program execution, and dynamically allocate resources based on that monitoring information in large-scale distributed computing systems. (*Id.* at 2:22-49.)

100. The '909 patent describes the Distributed Program Execution Service System Manager (“DPESSM”) that provides a service for distributed program execution by “executing multiple programs on behalf of multiple customers.” (*Id.* at 2:48-63.) The DPESSM manages the computing nodes allocated to customers, allows customers to communicate program information

and execution requirements, and configures and initiates execution of the programs in the distributed system. (*Id.* at 8:5-9:5.) For example, claim 21 recites “one or more components of a distributed execution service that are configured to, when executed by at least one of the one or more processors, and for each of multiple programs to be executed: select a subset of a plurality of available computing nodes to use for executing the program in a distributed manner.” (*Id.* at claim 21.) The ability to receive specific program execution and configuration information related to computing resources for virtual machines on a distributed system was a concrete, technological improvement over the prior art. This type of personalized configuration allows users to specify the computing resources needed to execute a program, but also allows users to limit the computing resources consumed to limit costs. This benefit over the prior art was recognized in the industry. For example, by using the AWS cloud companies were able to efficiently use “many GPU cores, CPU cores, and AWS instances” to train “multiple models either from different datasets or configurations.” (Alex Chen, et al., Distributed Neural Networks with GPUs in the AWS Cloud, Medium (Feb. 10, 2014), <https://netflixtechblog.com/distributed-neural-networks-with-gpus-in-the-aws-cloud-ccf71e82056b>.)

101. The DPESM also solved the technological problem of programs failing to execute due to insufficient resources by monitoring all computing nodes, storing information related to physical resources of each computing node, and dynamically allocating additional computing nodes to allow the program to complete execution. ('909 patent at Figs. 2A-2C, 13:16-57, 14:19-46, 16:52-17:11, 18:25-38, 20:59-21:28.) For example, claim 21 recites “determine, while at least one of the subset of computing nodes is still executing at least one of the jobs of the program, that an actual amount of computing resources being used by the multiple computing nodes to execute the indicated program differs from an expected amount of computing resources.” (*Id.* at claim 21.)

The '909 patent explains that “[s]uch monitoring of the execution of execution jobs may provide various benefits, such as to determine when to later initiate execution of other execution jobs.” (*Id.* at 19:25-32.) This monitoring and tracking of performance of the distributed program execution also “provide[s] various benefits, such as to enable the ongoing intermediate execution and data state from the partial execution of the execution job to be tracked and used” in re-allocating node clusters to increase or decrease computing resources within the distributed system. (*Id.* at 4:10-49, 19:41-46.)

102. The claimed elements of the '909 patent also provide an inventive concept individually and as an ordered combination. The technological solutions provided by the '909 patent solve problems specific to large-scale distributed systems, which must effectively provision, administer, and manage distributed program execution, by dynamically modifying clusters of computing nodes. For example, claim 21 recites “based on the determining, initiate a change in a quantity of the computing nodes of the subset being used for the executing of the program.” (*Id.* at claim 21.) It was an improvement over the prior art to modify a quantity of computing nodes for use in further execution of a program by adding one or more additional computing nodes while the execution of the program is ongoing and furthermore, using the additional computing nodes for further execution of the program, and allocating computing resources of physical computing systems to the additional computing nodes. This technical advantage of dynamically adjusting computing resources based on configuration or demand was, and is, widely recognized as one of the main benefits of using the AWS Cloud over other solutions. The industry widely recognized this benefit in “us[ing] a reasonable number of machines to implement a powerful machine learning solution” through program customization and by “leverag[ing] the full, on-demand computing power we can obtain from AWS.” (Alex Chen, et al., Distributed Neural Networks

with GPUs in the AWS Cloud, Medium (Feb. 10, 2014), <https://netflixtechblog.com/distributed-neural-networks-with-gpus-in-the-aws-cloud-ccf71e82056b>.) This claim element, and the ordered combination of each of the elements of the claim, provide an inventive technological solution for the particular technological environment of computing nodes in a distributed computing, and solved critical shortcomings in the prior art.

J. The '211 patent

103. On February 2, 2016, the USPTO issued the '211 patent, titled “Managing communications between computing nodes.” The '211 patent issued from U.S. patent application number 13/843,287, filed on March 15, 2013. A true and correct copy of the '211 patent is attached as Exhibit J.

104. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '211 patent, including the right to assert all causes of action arising under the '211 patent and the right to sue and obtain remedies for past, present, and future infringement.

105. The '211 patent concerns “managing the execution of programs on multiple computing systems,” for example, by “managing communications between multiple intercommunicating computing nodes.” ('211 patent at Technical Field, Abstract.) The '211 patent improves virtualization technologies by allowing users to “manag[e] execution of programs on multiple computing systems.” (*Id.* at 2:58-59.) For example, a “program execution service [uses] a variety of factors to select an appropriate computing system to execute an instance of a program.” (*Id.* at 2:62-65.) These factors include “geographical and/or logical location, such as in one of multiple data centers that house multiple computing machines available for use.” (*Id.* at 4:56-67.) Such locations may include “one of multiple data centers that house multiple computing machines available for use, on multiple computing systems that are proximate to each other” or

“computing system[s] that are proximate to computing systems executing one or more other indicated program instances.” (*Id.*) The techniques disclosed in the ’211 patent provide benefits, for example, “enhanced reliability” in case of “network outages or other problems,” “failure of [a] computing system,” or “high network bandwidth for communications between instances of [an] executing program.” (*Id.* at 5:50-6:7.)

106. The ’211 patent describes techniques for “managing the execution of programs on multiple computing systems.” (*Id.* at 2:58-59.) The ’211 patent specification describes that although “[d]ata centers housing significant numbers of interconnected computing systems have become commonplace,” “the task of provisioning, administering, and managing the physical computing resources has become increasingly complicated.” (*Id.* at 1:18-19, 1:30-32.) While “virtualization technologies for commodity hardware” provided a “partial solution,” a problem still arose in how to “allow communications between [systems operated by or on behalf of each customer] (if desired by the customer) while restricting undesired communications to those systems from other systems.” (*Id.* at 1:33-35, 1:53-60.)

107. Although traditional firewalls could be used to address some of these kinds of problems in part, problems persisted. Traditional firewalls still permitted malicious actors to create resource outages. (*Id.* at 1:62-67.) Traditional firewalls lacked the ability to dynamically provision resources. (*Id.* at 1:67-2:4.) And traditional firewalls could not “dynamically determine appropriate filtering rules required to operate correctly.” (*Id.* at 2:6-8.)

108. The ’211 patent provides a technological solution to these problems that firewalls could not solve. The ’211 patent’s techniques “allow users to efficiently specify communications policies that are automatically enforced via management of data transmissions for multiple

computing nodes, such as for multiple hosted virtual machines operating in one or more data centers or other computing resource facilities.” (*Id.* at 2:10-15.)

109. The ’211 patent describes techniques for “executing multiple programs on behalf of multiple users.” (*Id.* at 2:58-62.) It recites a “variety of factors” for provisioning, administering, and managing the appropriate computing systems on which to execute programs. (*Id.* at 2:63-3:2.) The ’211 patent identifies factors such as “the location of one or more previously stored copies of [a] program” and which computing systems are “proximate” geographically or logically. (*Id.* at 2:62-3:13.) The ’211 patent claims recite multiple criteria for use in determining which computing nodes to use for execution of the virtual machine: the nodes’ geographic location; the number of nodes required; the minimum or maximum number of instances to execute; initiation and termination times; user-specified resource criteria; amount of memory, processor usage, network bandwidth, disk space or swap space required; and minimum or maximum resources to be used.

110. The ’211 patent techniques provide the technical advantage of enhanced reliability. For example, distributing multiple instances that are to be run at the same time, among computing programs “that are members of different groups” provides “enhanced reliability in the face of group-specific network outages or other problems.” (*Id.* at 5:50-59.) In another example, executing multiple instances on multiple computing systems provides “enhanced reliability in the face of failure” or “loss of connection” to a single system. (*Id.* at 5:63-67.)

111. The ’211 patent techniques provide the technical advantage of increased network bandwidth for transmission of data. The service may execute multiple instances in a single data center to provide “relatively high network bandwidth for communications between instances.” (*Id.* at 6:2-7.)

112. The claimed elements of the '211 patent also provide an inventive concept individually and as an ordered combination. Early data centers with fewer computing resources were simple enough to function without the advanced scalable management technology afforded by the '211 patent. As the scale and scope of data centers increased, the need arose for new methods to manage a vast fleet of computing resources. (*Id.* at 1:18-32.) The elements in the claims of the '211 patent established multiple new techniques, uniquely advantageous to the newly-expanded scale of data centers, that permitted users to coordinate program execution configurations remotely and efficiently. (*Id.* at 1:6-14.) A user specifies configuration information to execute a specific program through a remote interface providing access to a program execution service. (*See id.* at claim 1.) After receiving the configuration information, the system does not merely execute a program using that information as was conventionally done. Instead, the system performs further actions to enhance the effectiveness of the user's configuration input without requiring the user to perform additional custom configurations to accommodate the increased scale of computational resources available. The actions of the system reduce the possibility of errors that could be introduced by forcing a user to engage in multiple configuration requests. (*See id.* at 1:67-2:9.) The first action selects computing nodes from among the available resources to use in executing potentially multiple instances of the specific program, mitigating the need for the user to specify specific system resources for program execution. (*See id.* at claim 1.) The second action manages the execution of potentially multiple instances of the specific program, applying at least the configuration information received from the user via the specialized configuration interface, removing the need for the user to provide resource-specific configurations. (*See id.* at 21:60-63.) In taking these actions the system considers, for example, communications, resource, and permissions capabilities of multiple pieces of physical hardware

without the need for additional configuration steps to be taken by the user. (*See id.* at 3:14-30.) For example, claim 23 provides the technical solution of using geographic location to determine choice of computing nodes, describing a service “having a plurality of computing nodes located in multiple geographic locations.” (*Id.* at claim 23.) The service uses configuration information that “indicates one or more geographical locations in which at least one indicated program is to be executed” to select multiple computing nodes to use for execution of the indicated program. (*Id.*) These actions and the configuration interface of the ordered combination together provided an advancement over the conventional technology at the time resulting in multiple benefits based on a user’s configuration specification.

K. The ’593 patent

113. On April 11, 2017, the USPTO issued the ’593 patent, titled “Managing execution of programs by multiple computing systems.” The ’593 patent issued from U.S. patent application number 14/928,659, filed on October 30, 2015. A true and correct copy of the ’593 patent is attached as Exhibit K.

114. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the ’593 patent, including the right to assert all causes of action arising under the ’593 patent and the right to sue and obtain remedies for past, present, and future infringement.

115. The ’593 patent concerns “managing the execution of programs on multiple computing systems,” for example, by “select[ing] appropriate computing systems to execute one or more instances of a program for a user, such as based in part on configuration information specified by the user.” (’593 patent at Abstract.)

116. The ’593 patent improves virtualization technologies by “managing the execution of programs on multiple computing systems, such as on virtual machine nodes executing on the

computing systems,” for example, to “select appropriate computing systems to execute ... instances of a program for a user, such as based ... on configuration information specified by the user.” (*Id.* at Abstract.) For example, techniques disclosed in the ’593 patent allow “data transmission management systems executing as part of the application execution service” to “control communications to and from the applications.” (*Id.* at 3:45-53.) As another example, the “new groups of computing nodes” may be created and “access policies for the groups” specified and the system is able to handle “changing conditions” such as “new application instances that are executed, previously executing application instances that are no longer executing, and/or new or adjusted access policies.” (*Id.* at 3:53-67.)

117. The ’593 patent describes techniques for “managing communications between multiple intercommunicating computing nodes.” (*Id.* at 2:66-67.) The ’593 patent specification describes that although “data centers housing significant numbers of interconnected computing systems have become commonplace,” “the task of provisioning, administering, and managing the physical computing resources has become increasingly complicated” due to the increase in “scale and scope” of data centers. (*Id.* at 1:22-23, 1:33-36.) While “virtualization technologies for commodity hardware” provided a “partial solution,” a problem still arose in how to “allow communications between [systems operated by or on behalf of each customer] (if desired by the customer) while restricting undesired communications to those systems from other systems.” (*Id.* at 1:37-43, 1:58-65.)

118. Although traditional firewalls could be used to address some of these kinds of problems in part, problems persisted. Traditional firewalls still permitted malicious actors to create resource outages. (*Id.* at 1:67-2:5.) Traditional firewalls lacked the ability to dynamically

provision resources. (*Id.* at 2:5-9.) And traditional firewalls could not “dynamically determine appropriate filtering rules required to operate correctly.” (*Id.* at 2:10-15.)

119. The ’593 patent provides a technological solution to these problems that firewalls could not solve. The ’593 patent’s techniques “allow users to efficiently specify communications policies that are automatically enforced via management of data transmissions for multiple computing nodes, such as for multiple hosted virtual machines operating in one or more data centers or other computing resource facilities.” (*Id.* at 2:17-22.)

120. The ’593 patent describes techniques for managing communications through “virtual machine nodes” hosted on machines or systems. (*Id.* at 3:1-3.) The ’593 patent also describes managing communications through “intermediary computing nodes” which are also virtual machines. (*Id.* at 3:6-12.) Management techniques include “analyzing outgoing data transmissions” to determine their authorizations, which may be based on “defined data transmission policies” that specify what groups of destination nodes are authorized. (*Id.* at 3:18-33.) Further, the ’593 patent describes the technique of using negotiations between intermediary nodes to determine authorization and storing the outcomes of those negotiations as rules. (*Id.* at 3:33-44.)

121. The ’593 patent also recites a “variety of factors” for selecting an appropriate computing system to execute an instance of a program, including “the location of one or more previously stored copies of [a] program” and which computing systems are “proximate” geographically or logically. (*Id.* at 20:35-46, 20:50-56.)

122. Many benefits result from the ’593 patent’s techniques. Utilizing the location of a previously-stored program results in “reduc[ing] the program execution startup latency.” (*Id.* at 23:16-22.) Using negotiated rules to authorize reply data transmissions “enable[s] some types of

transmission protocols (e.g., TCP) to function effectively.” (*Id.* at 13:61-64.) Distributing multiple instances that are to be run at the same time, among computing programs “that are members of different groups” provides “enhanced reliability in the fact of group-specific network outages or other problems.” (*Id.* at 23:33-40.) In another example, executing multiple instances on multiple computing systems provides “enhanced reliability in the face of failure” or “loss of connection” to a single system. (*Id.* at 23:44-48.)

123. The ’593 patent techniques also provide the benefit of increased network bandwidth for transmission of data. For example, the service may execute multiple instances in a single data center to provide “relatively high network bandwidth for communications between instances.” (*Id.* at 23:49-56.)

124. The claimed elements of the ’593 patent also provide an inventive concept individually and as an ordered combination. Early data centers with fewer computing resources were simple enough to function without the advanced scalable management technology afforded by the ’593 patent. As the scale and scope of data centers increased, the need arose for new methods to manage a vast fleet of computing resources. (*Id.* at 1:22-36.) The elements in the claims of the ’593 patent established multiple new techniques, uniquely advantageous to the newly-expanded scale of data centers, that permitted users to coordinate program execution configurations remotely and efficiently. (*Id.* at 2:15-21.) A user could specify configuration information to execute a specific program through a remote interface providing access to a program execution service. (*See id.* at 21:40-59.) After receiving the configuration information, the system would not merely execute a program using that information as was conventionally done. Instead, the system would perform further actions to enhance the effectiveness of the user’s configuration input without requiring the user to perform additional custom configurations to accommodate the

increased scale of computational resources available. The actions of the system reduced the possibility of errors that could be introduced by forcing a user to engage in multiple configuration requests. (*See id.* at 2:9-14.) The first action would select computing nodes from among the available resources to use in executing potentially multiple instances of the specific program, mitigating the need for the user to specify specific system resources for program execution. (*See id.* at 20:34-56, claim 23.) The second action would manage the execution of potentially multiple instances of the specific program, applying at least the configuration information received from the user via the specialized configuration interface, removing the need for the user to provide resource-specific configurations. (*See id.* at 21:60-63, claim 23.) In taking these actions the system could consider, for example, communications, resource, and permissions capabilities of multiple pieces of physical hardware without the need for additional configuration steps to be taken by the user. (*See id.* at 20:57-21:10, claim 23.) These actions and the configuration interface of the ordered combination together provided an advancement over the conventional technology at the time resulting in multiple benefits based on a user's configuration specification.

L. The '419 patent

125. On October 23, 2012, the USPTO issued the '419 patent, titled "Dynamically modifying a cluster of computing nodes used for distributed execution of a program." The '419 patent issued from U.S. patent application number 12/415,725, filed on March 31, 2009. A true and correct copy of the '419 patent is attached as Exhibit L.

126. Amazon Technologies, Inc. is the owner and assignee of all right, title, and interest in and to the '419 patent, including the right to assert all causes of action arising under the '419 patent and the right to sue and obtain remedies for past, present, and future infringement.

127. The '419 patent concerns “distributed execution of programs.” ('419 patent at Abstract.)

128. The '419 patent improves virtualization technologies by allowing users to “manag[e] the distributed execution of programs” by “dynamically monitoring the ongoing distributed execution of a program on a cluster of multiple computing nodes, and dynamically modifying the distributed program execution.” (*Id.* at 2:15-27.) The “dynamic modifying of the distributed program execution” provides several advantages, for example “[c]luster expansion ... to enable program execution to complete sooner” or “[c]luster shrinking ... to more efficiently use resources,” and so on. (*Id.* at Abstract, 4:5-35.)

129. Prior to the '419 patent, there had been rapid growth in distributed computing and that “data centers housing significant numbers of interconnected computing systems ha[d] become commonplace.” (*Id.* at 1:20-22.) The '419 patent explains that “as the scale and scope of typical data centers has increased, the task of provisioning, administering, and managing the physical computing resources has become increasingly complicated.” (*Id.* at 1:31-35.) Without proper management of the physical computing resources, programs could use more resources than they should be allocated, becoming bottlenecks for the entire system. (*Id.* at 2:27-39, 30:39-46.) In other scenarios, programs could fail to execute if “insufficient quantity of computing nodes of the cluster are available to perform execution” or complete execution in the allocated timeframe. (*Id.* at 2:37-39, 33:62-34:3.) These problems were compounded by virtualization technologies that allowed “a single physical computing machine to be shared among multiple users” or “a single virtual machine with multiple virtual processors that actually spans multiple distinct physical computing systems.” (*Id.* at 1:42-44, 1:53-55.) There was no well-understood, routine, and

conventional way to effectively provision, administer, and manage the physical computing resources in large-scale distributed systems to address these growing issues.

130. The '419 patent provides an inventive technological solution to communicate program configuration information, monitor the status of program execution, and dynamically allocate resources based on that monitoring information in large-scale distributed computing systems. (*Id.* at 2:15-42.)

131. The '419 patent describes the Distributed Program Execution Service System Manager (“DPESSM”) that provides a service for distributed program execution by “executing multiple programs on behalf of multiple customers.” (*Id.* at 2:43-57.) The DPESSM manages the computing nodes allocated to customers, allows customers to communicate program information and execution requirements, and configures and initiates execution of the programs in the distributed system. (*Id.* at 7:66-8:28, 8:29-67.) For example, claim 4 recites “receiving, by one or more computing systems configured to provide a distributed program execution service having a plurality of computing nodes, configuration information regarding executing an indicated program on an indicated quantity of multiple of the plurality of computing nodes, wherein the executing of the indicated program causes a plurality of jobs to be executed.” (*Id.* at claim 4.) The ability to receive specific program execution and configuration information related to computing resources for virtual machines on a distributed system was a concrete, technological improvement over the prior art. This type of personalized configuration allows users to specify the computing resources needed to execute a program, but also allows users to limit the computing resources consumed to limit costs. This technical advantage was recognized in the industry. For example, by using the AWS cloud companies were able to efficiently “many GPU cores, CPU cores, and AWS instances” to train “multiple models either from different datasets or configurations.” (Alex Chen, et al.,

Distributed Neural Networks with GPUs in the AWS Cloud, Medium (Feb. 10, 2014), <https://netflixtechblog.com/distributed-neural-networks-with-gpus-in-the-aws-cloud-ccf71e82056b>.)

132. The DPESM also solved the technological problem of programs failing to execute due to insufficient resources by monitoring all computing nodes, storing information related to physical resources of each computing node, and dynamically allocating additional computing nodes to allow the program to complete execution. ('419 patent at Figs. 2A-2C, 13:13-54, 14:16-43, 16:49-17:8, 18:22-35, 20:56-21:25.) For example, claim 4 recites “determining, by the one or more configured computing systems at a second time subsequent to the first time, whether a minimum subset of the multiple computing nodes have begun to execute the jobs of the indicated program as expected.” (*Id.* at claim 4.) The '419 patent explains that “[s]uch monitoring of the execution of execution jobs may provide various benefits, such as to determine when to later initiate execution of other execution jobs.” (*Id.* at 19:22-29.) This monitoring and tracking of performance of the distributed program execution also “provide[s] various benefits, such as to enable the ongoing intermediate execution and data state from the partial execution of the execution job to be tracked and used” in re-allocating node clusters to increase computing resources within the distributed system. (*Id.* at 4:5-43, 19:38-43.)

133. The claimed elements of the '419 patent also provide an inventive concept individually and as an ordered combination. The technological solutions provided by the '419 patent solve problems specific to large-scale distributed systems that effectively provision, administer, and manage distributed program execution by dynamically modifying clusters of computing nodes. For example, claim 4 recites “in response to the determining, initiating a change in a quantity of the multiple computing nodes that are used for executing the jobs of the indicated

program.” (*Id.* at claim 4.) It was an improvement over the prior art to modify a quantity of computing nodes for use in further execution of a program by adding one or more additional computing nodes while the execution of the program is ongoing and furthermore, using the additional computing nodes for further execution of the program, and allocating computing resources of physical computing systems to the additional computing nodes. This technical advantage of dynamically adjusting computing resources based on configuration or demand was, and is, widely recognized as one of the main benefits of using the AWS Cloud over other solutions. The industry widely recognized this benefit as “us[ing] a reasonable number of machines to implement a powerful machine learning solution” through program customization and by “leverag[ing] the full, on-demand computing power we can obtain from AWS.” (Alex Chen, et al., Distributed Neural Networks with GPUs in the AWS Cloud, Medium (Feb. 10, 2014), <https://netflixtechblog.com/distributed-neural-networks-with-gpus-in-the-aws-cloud-ccf71e82056b>.) This claim element, and the ordered combination of each of the elements of the claim, provide a concrete, inventive technological solution for the particular technological environment of computing nodes in a distributed computing, that solved critical shortcomings in the prior art.

COUNT I: PATENT INFRINGEMENT OF U.S. PATENT NO. 11,516,080

134. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

135. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 1) of the '080 patent in violation of 35 U.S.C. § 271, and will continue to do so.

136. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to the Nuage Networks Virtualized Services Platform, in violation of 35 U.S.C. § 271(a).

137. By way of example only, the Nuage Network Virtualized Services Platform meets all the limitations of at least independent claim 1 of the '080 patent, either literally or under the doctrine of equivalents.

138. Exemplary claim 1 of the '080 patent recites:

1. A method, comprising:

performing, by one or more computing systems of a configurable network service:

receiving first configuration information for a first virtual computer network of computing nodes to be provided for a first client, wherein the first configuration information indicates a first range of network addresses to be assigned to the computing nodes of the first virtual computer network;

providing the first virtual computer network to the first client according to the first configuration information, wherein the first virtual computer network is overlaid on a substrate network of the configurable network service;

assigning one of the network addresses in the first range to one of the computing nodes in the first virtual computer network;

receiving second configuration information for a second virtual computer network of computing nodes to be provided for a second client, wherein the second configuration information indicates a second range of network addresses to be assigned to the computing nodes of the second virtual computer network;

providing the second virtual computer network to the second client according to the second configuration information,

wherein the second virtual computer network is overlaid on the substrate network of the configurable network service;

assigning one of the network addresses in the first range to one of the computing nodes in the first virtual computer network;

providing a virtual peering router configured to manage an interconnection between the first virtual computer network and the second virtual computer network; and

routing, using the virtual peering router, network traffic from the computing nodes of the first virtual computer network to the computing nodes of the second virtual computer network over the substrate network.

139. For the preamble of claim 1, to the extent the preamble is determined to be limiting, the Nuage Network Virtualized Services Platform practices a “method, comprising: performing, by one or more computing systems of a configurable network service.”

140. For example, the Nuage Networks Virtualized Services Platform enables “Virtualized Cloud Services (VCS) [as] the data center and cloud networking framework” that allows “customers to automate the configuration, management and optimization of virtual networks, as seen below.”²⁹

²⁹Nuage Networks from Nokia, Nuage Networks Virtualized Services Platform, <https://www.nuagenetworks.net/platform/virtualized-services-platform/>.

Virtualized Cloud Services

Virtualized Cloud Services (VCS) is the data center and cloud networking framework enabled by Nuage Networks Virtualized Services Platform (VSP). It enables CSPs and their Telco Cloud deployments as well as large enterprise data center customers to automate the configuration, management and optimization of virtual networks, including security services that provide tenant isolation and access controls to individual applications and workloads.

VCS is a non-disruptive software-defined networking (SDN) overlay for all virtualized and non-virtualized server and network resources. It is transparent to the underlying physical infrastructure and provides a complete cloud networking framework that requires no specialized hardware. VCS can be deployed in any mixed environment including Docker containers, multi-hypervisor Virtual Machines (VMs), or bare metal servers.

[Learn More](#)

141. Claim 1 of the '080 patent further recites “receiving first configuration information for a first virtual computer network of computing nodes to be provided for a first client, wherein the first configuration information indicates a first range of network addresses to be assigned to the computing nodes of the first virtual computer network.”

142. The Nuage Networks Virtualized Services Platform practices this limitation. For example, the Nuage Networks Virtualized Services Platform (VSP) implements a VSP solution that includes a Virtualized Services Directory (VSD) component, as seen below.³⁰

³⁰Nuage Networks from Nokia, TECHNICAL DESCRIPTION, Nuage Networks Virtualized Services Platform: Service Chaining (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Service_Chaining_Technical_Description_Document_EN-compressed.pdf, p. 1.

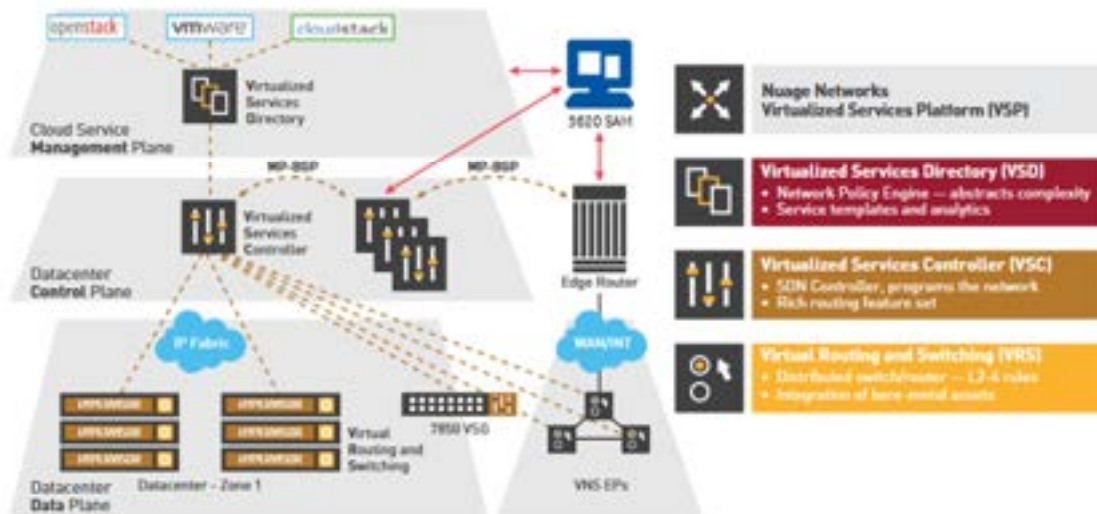
VSP solution overview

The Nuage Networks Virtualized Services Platform (VSP) is an industry-leading SDN solution that leverages Network Virtualization Overlay (NVO) technologies to make the network as readily consumable as the compute resources in a cloud environment. The VSP achieves this by ensuring rapid and efficient delivery of highly customizable application services in and across multitenant datacenters. The VSP enables the deployment of massively scalable cloud-based services over an existing IP network fabric with the agility and performance demanded by highly dynamic application environments.

The main components of the VSP solution are:

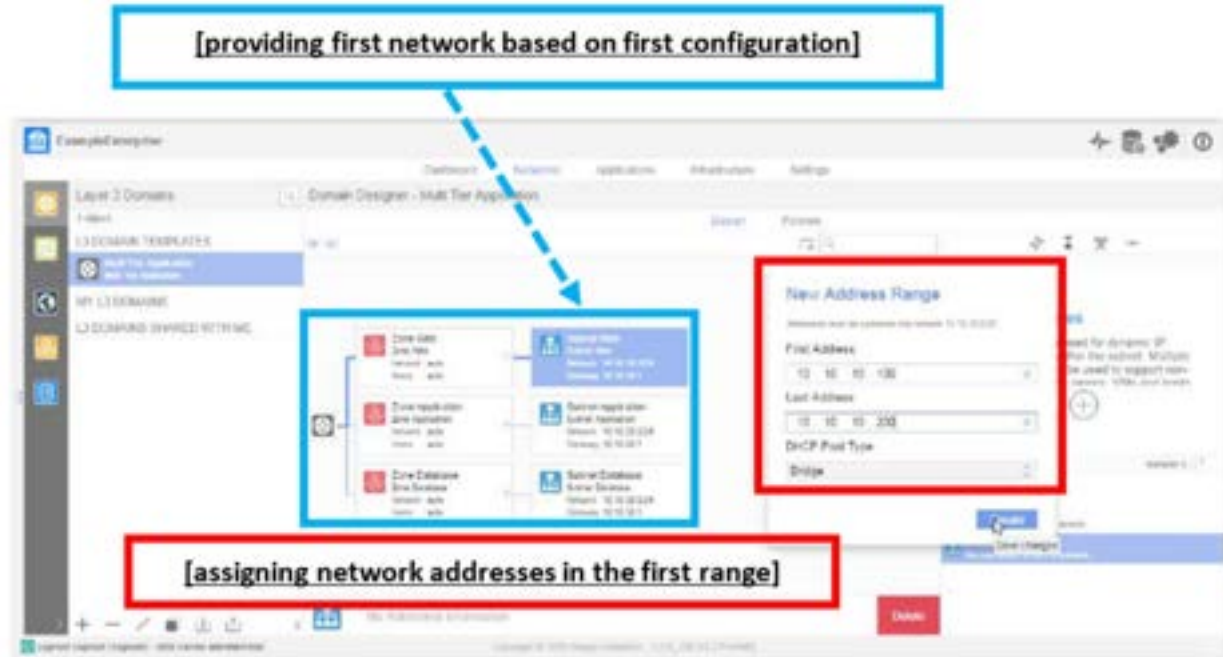
- Virtualized Services Directory (VSD)
- Virtualized Services Controller (VSC)
- Virtual Routing and Switching (VRS)

FIGURE 1. Nuage Networks VSP solution components



The VSD component allows configuration of a virtual computer network using a VSD Architect tool that receives a (first) range of network addresses as input.³¹

³¹*Id.* at pp.15-21.

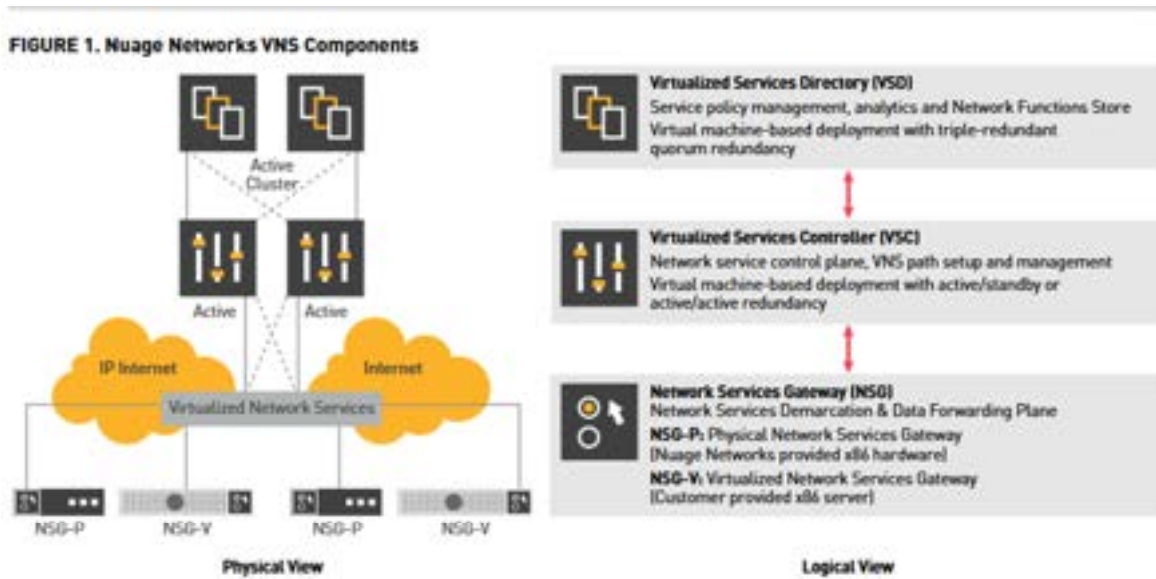


(Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2 (notation added).)

143. Claim 1 of the '080 patent further recites “providing the first virtual computer network to the first client according to the first configuration information, wherein the first virtual computer network is overlaid on a substrate network of the configurable network service” and “assigning one of the network addresses in the first range to one of the computing nodes in the first virtual computer network.”

144. The Nuage Networks Virtualized Services Platform practices this limitation. For example, the Nuage Networks Virtualized Services Platform implements Nuage Networks Virtualized Network Services (VNS) that are “based on an overlay model that uses any IP network to provide underlay connectivity between sites. This gives you maximum flexibility for your

locations and the support of multiple access/last-mile technologies including copper, fiber or mobile broadband,” as seen below.³²



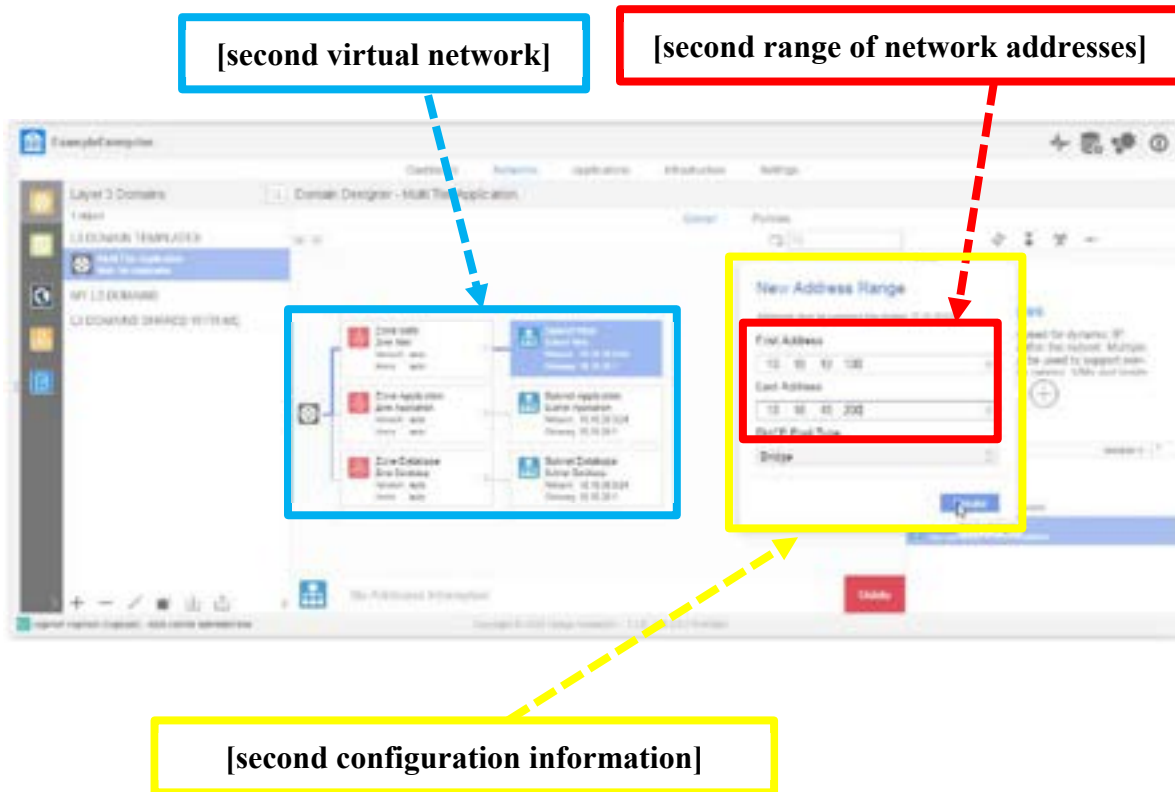
Traditional VPN services are tightly connected to dedicated network infrastructure from a service provider. Nuage Networks VNS are based on an overlay model that uses any IP network to provide underlay connectivity between sites. This gives you maximum flexibility for your locations and the support of multiple access/last-mile technologies including copper, fiber or mobile broadband.

145. Claim 1 of the '080 patent further recites “receiving second configuration information for a second virtual computer network of computing nodes to be provided for a second client, wherein the second configuration information indicates a second range of network addresses to be assigned to the computing nodes of the second virtual computer network.”

146. The Nuage Networks Virtualized Services Platform practices this limitation. For example, the VSD Architect of the Nuage Networks Virtualized Services Platform receives a range of network addresses from a user interface for configuring the second virtual network. The range of network addresses is included in the second configuration information that is received by the

³²Nuage Networks from Nokia, FLEXIBLE NETWORK SERVICES TO DRIVE YOUR ENTERPRISE AT CLOUD SPEED: Solution Primer (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Drive_your_enterprise_with_VNS_Solution_Sheet_EN.pdf, p. 6.

Nuage Networks Virtualized Services Platform for assigning network addresses to computing nodes of the second virtual computer network.³³

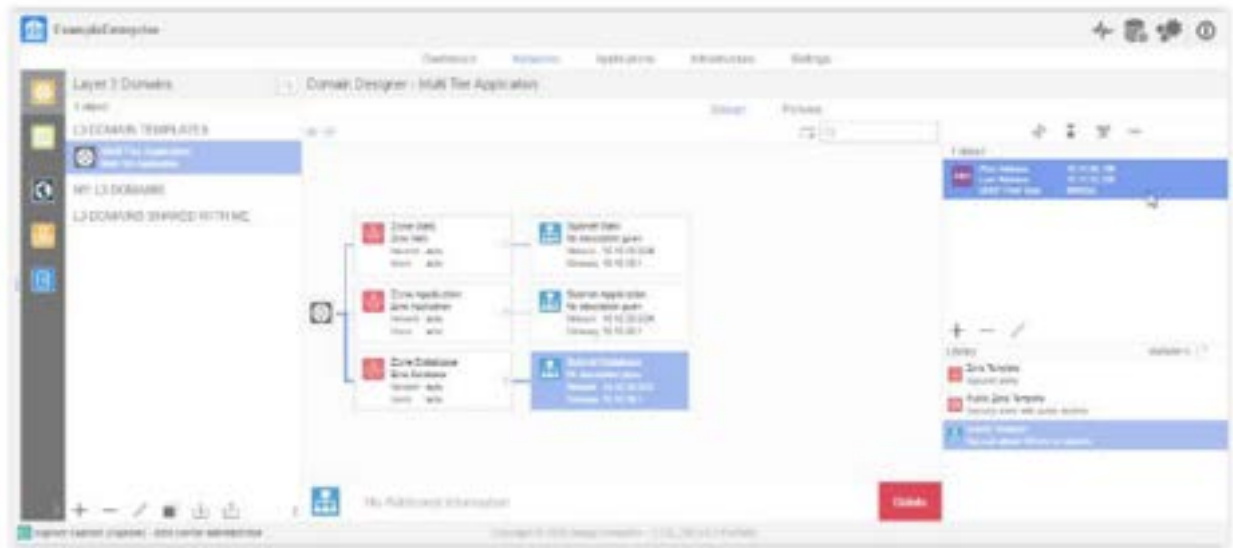


(Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2 (notation added).)

According to Nuage, the procedure may then be repeated for all subnets:

³³Nuage Networks from Nokia, TECHNICAL DESCRIPTION, Nuage Networks Virtualized Services Platform: Service Chaining (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Service_Chaining_Technical_Description_Document_EN-compressed.pdf, pp. 15-21.

15. Repeat the procedure for all subnets.

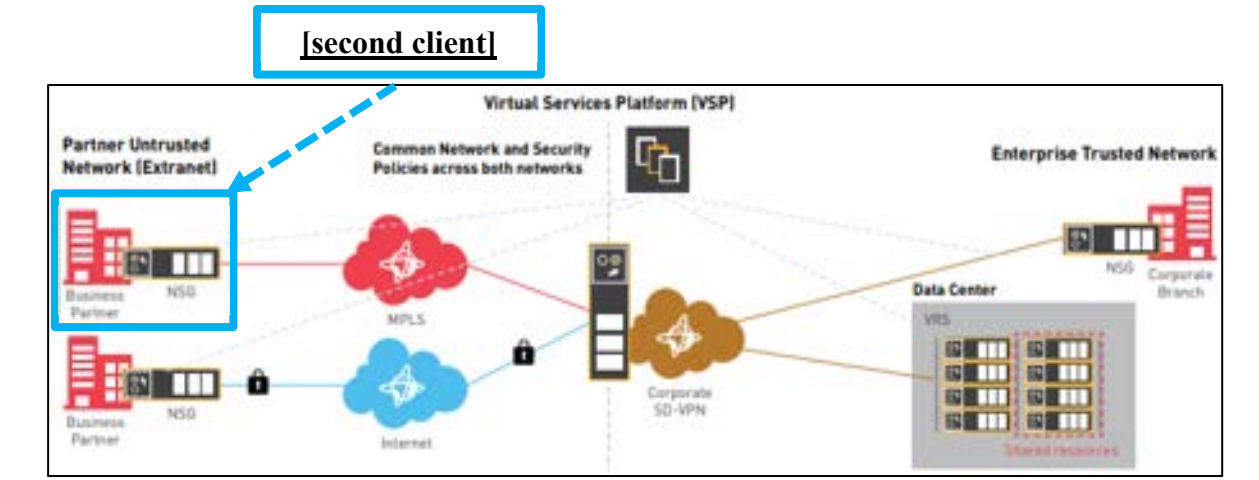


(Nuage Networks from Nokia, TECHNICAL DESCRIPTION, Nuage Networks Virtualized Services Platform: Service Chaining (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Service_Chaining_Technical_Description_Document_EN-compressed.pdf, p. 20, step 15.)

147. Claim 1 of the '080 patent further recites “providing the second virtual computer network to the second client according to the second configuration information, wherein the second virtual computer network is overlaid on the substrate network of the configurable network service” and “assigning one of the network addresses in the first range to one of the computing nodes in the first virtual computer network.”

148. The Nuage Networks Virtualized Services Platform practices this limitation. For example, the Nuage Networks Virtualized Services Platform provides the second configured virtual computer network to a (second) client.³⁴

³⁴Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2.



(Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuage networks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2 (notation added).)

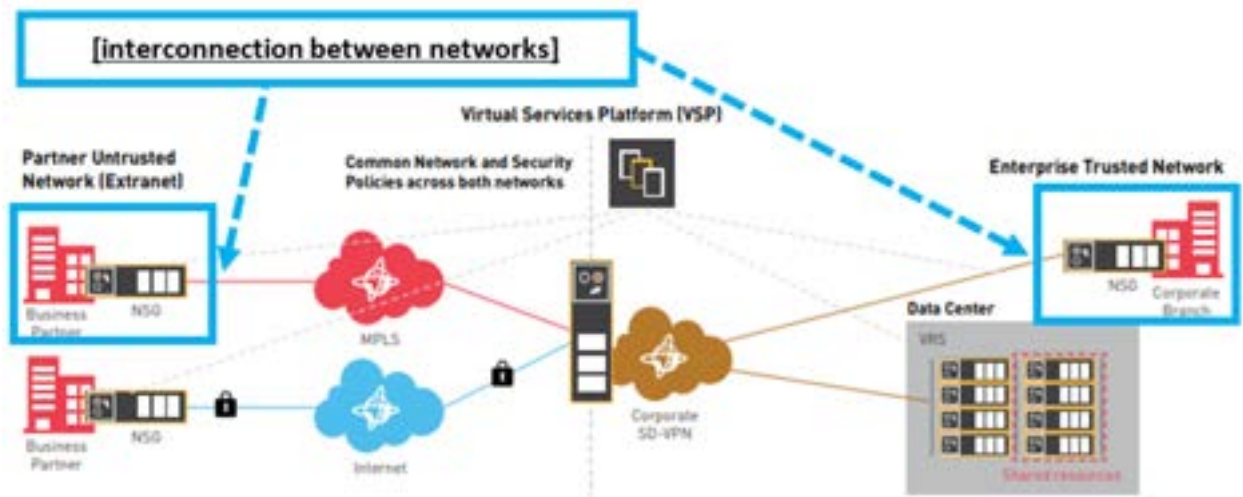
149. The Nuage Networks Virtualized Services Platform practices “wherein the second virtual computer network is overlaid on the substrate network of the configurable network service.” For example, the Nuage Networks Virtualized Services Platform implements Nuage Networks Virtualized Network Services (VNS) that are “based on an overlay model that uses any IP network to provide underlay connectivity between sites.”³⁵ Any IP network may be external or internal, as shown in the figure above. The platform “unifies the management and provisioning of virtual networks regardless of the underlying network technology, allowing seamless integration of IP/MPLS and broadband internet-connected sites and partners.”³⁶ According to Nuage Networks, “[t]his technology is now providing the same benefits to multi-party private Extranets.

³⁵Nuage Networks from Nokia, FLEXIBLE NETWORK SERVICES TO DRIVE YOUR ENTERPRISE AT CLOUD SPEED: Solution Primer (2016), https://www.nuage networks.net/wp-content/uploads/2020/06/Nuage_Networks_Drive_your_enterprise_with_VNS_Solution_Sheet_EN.pdf, p. 6.

³⁶Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuage networks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2.

Nuage Networks Virtualized Network Services (VNS) SD-WAN solution provides policy-based control to simplify the management and automation of Extranets, while maintaining security controls and accelerating the time to integrate new Extranet partners and sites.”³⁷

150. The Virtualized Services Platform provides an interconnection between the first and second configured virtual computer networks, as shown below, in an interconnection or “network sequencing of service functions” called “service chaining.”³⁸



(Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2 (notation added).)

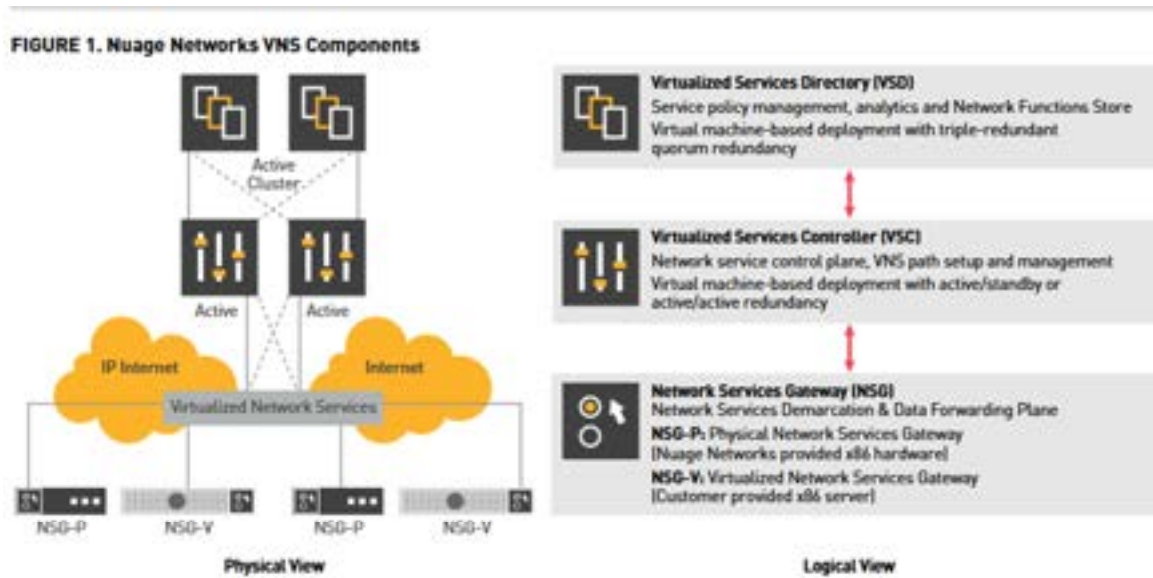
The Nuage VSP “supports service chaining in both virtual and physical environments within a datacenter, as well as Layer 3 and Layer 2 SDN networking.”³⁹

³⁷*Id.*

³⁸Nuage Networks from Nokia, TECHNICAL DESCRIPTION, Nuage Networks Virtualized Services Platform: Service Chaining (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Service_Chaining_Technical_Description_Document_EN-compressed.pdf, p. 9.

³⁹*Id.*

151. The “Virtualized Services Controller” of the “Nuage Networks Virtualized Network Services” implements a virtual computer network using an “overlay model” based on “network overlay paths to form the topology for the network service.”⁴⁰ Furthermore, Nuage Networks provides a “solution [that] correlates SDN-based overlay services to the underlying physical network, giving network operators better visibility.”⁴¹ The result is a single virtual network, with VNS providing a “single policy for both internet and Extranet domains,” with “security routing control” and “centrally managed policies.”⁴²



Traditional VPN services are tightly connected to dedicated network infrastructure from a service provider. Nuage Networks VNS are based on an overlay model that uses any IP network to provide underlay connectivity between sites. This gives you maximum flexibility for your locations and the support of multiple access/last-mile technologies including copper, fiber or mobile broadband.

⁴⁰*Id.* at p. 3.

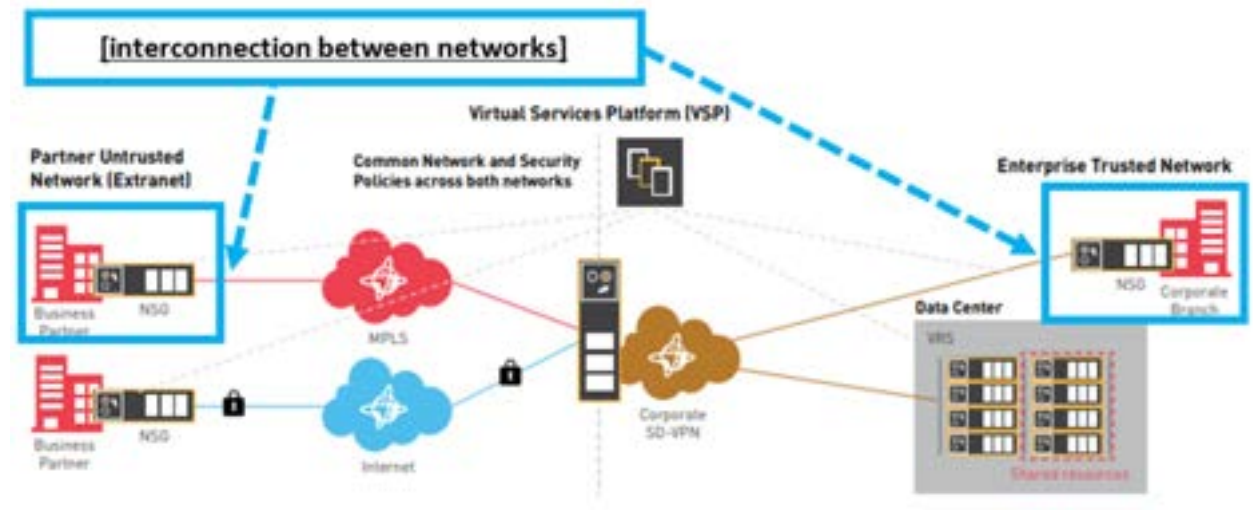
⁴¹Nuage Networks from Nokia, Technology White Paper: Correlating SDN Overlays and the Physical Network with Nuage Networks Virtualized Services Assurance Platform (2016), https://www.nuage networks.net/wp-content/uploads/2020/06/Nuage_Networks_Correlating_SDN_Overlays_and_VSAP_White_Paper_EN-compressed.pdf, p. 4.

⁴²Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuage networks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 1.

(Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuage networks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2.)

152. Claim 1 of the '080 patent further recites “providing a virtual peering router configured to manage an interconnection between the first virtual computer network and the second virtual computer network.”

153. The Nuage Networks Virtualized Services Platform practices this limitation. For example, the Nuage Networks Virtualized Services Platform includes a “Nuage Networks NSG-BR (border router) [that] extends seamless connectivity between disparate networks. It connects untrusted domains (partner Extranet networks) to the core enterprise trusted network and provides the gateway functionality for partner Extranet connectivity.”⁴³



(*Id.* at p. 2 (notation added).)

The NSG Border Router is “[a] key complementary component that Nuage Networks has developed for SD-WAN deployments in release 4.0 NSG is Nuage Networks’ SD-WAN branch router or customer premises equipment (CPE) based on an open x86 architecture. NSG

⁴³*Id.* at pp. 1, 2.

Border Router is a software function running on CPE device that typically acts as an overlay network gateway between the on-premises datacenter/cloud network and the WAN network.”⁴⁴

The NSG Border Router permits multi-tenancy by “enabl[ing] connectivity between multiple partner domains with a single appliance.”⁴⁵

The Nuage Networks VSP thus provides a seamless layer to users: “Organizations want the ability to specify Virtual Network Functions (VNFs) or Physical Network Functions (PNFs) and their sequence, so service functions can be added or removed seamlessly without requiring changes to the underlying network infrastructure.” Its service chaining function supports a single virtual environment: “The Nuage Networks VSP supports service chaining in both virtual and physical environments within a datacenter, as well as Layer 3 and Layer 2 SDN networking.”⁴⁶

The VSP is touted as “providing the same benefits to multi-party private Extranets. Nuage Networks Virtualized Network Services (VNS) SD-WAN solution provides policy-based control to simplify the management and automation of Extranets, while maintaining security controls and accelerating the time to integrate new Extranet partners and sites.”⁴⁷

⁴⁴Saurabh Sandhir, Nuage Networks Evolves SDN/SD-WAN Platform to Broader Use Cases, Diverse Cloud Environments in Release 4.0 – Part 1 (2016), <https://www.nuagenetworks.net/blog/rel4-1/>.

⁴⁵Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2.

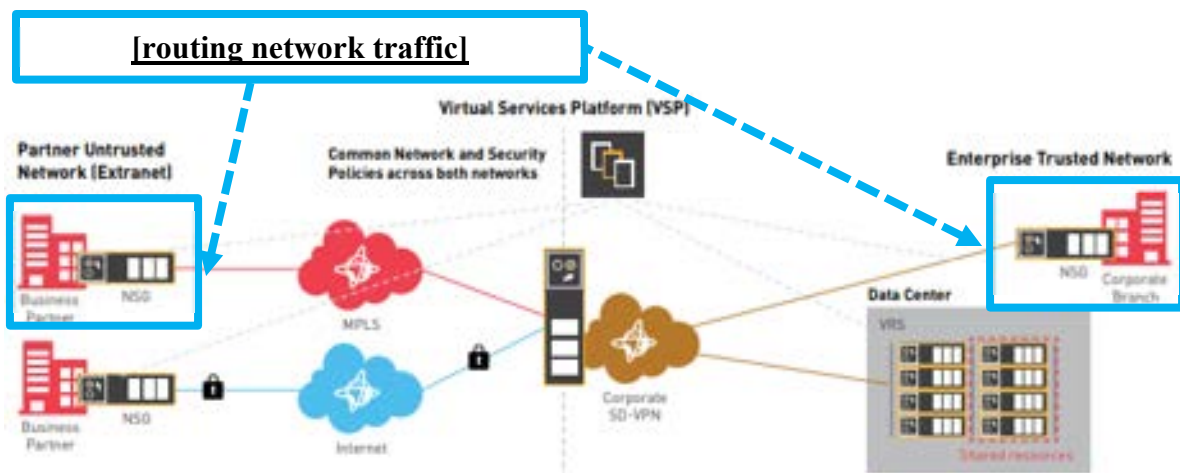
⁴⁶Nuage Networks from Nokia, TECHNICAL DESCRIPTION, Nuage Networks Virtualized Services Platform: Service Chaining (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Service_Chaining_Technical_Description_Document_EN-compressed.pdf, p. 9.

⁴⁷Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2.

The VSP also governs common network and security policies for other parties to access shared resources in the data center of the Enterprise Trusted Network.⁴⁸

154. Claim 1 of the '080 patent further recites “routing, using the virtual peering router, network traffic from the computing nodes of the first virtual computer network to the computing nodes of the second virtual computer network over the substrate network.”

155. The Nuage Networks Virtualized Services Platform practices this limitation. For example, the Nuage Networks VNS includes a “Nuage Networks NSG-BR (border router) functionality [that] extends seamless connectivity between disparate networks. It connects untrusted domains (partner Extranet networks) to the core enterprise trusted network and provides the gateway functionality for partner Extranet connectivity.”⁴⁹



(Nuage Networks from Nokia, Nuage Networks Enterprise Extranet Solution Brief (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Extranet_Solution_Brief_Document.pdf, p. 2 (notation added).)

156. Nokia is and has been on notice of the infringement of the '080 patent at least as of the time Amazon filed and provided notice of this Complaint.

⁴⁸*Id.*

⁴⁹*Id.* at p. 2.

157. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '080 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, the Nuage Networks Virtualized Services Platform.

158. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 1 of the '080 patent.

159. Nokia will sell the Nuage Networks Virtualized Services Platform with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '080 patent.

160. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts will cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

161. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '080 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of the Nuage Networks Virtualized Services Platform, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '080 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT II: PATENT INFRINGEMENT OF U.S. PATENT NO. 11,425,194

162. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

163. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 1) of the '194 patent in violation of 35 U.S.C. § 271, and will continue to do so.

164. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia CloudBand Application Manager, in violation of 35 U.S.C. § 271(a).

165. By way of example only, Nokia CloudBand Application Manager meets all the limitations of at least independent claim 1 of the '194 patent, either literally or under the doctrine of equivalents.

166. Exemplary claim 1 of the '194 patent recites:

1. A computer-implemented method, comprising:

receiving, by one or more configured computing systems of a program execution service that provides computing resources available to multiple users of the program execution service, instructions from a first user to execute a program using specified configuration information, the specified configuration information including a first number of virtual machines (VMs) to use to execute the program and a set of instructions specifying how to modify a number of VMs used during execution of the program based on resource utilization metrics, the set of instructions specifying a resource utilization threshold, the specified configuration information further specifying a maximum number of VMs that can be used to execute the program;

selecting, by the one or more configured computing systems, a group of multiple computing VMs from the computing resources based on the first number of VMs, and initiating execution of the program by the group of VMs based on the specified configuration information;

monitoring, by the one or more configured computing systems during the execution of the program, a resource utilization of the group of VMs, wherein the resource utilization is based on a measured amount of a resource used by the group of VMs;

determining, by the one or more configured computing systems based on the resource utilization of the group of VMs over a time interval exceeding the resource utilization threshold, a second number of VMs to use for executing the program according to the set of instructions, wherein the second number of VMs is greater than the first number of VMs and is less than or equal to the maximum number of VMs; and

modifying, by the one or more configured computing systems, a quantity of VMs in the group of VMs for use in further execution of the program, wherein the modifying includes adding one or more additional VMs to the group of VMs while the execution of the program is ongoing and using the one or more additional VMs for further execution of the program, wherein adding the one or more additional VMs to the group of VMs includes allocating computing resources of one or more physical computing systems to the one or more additional VMs.

167. Nokia infringes claims of the '194 patent, for example, claim 1 of the '194 patent since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including but not limited to Nokia CloudBand Application Manager (CBAM).

168. Claim 1 of the '194 patent begins, "A computer-implemented method."

169. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager "CloudBand Application Manager automates lifecycle management by providing an open templating system, managing resources and applying associated workflows. It executes lifecycle management actions more easily and predictably than manual methods," as seen below.⁵⁰

⁵⁰Nokia, CloudBand Application Manager: Get a VNFM for OpenStack- and VMware-based Virtual Machines, <https://www.nokia.com/networks/core-networks/cloudband/application-manager/>.



Ready-to-use VNF lifecycle management

CloudBand Application Manager as an ETSI NFV-compliant Virtualized Network Function Manager (VNFM) automates VNF lifecycle management and cloud resource management, and its standards-based APIs make it easy to work with any vendor's VNF, Element Management System (EMS), Virtualized Infrastructure Manager (VIM), and NFV Orchestrator (NFVO). For the Nokia VNF portfolio, CloudBand Application Manager offers a ready-to-use, pre-integrated, one-click solution for lifecycle management.

CloudBand Application Manager automates lifecycle management by providing an open templating system, managing resources and applying associated workflows. It executes lifecycle management actions more easily and predictably than manual methods. It supports Nokia and third-party VNFs. Using OpenStack Heat orchestration templates, VMware Open Virtualization Format templates (OVF), Mistral workflows and Ansible playbooks, CloudBand Application Manager is open to a broad range of VNF onboarding options. It visualizes the structure and status of applications and performs lifecycle management, including basic functions (create, instantiate, scale, terminate, delete, operate, query and modify VNF), advanced functions (such as healing, update/patching, upgrades, backup and restore) and fault management for virtualized resources.

170. Claim 1 of the '194 patent further recites "receiving, by one or more configured computing systems of a program execution service that provides computing resources available to multiple users of the program execution service, instructions from a first user to execute a program using specified configuration information."

171. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager supports "[n]etwork function

virtualization (NFV)” that allows “[n]etwork administrators [to] dynamically deploy network elements and services without needing to physically provision the underlying routers,” as seen below.⁵¹

Managing VNFs using CBAM
What is network function virtualization?

NSP

2 Managing VNFs using CBAM

2.1 What is network function virtualization?

Network function virtualization (NFV) allows network administrators to uncouple network functions from underlay hardware NEs so that the functions can run as software images. These network functions include load balancers, firewalls, and NAT. The purpose of NFV is to provide a simpler way to deliver and manage the network components required for a virtualized infrastructure. Network administrators are able to dynamically deploy network elements and services without needing to physically provision the underlying routers. The virtualized network element that represents the physical NE is called a virtualized network function (VNF).

VNF management is provided by the external application CBAM, to which the NFM-P provides an interface using the Network Supervision application.

172. Furthermore, Nokia CloudBand Application Manager allows a “VNF network” to be configured using “[t]he VNF Descriptor [that] is a package that describes the configuration of the VNF network [and] consists of OpenStack Heat templates which define VNF specifications.”⁵²

173. Claim 1 of the ’194 patent further recites “the specified configuration information including a first number of virtual machines (VMs) to use to execute the program and a set of instructions specifying how to modify a number of VMs used during execution of the program based on resource utilization metrics.”

174. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, a VNF instance is configured using a “HEAT stack [that] describes the

⁵¹NSP Network Services Platform Release 22.9 Network Supervision Application Help (2022), https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE18136AAADTQZZA_V1_NSP%2022.9%20Network%20Supervision%20Application%20Help.pdf.

⁵²*Id.*

resources needed to create a VNF instance including the type and number of VMs, the VM image, storage, networks, routers, security groups and so on,” as seen below.⁵³

5.3.5 Deploying a VSR Instance Using OpenStack HEAT

Perform [Create the HEAT Stack](#) to deploy a VSR instance using OpenStack HEAT.

5.3.5.1 Introduction to OpenStack HEAT

VSR instances can also be deployed using OpenStack HEAT. HEAT provides template-based orchestration within OpenStack. A HEAT Orchestration Template (HOT) defines a HEAT stack. In an NFV context, the HEAT stack describes the resources needed to create a VNF instance including the type and number of VMs, the VM image, storage, networks, routers, security groups and so on. The template is defined in YAML format and it is reusable; that is, a template can be invoked multiple times to create several instances of a VNF.

175. Claim 1 of the '194 patent further recites “the set of instructions specifying a resource utilization threshold.”

176. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager supports “VNF threshold policies and policy templates,” as seen below.⁵⁴

⁵³Nokia VSR Installation and Setup Guide 20.2.R1 (2020), https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE15837AAAATQZZA01_V1_VSR%20Installation%20and%20Setup%20Guide%2020.2.R1.pdf.

⁵⁴NSP Network Services Platform Release 22.9 Network Supervision Application Help (2022), https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE18136AAADTQZZA_V1_NSP%2022.9%20Network%20Supervision%20Application%20Help.pdf.

1 Network Supervision

1.1 What is Network Supervision?

The Network Supervision application is a browser-based tool for monitoring the health of network equipment. It allows you to identify, investigate, and resolve problems with objects such as routers, cards, ports, physical links, and virtual network functions.

Network Supervision provides high-level visibility of equipment problems in the network, and allows quick navigation to specific objects for closer inspection and analysis. It also provides visualization tools and task flows for troubleshooting, to determine the root causes and impacts of equipment issues. A variety of display formats allows you to view and manage large amounts of information according to your needs and methods.

Network Supervision provides flexible task flows; there are multiple ways to identify and investigate problems. For example, you may decide to investigate issues based on active monitoring, Analytics reports, customer tickets, or historical problems. You can monitor and investigate links as well as NEs.

Network Supervision also allows you to monitor virtualized network functions (VNFs), such as load balancers, firewalls, and NAT, in a virtualized or datacenter environment where network functions are uncoupled from the underlying hardware. VNF management is provided by the external application CBAM, using access points available to Network Supervision. VNF threshold policies and policy templates are also supported.

177. Claim 1 of the '194 patent further recites “the specified configuration information further specifying a maximum number of VMs that can be used to execute the program.”

178. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, “VNFs can be scaled in or scaled out in CBAM” such that “[w]hen performing a scaling operation, you must specify a scaling aspect and a new level. The scaling level cannot exceed the maximum scaling level specified in the CBAM VNFD,” as seen below.⁵⁵

⁵⁵*Id.*

VNFs can be deleted in CBAM and advertised to the NFM-P via an LCN. When the NFM-P receives information on a deleted VNF, it removes the VNF from its database and unmanages the associated network element.

VNFs can be healed to trigger a reboot in CBAM or the Network Supervision application. Healing must be enabled in the CBAM VNFD before this operation can be performed in the GUI. If the VNFD requires additional parameters for VNF healing, the parameters are visible in the Network Supervision application.

VNFs can be scaled in or scaled out in CBAM or the Network Supervision application. Scaling must be enabled in the CBAM VNFD before this operation can be performed in the GUI. When performing a scaling operation, you must specify a scaling aspect and a new level. The scaling level cannot exceed the maximum scaling level specified in the CBAM VNFD. If the VNFD requires additional parameters for VNF scaling, the parameters are visible in the Network Supervision application.

179. Claim 1 of the '194 patent further recites “selecting, by the one or more configured computing systems, a group of multiple computing VMs from the computing resources based on the first number of VMs, and initiating execution of the program by the group of VMs based on the specified configuration information.”

180. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, the “VNFD - node templates” “[d]escribes the following virtual resources needed for VNF instantiation: VMs—virtual deployment units (VDUs) associated with a software image, flavor (vCPU and memory), and storage requirements,” as seen below.⁵⁶

⁵⁶Nokia VSR Installation and Setup Guide 20.2.R1 (2020), https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE15837AAAATQZZA01_V1_VSR%20Installation%20and%20Setup%20Guide%2020.2.R1.pdf, p. 232.

Table 37 V (Continued)

Term	Definition
VNFD - extensions	Modifiable parameters defined per VNF instance.
VNFD - interface	The supported lifecycle operations for a VNF that are implemented either through the CBAM built-in Mistral workflows or by custom Mistral workflows included in the VNF package.
VNFD - node templates	Describes the following virtual resources needed for VNF instantiation: <ul style="list-style-type: none"> • VMs—virtual deployment units (VDUs) associated with a software image, flavor (vCPU and memory), and storage requirements • internal networks—virtual links • VM connectivity—a set of virtual connection points that are associated with internal or external networks
VNFD - heat mapping	The HEAT mapping section of the VNFD associates the abstract resource descriptions with their corresponding implementations in HOT.
VNFM	VNF Manager The MANO component responsible for lifecycle management of VNF instances. Coordinates with EMS/NMS. This role is provided by CloudBand CBAM for VSR instances.
VxLAN	Virtual eXtensible Local Area Network A method of encapsulating Ethernet frames inside IP/UDP packets to create a tenant-specific overlay network within a data center.

181. Claim 1 of the '194 patent further recites “monitoring, by the one or more configured computing systems during the execution of the program, a resource utilization of the group of VMs, wherein the resource utilization is based on a measured amount of a resource used by the group of VMs.”

182. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager allows monitoring VNFs using

“threshold policy [that] allows you to monitor a set of pre-defined KPIs and create rules to define when the application indicates an overload, underload, or healing condition,” as seen below.⁵⁷

Network Supervision web application
Network Supervision web application

NFM-P

VNF threshold policies

You can assign a threshold policy to a VNF to allow the NFM-P to trigger automatic lifecycle management operations based on defined KPIs or alarms. A threshold policy allows you to monitor a set of pre-defined KPIs and create rules to define when the application indicates an overload, underload, or healing condition. The policy also allows you define an automatic triggered action to be performed when any of these conditions is met. These corrective actions include performing a scaling operation, performing a healing operation, or raising an alarm. When a lifecycle management action is triggered, the NFM-P automatically sends a lifecycle change notification to CBAM. You can assign only one VNF threshold policy to a VNF.

You can assign a CMM or CMG template to define a list of conditions and specify an action to be automatically performed when those conditions are met. The template can be used to create a VNF threshold policy, but you can modify the default conditions and actions imported from the template each time you create a new policy.

Use the policy agent in the Network Supervision application to create a VNF threshold policy. The policy agent includes default policy templates for the CMM and CMG. After you select a policy template, you can review the overload, underload, and healing settings to select the actions for each condition. You can also configure the hold time to specify how long the NFM-P should wait before performing the specified action.

183. Claim 1 of the '194 patent further recites “determining, by the one or more configured computing systems based on the resource utilization of the group of VMs over a time interval exceeding the resource utilization threshold, a second number of VMs to use for executing the program according to the set of instructions, wherein the second number of VMs is greater than the first number of VMs and is less than or equal to the maximum number of VMs.”

184. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager supports “[m]onitoring a VNF threshold policy” that determines “[w]hen a threshold is crossed and the VNF threshold policy

⁵⁷NSP Network Services Platform Release 18.6 NFV Solutions Guide (2018), https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14080AAABTQZZA01_V1_NSP%20NFM-P%2018.6%20NFV%20Solutions%20Guide.pdf.

triggers a scaling or healing operation” causing a “policy monitoring icon [to change] to red and the operation ... shown in-progress on the VNF tile,” as seen below.⁵⁸

Monitoring a VNF threshold policy

When a policy has been applied to a VNF, the policy monitoring icon appears on the VNF tile in the matrix view. You can click on the policy monitoring icon to view the status of each of the overload, underload, and healing threshold rules specified in the policy. The policy info view shows the

Release 18.6
June 2018
Issue 1

3HE-14080-AAAB-TQZZA

Network Supervision web application
Network Supervision web application

NFM-P

conditions, thresholds, and actions for each rule. It also shows the current value of the condition, allowing you to assess how close the condition is to reaching the specified threshold.

Figure 5-7 VNF threshold policy status



When a threshold is crossed and the VNF threshold policy triggers a scaling or healing operation, the policy monitoring icon changes to red and the operation is shown in-progress on the VNF tile. While the operation is being performed, you can disable the policy to cancel the action.

185. Furthermore, “VNFs can be scaled” such that “[w]hen performing a scaling operation, you must specify a scaling aspect and a new level. The scaling level cannot exceed the maximum scaling level specified in the CBAM VNFD.”⁵⁹

⁵⁸*Id.*

⁵⁹*Id.*

186. Claim 1 of the '194 patent further recites “modifying, by the one or more configured computing systems, a quantity of VMs in the group of VMs for use in further execution of the program, wherein the modifying includes adding one or more additional VMs to the group of VMs while the execution of the program is ongoing and using the one or more additional VMs for further execution of the program, wherein adding the one or more additional VMs to the group of VMs includes allocating computing resources of one or more physical computing systems to the one or more additional VMs.”

187. As a non-limiting example, Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager supports “Virtual Machine Manager (VMM)” that allows “automatic scale-out and automatic healing” that “automatically increases the processing capacity of the VMM,” as seen below.⁶⁰

NFV use cases
Automatic scale-out and healing

NFM-P

10 NFV use cases

10.1 Automatic scale-out and healing

Overview

You can configure the NFM-P to perform automatic scale-out and healing operations based on system-defined and user-defined processing thresholds and network events.

A scale-out operation increases the call processing capacity of a VNF. Processing capacity is increased by automatically creating and provisioning a VNFC. A healing operation reboots a VNFC. Automatic healing is triggered by the NFM-P when an alarm is raised against a VNFC. The NFM-P performs additional tasks based on the results of the attempted healing operation. You can enable automatic scale-out and automatic healing for the VMM or CMG. This configuration must be performed by a system administrator.

VMM automatic scale-out

You can enable automatic scale-out on the VMM and configure an automatic scale-out threshold. The automatic scale-out threshold defines the point at which an automatic scale-out operation can be triggered, where the threshold is the total UE capacity of the VMM multiplied by the scale-out factor. When the threshold is reached and/or one or more MAFs is affected by a resource overload node alarm, the NFM-P automatically increases the processing capacity of the VMM by creating an additional virtual MAF on the NE.

⁶⁰*Id.*

188. Nokia is and has been on notice of the infringement of the '194 patent at least as of the time Amazon filed and provided notice of this Complaint.

189. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '194 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, Nokia CloudBand Application Manager.

190. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 1 of the '194 patent.

191. Nokia will sell Nokia CloudBand Application Manager with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '194 patent.

192. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts will cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

193. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '194 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia CloudBand Application Manager, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '194 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT III: PATENT INFRINGEMENT OF U.S. PATENT NO. 9,253,211

194. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

195. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 23) of the '211 patent in violation of 35 U.S.C. § 271, and will continue to do so.

196. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software in violation of 35 U.S.C. § 271(a).

197. By way of example only, Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software meet all the limitations of at least independent claim 23 of the '211 patent, either literally or under the doctrine of equivalents.

198. Exemplary claim 23 of the '211 patent recites:

23. A system, comprising:

one or more computing systems each having one or more processors; and

at least one memory, the memory including instructions that, upon execution by at least one of the one or more processors, cause the system to:

receive, via an interface provided for use in configuring execution of programs by a program execution service having a plurality of computing nodes located in multiple geographic locations, a request from a client, the received request including configuration information that indicates one or more geographical locations in which at least one indicated program is to be executed;

select, based at least in part on the indicated one or more geographical locations, multiple computing nodes of the program execution service to use for execution of the indicated program; and

manage execution of one or more instances of the indicated program by the selected multiple computing nodes on behalf of the client, the managing of the execution being based at least in part on the received configuration information.

199. Nokia infringes claims of the '211 patent, for example, claim 23 of the '211 patent since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including but not limited to Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software.

200. For the preamble of claim 23, to the extent the preamble is determined to be limiting, Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software practice a “computing system comprising: one or more processors; and; a manager module of a configurable network service that is configured to, when executed by at least one of the processors, provide a virtual computer network for each of one or more clients.”

201. Nokia AirFrame Data Center and CloudBand Infrastructure Software practice this limitation. For example, Nokia AirFrame Data Center allows users to “build compact and efficient cloud computing data center” using “scalable processors,” as seen below.⁶¹

⁶¹Nokia, AirFrame Data Center: Adapt to Any Cloud-Based Application, <https://www.nokia.com/networks/data-center/airframe-data-center/>.

AirFrame data center solution elements

The Nokia AirFrame data center solution gives operators all the benefits and flexibility to build compact and efficient cloud computing data center that meet both Telco and IT requirements.

Nokia Data center solution is complemented with OpenStack compatible, OpenStack and Container distribution.

Nokia data center solution provides solution from far edge to central data centers with common automation and managements system with needed workflow tools.

Nokia AirFrame Rackmount server

Nokia AirFrame Rackmount is a high-performance, energy-efficient 2-socket solution for the heavy workload data centers.

The latest generation Intel Xeon scalable processors maximize platform's speed and crucial computing performance. Utilizes the latest NVMe SSDs in a wide range of forms to minimize energy use and to deliver the fastest bandwidth, higher IOPS and lower latency.

Complete portfolio of servers, storage, networking, accelerators and commodities help communication service providers quickly move to the cloud-defined data center. With a native support for Nokia Virtual Network Functions (VNFs) and Cloud Native Network Functions (CNFs), Nokia AirFrame Rackmount is a very important building block of 5G network deployments.



Product

AirFrame Rackmount
server →

202. Claim 23 of the '211 patent further recites “receive, via an interface provided for use in configuring execution of programs by a program execution service having a plurality of computing nodes located in multiple geographic locations, a request from a client, the received request including configuration information that indicates one or more geographical locations in which at least one indicated program is to be executed.”

203. Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software practice this limitation. For example, Nokia CloudBand Infrastructure Software supports a “Cloud Operations Manager [that] offers single pane of glass management of hybrid, geographically distributed Kubernetes and OpenStack cloud infrastructures, including dynamic cluster creation as well as management of cluster resources,” as seen below.⁶²

⁶²Cloud Operations Manager, <https://www.nokia.com/networks/core-networks/cloudband/cloud-operations-manager/>.



What is Nokia Cloud Operations Manager?

Nokia Cloud Operations Manager (previously [Nokia CloudBand Network Director](#)) automates lifecycle operations of containerized network functions (CNFs), virtual network functions (VNFs), and network services (NSs). It optimizes and governs platform resource usage and is designed for distributed, multi-tenant, multi-vendor cloud infrastructures.

Telco applications are increasingly leveraging cloud-native principles, which will often co-exist with VNFs. Multiple integrated applications provide an inter-related service as application suites or engineered systems such as [VoLTE](#) or [core networks](#). Cloud Operations Manager automates the lifecycle management of such [hybrid networks](#) to achieve agility and cost savings.

Cloud Operations Manager provides unified lifecycle management, centralized monitoring and visualization of CNFs, VNFs and NSs. It also offers single pane of glass management of hybrid, geographically distributed Kubernetes and OpenStack cloud infrastructures, including dynamic cluster creation as well as management of cluster resources.

Today, more than 50 customers around the world rely on Cloud Operations Manager for automated cloud lifecycle operations.

On this page

- ↓ [What is Nokia Cloud Operations Manager?](#)
- ↓ [Benefits and features](#)
- ↓ [Related topics](#)
- ↓ [Resources](#)
- ↓ [Related solutions and products](#)
- ↓ [Learn more](#)

204. Nokia CloudBand Application Manager supports “virtual network functions ... referred to as VNFs” and supports use of “TOSCA template” for “VNF lifecycle management,” as seen below.⁶³

⁶³Nokia, Nokia NSP Network Services Platform, https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE15164AAAETQZZA_V1_NSP%2019.11%20Network%20Supervision%20Application%20Help.pdf.

2.14 VNF custom actions

If there are custom actions defined in the TOSCA template for a VNF, you can trigger them using the Network Supervision application. Click on the More button on a VNF and then choose VNF > Custom Actions to view a list of any custom actions that have been defined for the VNF.

205. The TOSCA requests support “placement” of “TOSCA nodes” “by region[s],” including “[g]eographic regions (e.g., cities, municipalities, states, countries, etc.),” as seen below.⁶⁴

12.5.1.2 Use Case 2: Controlled placement by region

12.5.1.2.1 Description

This use case demonstrates the use of named “containers” which could represent the following:

- Datacenter regions
- Geographic regions (e.g., cities, municipalities, states, countries, etc.)
- Commercial regions (e.g., North America, Eastern Europe, Asia Pacific, etc.)

12.5.1.2.2 Features

This use case introduces the following policy features:

- Separation of resources (i.e., TOSCA nodes) by logical regions, or zones.

12.5.1.2.3 Sample YAML: Region separation amongst named set of regions

```
failover_policy_2:
  type: tosca.policy.placement
  description: My failover policy with allowed target regions (logical containers)
  properties:
    container type: region
    container_number: 3
    # If "containers" keyname is provided, they represent the allowed set
    # of target containers to use for placement for .
    containers: [ region1, region2, region3, region4 ]
```

206. Claim 23 of the '211 patent further recites “select, based at least in part on the indicated one or more geographical locations, multiple computing nodes of the program execution service to use for execution of the indicated program.”

207. Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software practice this limitation.

⁶⁴TOSCA Simple Profile in YAML Version 1.1, <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/csprd01/TOSCA-Simple-Profile-YAML-v1.1-csprd01.html>.

208. Nokia CloudBand supports Network Functions Virtualization (NFV), for example, by “support[ing] distributed NFV cloud infrastructures in a variety of ways” including “Placement Zones, which span geo-distributed datacenters,” as seen below.⁶⁵

CloudBand supports distributed NFV cloud infrastructures in a variety of ways. CloudBand provides aggregated northbound APIs that allow NFV applications and BSS/OSS to deal with the different locations as a single cloud. CloudBand provides a policy-based placement algorithm that computes a quasi-optimal location based on server utilization at the different locations, affinity and anti-affinity rules, and other parameters. In addition, CloudBand supports Placement Zones, which span geo-distributed datacenters by “aggregating” CloudBand Node OpenStack availability zones. Virtual machines belonging to vNFs can then be instantiated across these Placement Zones according to pre-defined business policies. CloudBand provides built-in load balancing services for scalable distributed network functions. The CloudBand graphical user portal gives users an aggregated view of the infrastructure according to their different roles and responsibilities. The image management component of CloudBand manages a single catalog of images automatically and assures that images are made available where they are needed. User accounts and key pairs are also managed at the global level.

209. Claim 23 of the '211 patent further recites “manage execution of one or more instances of the indicated program by the selected multiple computing nodes on behalf of the client, the managing of the execution being based at least in part on the received configuration information.”

210. Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software practice this limitation. For example, Nokia CloudBand Application Manager interacts with Nokia Cloud Operations Manager, previously Nokia CloudBand Network Director which automates lifecycle operations of containerized network functions (CNFs), virtual network functions (VNFs), and network services (NSs). It optimizes and governs platform resource usage and is designed for distributed, multi-tenant, multi-vendor cloud infrastructures.”⁶⁶ Nokia CloudBand also includes Nokia Container Services, which provides functionalities similar to those offered earlier by Nokia

⁶⁵Alcatel-Lucent, Cloudband with OpenStack as NFV Platform (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>.

⁶⁶Cloud Operations Manager, <https://www.nokia.com/networks/core-networks/cloudband/cloud-operations-manager/>.

CloudBand Network Director. Nokia Container Services is an “offering for deploying, orchestrating, monitoring and managing containers and container-based applications” and supports “complete infrastructure life cycle management functions,” “infrastructure scaling operations which allow addition or removal of nodes from a cluster,” and “heal[ing]” to “recover a failed node.”⁶⁷

⁶⁷Nokia Container Services, https://onestore.nokia.com/asset/207821?_gl=1*_1eyhxl*_gcl_au*MTQ5MzMwNjczOC4xNzE2OTY1OTE5*_ga*MTg4NDkzNjMxMi4xNzE2OTY1OTE3*_ga_D6GE5QF247*MTcyMTE2MTkxOS4yMC4xLjE3MjExNjQ5MTIuMC4wLjA.&_ga=2.19237244.578210260.1721069959-1884936312.1716965917, p. 2.

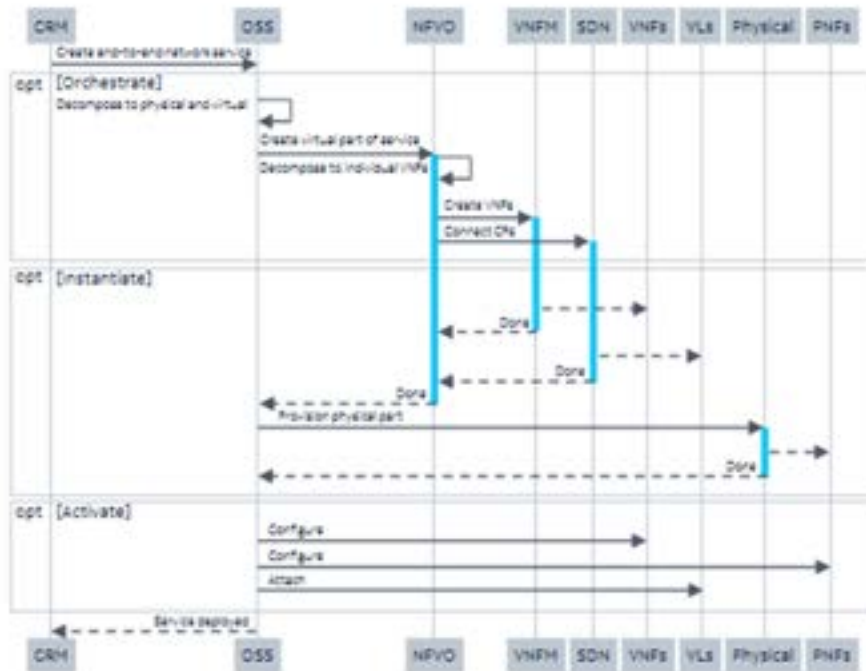


Figure 3: High level deployment flow

During the lifetime of the network service, the Network Director captures network service and VNF topology, monitors the health status of the service and scales, heals and updates the service according to policies described in the network service descriptor.

Monitoring, assurance and disaster recovery

CloudBand Network Director facilitates fault monitoring by tapping into capabilities of various fault management applications such as EMS/ NMS, VIM and VNFM. Faults received from these sources are enriched with information about associated network services to simplify network troubleshooting. Troubleshooting and automatic repair use a common topology and correlation engine based on a root cause analysis tool.

CloudBand Network Director, with its lifecycle management engine, supports several automated disaster recovery models that can be described and configured in the network service/VNF descriptor.

(Exhibit M, Nokia CloudBand Network Director, Product Information, <https://docplayer.net/19792521-Nokia-cloudband-network-director.html>, pp. 3-4.)

211. Nokia is and has been on notice of the infringement of the '211 patent at least as of the time Amazon filed and provided notice of this Complaint.

212. Nokia will infringe indirectly and continue to infringe indirectly one or more claims of the '211 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software.

213. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 23 of the '211 patent.

214. Nokia will sell Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '211 patent.

215. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts will cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

216. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '211 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '211 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT IV: PATENT INFRINGEMENT OF U.S. PATENT NO. 8,117,289

217. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

218. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 20) of the '289 patent in violation of 35 U.S.C. § 271, and will continue to do so.

219. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia AirFrame Data Center and CloudBand Infrastructure Software, in violation of 35 U.S.C. § 271(a).

220. By way of example only, Nokia AirFrame Data Center and CloudBand Infrastructure Software meet all the limitations of at least independent claim 20 of the '289 patent, either literally or under the doctrine of equivalents.

221. Exemplary claim 20 of the '289 patent recites:

20. A computing system comprising:

one or more processors; and;

a manager module of a configurable network service that is configured to, when executed by at least one of the processors, provide a virtual computer network for each of one or more clients, the providing of the virtual computer network for each of the one or more clients including:

receiving, by one or more configured computing systems of receiving information from the client for use in configuring the virtual computer network for the client, the configuring including specifying interconnections between multiple computing nodes of the virtual computer network and including providing an indicated type of functionality for

handling at least some communications between the multiple computing nodes; and

automatically providing the configured virtual computer network for the client in accordance with the configuring by overlaying the virtual computer network on a distinct substrate network, the multiple computing nodes being connected to the substrate network, and the automatic providing of the virtual computer network including:

selecting one or more network devices that are accessible via the substrate network and that are configured to provide the indicated type of functionality; and

forwarding multiple communications between the multiple computing nodes in accordance with the configuring, the forwarding including routing at least one of the multiple communications to at least one of the selected devices to enable the at least one selected device to provide the indicated type of functionality for the at least one communication.

222. Nokia infringes claims of the '289 patent, for example, claim 20 of the '289 patent since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including but not limited to Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software.

223. Claim 20 of the '289 patent begins, "A computing system comprising: one or more processors; and; a manager module of a configurable network service that is configured to, when executed by at least one of the processors, provide a virtual computer network for each of one or more clients."

224. Nokia AirFrame Data Center and CloudBand Infrastructure Software practice this limitation. For example, Nokia AirFrame Data Center includes the necessary hardware, software

and services that can adapt to any cloud-based application and supports creating a “scalable and distributed cloud-based architecture,” as seen below.⁶⁸



The acceleration of telco and IT convergence and the need to support a diverse range of demanding applications requires an innovative solution that takes all the benefits from the IT and open source domains to create a scalable and distributed cloud-based architecture.

More and more services, including 5G need network functionalities and capabilities located at the most efficient point within a network. This is necessary in order to address strict latency constraints and to process huge data demands that will be critical in delivering services with real-time responsiveness.

Network architectures need a re-think, with layered and distributed network topologies containing optimized hardware to deliver unparalleled performance and the greatest flexibility.

Nokia AirFrame data center solution is designed for running demanding virtualized and cloud-native software workloads. Security of the hardware and firmware is uncompromised and fulfills increased privacy and security needs for carrier grade networks. Enhancements including advanced packet, crypto, GPU and workload-specific acceleration make AirFrame to perform better than any traditional IT servers.

225. Nokia CloudBand software is a configurable network service for virtual network functions (VNFs) that allows “MANO (management and orchestration)” of NFV (Network

⁶⁸Nokia, AirFrame data center solution: Adapt to any cloud-based application, <https://networks.nokia.com/solutions/airframe-data-center-solution>.

functions virtualization) and “makes it simple to host, orchestrate, automate, and manage VNFs,” as seen below.⁶⁹



The image shows a dark teal rectangular box with the word "CloudBand" in large white font at the top. Below it, in a smaller white font, is the text "Deploy cloud infrastructure, management and orchestration for Virtual Machines and Container-based software".

Transition to cloud-native

The Telco cloud is in transition. Networks are shifting from monolithic Virtualized Network Functions (VNFs) running on Virtual Machines (VMs) to Cloud-native Network Functions (CNFs) running as micro-services in containers.

The shift promises greater choice in deployment environments (private or public cloud), locations (data center or edge), and a more efficient use of computing resources. It also promises simpler, less costly and less risky application and infrastructure enhancements and upgrades through fully automated software onboarding, installation and verification testing functions. These characteristics are practically pre-requisites to support 5G's use cases.

ETSI NFV deployments

Nokia is well positioned to assist you with this transition. For starters, Nokia's CloudBand software is a widely deployed system for ETSI NFV MANO (management and orchestration), with commercially proven reliability, automation, repeatability and security.

Our CloudBand portfolio makes it simple to host, orchestrate, automate, and manage VNFs and it manages tens of thousands of servers across more than 200 service providers around the globe.

⁶⁹CloudBand Deploy cloud infrastructure, management and orchestration for Virtual Machines and Container-based software, <https://www.nokia.com/networks/core-networks/cloudband/>.

226. Claim 20 of the '289 patent further recites “the providing of the virtual computer network for each of the one or more clients including: receiving, by one or more configured computing systems of receiving information from the client for use in configuring the virtual computer network for the client, the configuring including specifying interconnections between multiple computing nodes of the virtual computer network and including providing an indicated type of functionality for handling at least some communications between the multiple computing nodes.”



227. Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software practice this limitation. For example, Nokia CloudBand supports configuration of a VNF (virtual network function) using a “TOSCA template” (Topology and Orchestration Specification for Cloud Applications template), as seen below.⁷⁰

2.14 VNF custom actions

If there are custom actions defined in the TOSCA template for a VNF, you can trigger them using the Network Supervision application. Click on the More button on a VNF and then choose VNF > Custom Actions to view a list of any custom actions that have been defined for the VNF.

2.15 How do I trigger a custom action on a VNF?

You can trigger custom actions that have been defined for a VNF using the Network Supervision application. The actions must be defined in a TOSCA template that has been configured on the CBAM.

- 1 
Click on the More button on a VNF panel and select VNF > Custom Actions. The Custom Actions window appears.
- 2 
Select a custom action from the drop-down list and click OK. The custom action is triggered.

228. Claim 20 of the '289 patent further recites “the configuring including specifying interconnections between multiple computing nodes of the virtual computer network and including

⁷⁰Nokia, Nokia NSP Network Services Platform, https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE15164AAAETQZZA_V1_NSP%2019.11%20Network%20Supervision%20Application%20Help.pdf, p. 28.

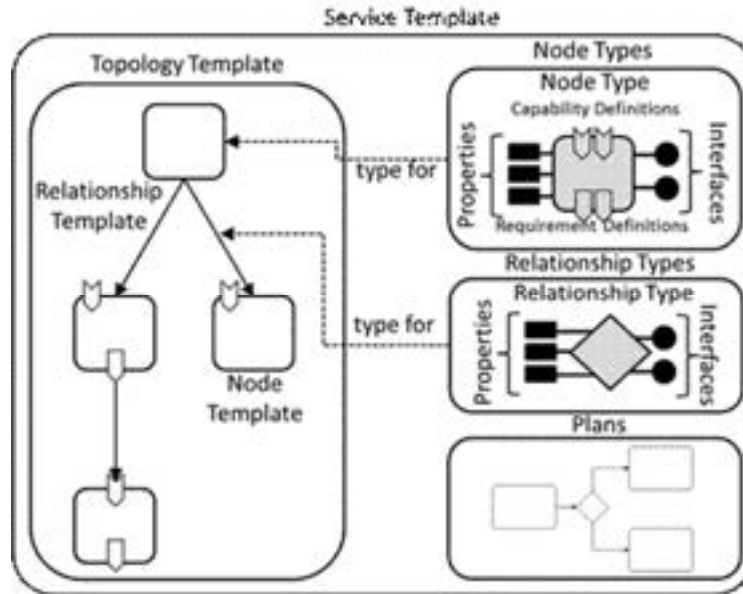
providing an indicated type of functionality for handling at least some communications between the multiple computing nodes.”

229. Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software practice this limitation. For example, the “TOSCA metamodel uses the concept of service templates to describe cloud workloads as a topology template, which is a graph of node templates modeling the components a workload is made up of and as relationship templates modeling the relations between those components. TOSCA further provides a type system of node types to describe the possible building blocks for constructing a service template, as well as relationship type to describe possible kinds of relations.”⁷¹ Furthermore the “core TOSCA specification provides a language to describe service components and their relationships using a service topology, and it provides for describing the management procedures that create or modify services using orchestration processes.”⁷² Furthermore, “[a] Topology Template consists of a set of Node Templates and Relationship Templates that together define the topology model of a service as [a] directed graph. A node in this graph is represented by a Node Template. A Node Template specifies the occurrence of a Node Type as a component of a service.”⁷³ Accordingly, Nokia CloudBand cloud infrastructure uses TOSCA templates to configure the virtual computer network by specifying interconnections between multiple computing nodes of the virtual computer network, as seen below. (*See id.* at FIG. 1.)

⁷¹TOSCA Simple Profile for Network Functions Virtualization (NFV) Version 1.0, Oasis Documentation, http://docs.oasis-open.org/tosca/tosca-nfv/v1.0/csd04/tosca-nfv-v1.0-csd04.html#_Toc482896040.

⁷²Topology and Orchestration Specification for Cloud Applications Version 1.0, <http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.html>.

⁷³*Id.*



230. Claim 20 of the '289 patent further recites “providing an indicated type of functionality for handling at least some communications between the multiple computing nodes.”

231. Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software practice this limitation. The '289 patent describes load balancing as a networking related functionality. (See '289 patent at 3:40-48.) Nokia CloudBand cloud infrastructure uses TOSCA templates to define a type of functionality for handling at least some communications between the multiple computing nodes, for example, load-balancing, as seen below.⁷⁴

⁷⁴TOSCA Simple Profile in YAML Version 1.1, <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/csprd01/TOSCA-Simple-Profile-YAML-v1.1-csprd01.html>.

5.9.12 tosca.nodes.LoadBalancer

The TOSCA **Load Balancer** node represents logical function that be used in conjunction with a Floating Address to distribute an application's traffic (load) across a number of instances of the application (e.g., for a clustered or scaled application).

Shorthand Name	LoadBalancer
Type Qualified Name	tosca:LoadBalancer
Type URI	tosca.nodes.LoadBalancer

5.9.12.1 Definition

```
tosca.nodes.LoadBalancer:
  derived_from: tosca.nodes.Root
  properties:
    algorithm:
      type: string
      required: false
      status: experimental
  capabilities:
    client:
      type: tosca.capabilities.Endpoint.Public
      occurrences: [0, UNBOUNDED]
      description: the Floating (IP) client's on the public network can connect to
  requirements:
    - application:
      capability: tosca.capabilities.Endpoint
      relationship: tosca.relationships.RoutesTo
      occurrences: [0, UNBOUNDED]
      description: Connection to one or more load balanced applications
```

5.9.12.2 Notes:

- A **LoadBalancer** node can still be instantiated and managed independently of any applications it would serve; therefore, the load balancer's **application** requirement allows for zero occurrences.

232. Claim 20 of the '289 patent further recites “automatically providing the configured virtual computer network for the client in accordance with the configuring by overlaying the virtual

computer network on a distinct substrate network, the multiple computing nodes being connected to the substrate network.”

233. Nokia CloudBand Infrastructure Software practices this limitation. For example, Nokia CloudBand Infrastructure Software allows configuring virtual network functions (VNFs) and supports “dynamic overlay network layer to address the needs of virtual network functions,” as seen below.⁷⁵

NETWORKING

Virtual network functions vary widely in their network demands. Due to their distribution throughout an NFV infrastructure, the baseline requirement for an NFV network is connectivity — not only within the datacenter but also across the wide area. For security reasons, different network functions should only be connected to each other if they need to exchange data, and the NFV control, data and management traffic should be separated. As network functions are decomposed, for example into data plane components and a centralized control plane component, the wide area connectivity between these components needs to remain as highly reliable as with traditional integrated architectures. Enough network resources should be available to ensure that surging traffic from other applications cannot adversely affect the NFV applications. The network should be resilient against equipment failures and force majeure disasters. Latency and jitter requirements vary from hundreds of milliseconds for some control and management systems to single digit milliseconds for mobile gateways and cloud radio access networks.

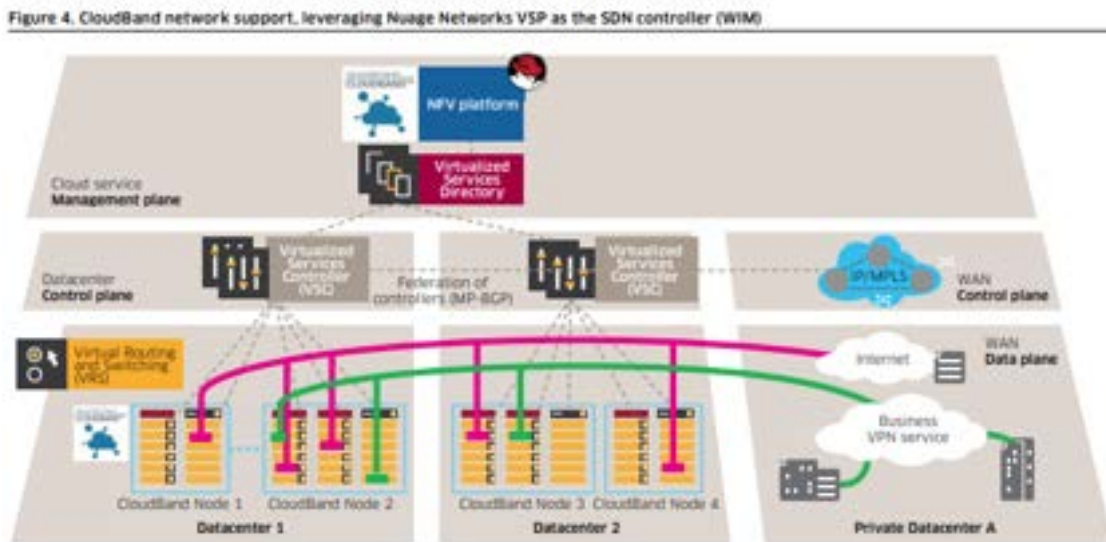
NFV networks will typically consist of a semi-static physical infrastructure along with a much more dynamic overlay network layer to address the needs of virtual network functions. The overlay layer needs to respond quickly to changing service demand, new service deployments and so on.

234. Claim 20 of the '289 patent further recites “the automatic providing of the virtual computer network including: selecting one or more network devices that are accessible via the substrate network and that are configured to provide the indicated type of functionality; and forwarding multiple communications between the multiple computing nodes in accordance with the configuring, the forwarding including routing at least one of the multiple communications to

⁷⁵Alcatel-Lucent, Cloudband with OpenStack as NFV Platform (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>, p. 5.

at least one of the selected devices to enable the at least one selected device to provide the indicated type of functionality for the at least one communication.”

235. Nokia CloudBand Infrastructure Software practices this limitation. For example, as described above, “[Network Functions Virtualization (NFV)] networks will typically consist of a semi-static physical infrastructure,”⁷⁶ which are “one or more network devices that are accessible via the substrate network.” Nokia CloudBand provides the dynamic overlay that connects to the physical infrastructure devices. CloudBand is connected to the physical infrastructure: it “delivers NFV networking abstractions that extend from the datacenter across the WAN toward multiple locations (Figure 4).”⁷⁷



(Alcatel-Lucent, Cloudband with OpenStack as NFV Platform (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>.)

⁷⁶*Id.*

⁷⁷*Id.* at pp. 5-6.

236. Further, “CloudBand Route Domains and Network Templates support flexible aggregations of networks. CloudBand is open to interface with any networking framework using standard OpenStack Neutron APIs and plugins. CloudBand can interface with existing MPLS networks and other legacy networks using a process called VPN stitching.”⁷⁸

237. Further, CloudBand Infrastructure Software overlay substrate “delivers” access and configuration “to provide the indicated functionality” for the physical network. For example, Nokia CloudBand Infrastructure Software includes the “Nokia Virtualized Service Router (VSR) [that] is a highly flexible virtualized IP edge router designed and optimized for telco cloud environments,” as seen below.⁷⁹

⁷⁸*Id.* at p. 5.

⁷⁹Virtualized service router, <https://www.nokia.com/networks/ip-networks/virtualized-service-router/>.



The Nokia Virtualized Service Router (VSR) is a highly flexible virtualized IP edge router designed and optimized for telco cloud environments. It enables rapid service innovation while lowering operating costs. With the VSR, service providers can extend service reach, open new markets, and accelerate time to market.

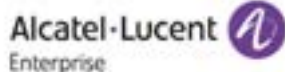
The VSR supports a broad and rich set of virtualized IP edge applications – virtualized network functions (VNFs) to support mobile, business, and residential services.

Based on the Service Routing Operating System (SR OS), the VSR delivers high performance and elastic scalability for x86 server environments.

The VSR supports the highest levels of reliability and resiliency. It also offers flexible management options – from working with open frameworks to delivering OpenStack®-integrated VNF management and element management capabilities with our Network Services Platform (NSP).

238. Nokia Virtualized Service Router (VSR) supports “forwarding” and “routing,” as seen below.⁸⁰

⁸⁰Nokia Virtualized Service Router, <https://www.al-enterprise.com/-/media/assets/internet/documents/ale-nokia-virtualized-service-router-ip-mpls-datasheet-en.pdf>, p. 4.



- Ease of interoperability: Using standardized and open interfaces allows for integration in a wide variety of deployment environments.
- Optimized use of resources and improved telecommunications and IT integration: Proven high performance with optimized use of resources on a standardized x86 compute platform for different applications, users and tenants enables rollout of profitable services based on measurable business models.

VSR architecture

Nokia has leveraged its leading expertise and innovation in service routing and has architected and optimized the Nokia VSR for the x86-based server architecture by applying advanced design concepts, principles and approaches, including:

- Separation of control plane and data plane tasks: Allows for independent scaling of control plane and data plane within the same VM.
- A virtual Forwarding Path (vFP): The vFP is the x86-optimized forwarding path that supports data path functions, including access control lists, QoS classification, policing, Forwarding Information Base (FIB) lookup, and related packet-processing functions.

As a result of these advanced design concepts, principles and approaches, Nokia's flexible and robust virtualized router implementation on the VSR allows:

- Optimal utilization of hypervisor (host) resources
- High performance for both control plane (routing) and data plane (packet forwarding) functions
- Separation of control plane and data plane CPU cores
- Advanced multi-system redundancy features
- Resilient cloud scaling
- Superior life-cycle management capabilities with a unique approach to consistent operations across physical and virtualized network elements.

VSR deployment

The Nokia VSR is deployed in an integrated model, where the VSR control plane and data plane functionality are implemented on a single VM. In this model, the virtual CPU and memory of the VM are shared among:

- Control tasks, including:

The VSR architecture provides forwarding by a “virtual Forwarding Path (vFP): The vFP is the x86-optimized forwarding path that supports data path functions, including access control lists, QoS classification, policing, Forwarding Information Base (FIB) lookup, and related packet-processing functions.”⁸¹ The VSR also provides the full range of routing capabilities: “Provides a full range of IP/MPLS edge services: business, residential, and mobile; Flexibly scales IP routing infrastructure by supporting a wide range of virtualized routing applications; Leverages a rich legacy of service robustness and stability from the Service Router portfolio, and is fully optimized for the x86 server operating environment.”⁸²

⁸¹*Id.*

⁸²Virtualized Service Router, <https://www.nokia.com/networks/ip-networks/virtualized-service-router/#benefits-features>.

239. Furthermore, Nokia CloudBand Infrastructure Software can be configured to use a “TOSCA Load Balancer node [that] represents logical function that be used in conjunction with a Floating Address to distribute an application’s traffic (load) across a number of instances of the application.”⁸³ As a logical function, the TOSCA Load Balancer is a layer on top of the substrate network which connects through the layer to the network devices accessible via the substrate network. Distribution of the application’s traffic across a number of instances forwards or routes application requests to the underlying network devices processing those application instances. Load balancing represents an example of an indicated type of functionality performed by the Nokia CloudBand Infrastructure Software using TOSCA Load Balancer, as seen below.⁸⁴

⁸³TOSCA Simple Profile in YAML Version 1.1, <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/csprd01/TOSCA-Simple-Profile-YAML-v1.1-csprd01.html>.

⁸⁴*Id.*

5.9.12 `tosca.nodes.LoadBalancer`

The TOSCA **Load Balancer** node represents logical function that be used in conjunction with a Floating Address to distribute an application's traffic (load) across a number of instances of the application (e.g., for a clustered or scaled application).

Shorthand Name	LoadBalancer
Type Qualified Name	tosca:LoadBalancer
Type URI	tosca.nodes.LoadBalancer

5.9.12.1 Definition

```
tosca.nodes.LoadBalancer:
  derived_from: tosca.nodes.Root
  properties:
    algorithm:
      type: string
      required: false
      status: experimental
  capabilities:
    client:
      type: tosca.capabilities.Endpoint.Public
      occurrences: [0, UNBOUNDED]
      description: the Floating (IP) client's on the public network can connect to
  requirements:
    - application:
      capability: tosca.capabilities.Endpoint
      relationship: tosca.relationships.RoutesTo
      occurrences: [0, UNBOUNDED]
      description: Connection to one or more load balanced applications
```

240. Nokia is and has been on notice of the infringement of the '289 patent at least as of the time Amazon filed and provided notice of this Complaint.

241. Nokia will infringe indirectly and continue to infringe indirectly one or more claims of the '289 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software.

242. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 20 of the '289 patent.

243. Nokia will sell Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '289 patent.

244. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts will cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

245. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '289 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '289 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT V: PATENT INFRINGEMENT OF U.S. PATENT NO. 8,296,419

246. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

247. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 4) of the '419 patent in violation of 35 U.S.C. § 271, and will continue to do so.

248. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia CloudBand, in violation of 35 U.S.C. § 271(a).

249. By way of example only, Nokia CloudBand meets all the limitations of at least independent claim 4 of the '419 patent, either literally or under the doctrine of equivalents.

250. Exemplary claim 4 of the '419 patent recites:

4. A computer-implemented method comprising:

receiving, by one or more computing systems configured to provide a distributed program execution service having a plurality of computing nodes, configuration information regarding executing an indicated program on an indicated quantity of multiple of the plurality of computing nodes, wherein the executing of the indicated program causes a plurality of jobs to be executed;

initiating at a first time, by the one or more configured computing systems, the executing of the indicated program in a distributed manner on the multiple computing nodes in such a manner that one or more of the jobs of the indicated program are attempted to be executed on each of the multiple computing nodes;

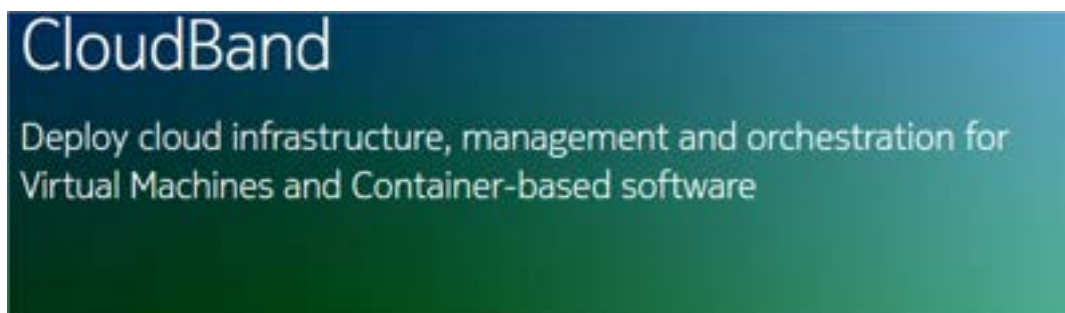
determining, by the one or more configured computing systems at a second time subsequent to the first time, whether a minimum subset of the multiple computing nodes have begun to execute the jobs of the indicated program as expected; and

in response to the determining, initiating a change in a quantity of the multiple computing nodes that are used for executing the jobs of the indicated program.

251. Nokia infringes claims of the '419 patent, for example, claim 4 of the '419 patent since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including but not limited to Nokia CloudBand.

252. For the preamble of claim 4, to the extent the preamble is determined to be limiting, Nokia CloudBand practices a “computer-implemented method comprising.”

253. For example, “Our CloudBand portfolio makes it simple to host, orchestrate, automate, and manage VNFs and it manages tens of thousands of servers across more than 200 service providers around the globe,” as seen below.⁸⁵



Transition to cloud-native

The Telco cloud is in transition. Networks are shifting from monolithic Virtualized Network Functions (VNFs) running on Virtual Machines (VMs) to Cloud-native Network Functions (CNFs) running as micro-services in containers.

The shift promises greater choice in deployment environments (private or public cloud), locations (data center or edge), and a more efficient use of computing resources. It also promises simpler, less costly and less risky application and infrastructure enhancements and upgrades through fully automated software onboarding, installation and verification testing functions. These characteristics are practically pre-requisites to support 5G's use cases.

ETSI NFV deployments

Nokia is well positioned to assist you with this transition. For starters, Nokia's CloudBand software is a widely deployed system for ETSI NFV MANO (management and orchestration), with commercially proven reliability, automation, repeatability and security.

Our CloudBand portfolio makes it simple to host, orchestrate, automate, and manage VNFs and it manages tens of thousands of servers across more than 200 service providers around the globe.

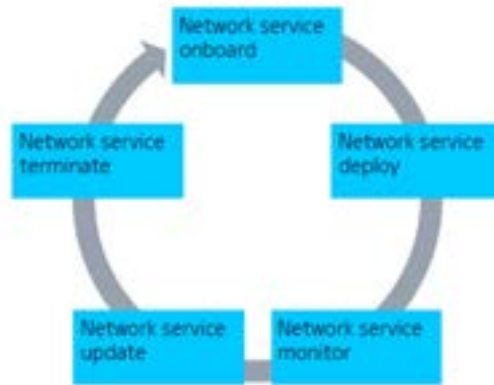
254. Claim 4 of the '419 patent further recites “receiving, by one or more computing systems configured to provide a distributed program execution service having a plurality of computing nodes, configuration information regarding executing an indicated program on an indicated quantity of multiple of the plurality of computing nodes, wherein the executing of the indicated program causes a plurality of jobs to be executed.”

⁸⁵CloudBand Deploy cloud infrastructure, management and orchestration for Virtual Machines and Container-based software, <https://www.nokia.com/networks/core-networks/cloudband/>.

255. Nokia CloudBand practices this limitation. For example, “Once a network service or function has been onboarded, higher level orchestrators or human operators can trigger the Network Director to deploy or update an instance of the network service as part of the service orchestration process (Figure 3). As a first step of this process, the set of VNFs and VNF forwarding graphs is determined based on a TOSCA template.”⁸⁶ Nokia CloudBand also includes Nokia Container Services, which provides functionalities similar to those offered earlier by Nokia CloudBand Network Director. Nokia Container Services is an “offering for deploying, orchestrating, monitoring and managing containers and container-based applications” that supports “complete infrastructure life cycle management functions,” “infrastructure scaling operations which allow addition or removal of nodes from a cluster,” and “heal[ing]” to “recover a failed node.”⁸⁷

⁸⁶Exhibit M, Nokia CloudBand Network Director, Product Information, <https://docplayer.net/19792521-Nokia-cloudband-network-director.html>, p. 4.

⁸⁷Nokia Container Services, https://onestore.nokia.com/asset/207821?_gl=1*_1eyhxl*_gcl_au*MTQ5MzMwNjczOC4xNzE2OTY1OTE5*_ga*MTg4NDkzNjMxMi4xNzE2OTY1OTE3*_ga_D6GE5QF247*MTcyMTE2MTkxOS4yMC4xLjE3MjExNjQ5MTIuMC4wLjA.&_ga=2.19237244.578210260.1721069959-1884936312.1716965917, p. 2.



Once a network service or function has been onboarded, higher level orchestrators or human operators can trigger the Network Director to deploy or update an instance of the network service as part of the service orchestration process (Figure 3). As a first step of this process, the set of VNFs and VNF forwarding graphs is determined based on a TOSCA template. Then, the Network Director identifies a generic VNF manager, such as CloudBand Application Manager, or VNF-specific VNF managers, to execute the lifecycle of the VNFs. If necessary, an appropriate VNF manager is instantiated. It then instructs the VNF manager to deploy or update and commission the VNFs and commission connected PNFs to work with the VNFs.

(Exhibit M, Nokia CloudBand Network Director, Product Information, <https://docplayer.net/19792521-Nokia-cloudband-network-director.html>, pp. 3-4.)

256. Claim 4 of the '419 patent further recites “initiating at a first time, by the one or more configured computing systems, the executing of the indicated program in a distributed manner on the multiple computing nodes in such a manner that one or more of the jobs of the indicated program are attempted to be executed on each of the multiple computing nodes.”

257. Nokia CloudBand practices this limitation. For example, “CloudBand provides a policy-based placement algorithm that computes a quasi-optimal location based on server utilization at the different locations, affinity and anti-affinity rules, and other parameters. In addition, CloudBand supports Placement Zones, which span geo-distributed datacenters by ‘aggregating’ CloudBand Node OpenStack availability zones. Virtual machines belonging to

vNFs can then be instantiated across these Placement Zones according to pre-defined business policies,” as seen below.⁸⁸

CloudBand supports distributed NFV cloud infrastructures in a variety of ways. CloudBand provides aggregated northbound APIs that allow NFV applications and BSS/OSS to deal with the different locations as a single cloud. CloudBand provides a policy-based placement algorithm that computes a quasi-optimal location based on server utilization at the different locations, affinity and anti-affinity rules, and other parameters. In addition, CloudBand supports Placement Zones, which span geo-distributed datacenters by “aggregating” CloudBand Node OpenStack availability zones. Virtual machines belonging to vNFs can then be instantiated across these Placement Zones according to pre-defined business policies. CloudBand provides built-in load balancing services for scalable distributed network functions. The CloudBand graphical user portal gives users an aggregated view of the infrastructure according to their different roles and responsibilities. The image management component of CloudBand manages a single catalog of images automatically and assures that images are made available where they are needed. User accounts and key pairs are also managed at the global level.

CloudBand with OpenStack as NFV Platform
COLLABORATIVE WHITE PAPER BETWEEN ALCATEL-LUCENT AND RED HAT
4

(Alcatel-Lucent, Cloudband with OpenStack as NFV Platform (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>, p. 4.)

258. Claim 4 of the '419 patent further recites “determining, by the one or more configured computing systems at a second time subsequent to the first time, whether a minimum subset of the multiple computing nodes have begun to execute the jobs of the indicated program as expected.”

259. Nokia CloudBand practices this limitation. For example, “You can assign a threshold policy to a VNF to allow the NFM-P to trigger automatic lifecycle management operations based on defined KPIs or alarms. A threshold policy allows you to monitor a set of

⁸⁸Alcatel-Lucent, Cloudband with OpenStack as NFV Platform (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>, p. 4.

pre-defined KPIs and create rules to define when the application indicates an overload, underload, or healing condition,” as seen below.⁸⁹

4.1.7 VNF threshold policies

You can assign a threshold policy to a VNF to allow the NFM-P to trigger automatic lifecycle management operations based on defined KPIs or alarms. A threshold policy allows you to monitor a set of pre-defined KPIs and create rules to define when the application indicates an overload, underload, or healing condition. The policy also allows you to define an automatic triggered action to be performed when any of these conditions is met. These corrective actions include performing a scaling operation, performing a healing operation, or raising an alarm. When a lifecycle management action is triggered, the NFM-P automatically sends a lifecycle change notification to CBAM.

You can create a CMM or CMG template to define a list of conditions and specify an action to be automatically performed when those conditions are met. The template can be used to create a VNF threshold policy, but you can modify the default conditions and actions imported from the template each time you create a new policy.

Release 17.3
March 2017
Issue 1

3HE-11999-AAAA-TQZZA

19

260. Claim 4 of the '419 patent further recites “in response to the determining, initiating a change in a quantity of the multiple computing nodes that are used for executing the jobs of the indicated program.”

261. Nokia CloudBand practices this limitation. For example, “The policy also allows you to define an automatic triggered action to be performed when any of these conditions is met. These corrective actions include performing a scaling operation, performing a healing operation, or raising an alarm. When a lifecycle management action is triggered, the NFM-P automatically sends a lifecycle change notification to CBAM.”⁹⁰

262. Nokia is and has been on notice of the infringement of the '419 patent at least as of the time Amazon filed and provided notice of this Complaint.

⁸⁹NSP Network Services Platform, Network Functions Manager - Packet (NFM-P) Release 17.3, NFV Solutions Guide, https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE11999AAAATQZZA01_V1_NSP%20NFM-P%2017.3%20NFV%20Solutions%20Guide.pdf, p. 19.

⁹⁰*Id.*

263. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '419 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, Nokia CloudBand.

264. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 4 of the '419 patent.

265. Nokia will sell Nokia CloudBand with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '419 patent.

266. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts would cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

267. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '419 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia CloudBand, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '419 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT VI: PATENT INFRINGEMENT OF U.S. PATENT NO. 9,106,540

268. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

269. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 4) of the '540 patent in violation of 35 U.S.C. § 271, and will continue to do so.

270. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia CloudBand Infrastructure Software, in violation of 35 U.S.C. § 271(a).

271. By way of example only, Nokia CloudBand Infrastructure Software meets all the limitations of at least independent claim 4 of the '540 patent, either literally or under the doctrine of equivalents.

272. Exemplary claim 4 of the '540 patent recites:

4. A computer-implemented method for providing logical networking functionality for computer networks, the method comprising:

receiving, by one or more configured computing systems of a configurable network service, configuration information for a first virtual computer network that includes multiple computing nodes arranged via a specified network topology, the configuration information indicating a specified virtual router device that logically interconnects at least first and second groups of the multiple computing nodes;

the multiple computing nodes for information from the specified virtual router device, wherein the responding includes emulating functionality of the specified virtual router device in generating response information for the one computing node without using any physical router devices to represent the specified virtual router device; and

forwarding, by the one or more configured computing systems, a communication between the first and second groups of computing nodes by emulating further functionality of the specified virtual router device for modifying the communication as part of the forwarding.

273. Nokia infringes claims of the '540 patent, including for example claim 4 of the '540 patent, since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including, but not limited to Nokia CloudBand.

274. Claim 4 of the '540 patent begins, "A computer-implemented method for providing logical networking functionality for computer networks." As a non-limiting example, Nokia performs this method using Nokia CloudBand to implement logical networking functionality for computer networks.

275. Nokia CloudBand practices this limitation. For example, Nokia CloudBand allows hosting, orchestrating, automating, and managing of virtualized network functions (VNFs) and services, as seen below.⁹¹

ETSI NFV deployments

Nokia is well positioned to assist you with this transition. For starters, Nokia's CloudBand software is a widely deployed system for ETSI NFV MANO (management and orchestration), with commercially proven reliability, automation, repeatability and security.

Our CloudBand portfolio makes it simple to host, orchestrate, automate, and manage VNFs and it manages tens of thousands of servers across more than 200 service providers around the globe.

Nokia can help

Our cloud management portfolio makes it simple to host, orchestrate, automate, and manage VNFs, CNFs and services. Our scalable and flexible software helps you reduce time to revenue for new services and use automation and optimization to make network operations agile and efficient. With Nokia, you benefit from the rapid pace of open-source innovation and get a carrier-grade product that is ready for deployment.

276. Claim 4 of the '540 patent further recites "receiving, by one or more configured computing systems of a configurable network service, configuration information for a first virtual computer network that includes multiple computing nodes arranged via a specified network

⁹¹Nokia, CloudBand: Deploy cloud infrastructure, management and orchestration for Virtual Machines and Container-based software, <https://networks.nokia.com/products/cloudband>.

topology, the configuration information indicating a specified virtual router device that logically interconnects at least first and second groups of the multiple computing nodes.”

277. Nokia CloudBand practices this limitation. For example, “CloudBand Application Manager as an ETSI NFV-compliant Virtualized Network Function Manager (VNFM) automates VNF lifecycle management and cloud resource management” and “It executes lifecycle management actions more easily and predictably than manual methods. It supports Nokia and third-party VNFs,” as seen below.⁹²



278. Furthermore, Nokia CloudBand supports “OpenStack Neutron” which is “the networking component offering abstractions, such as Layer 2 and Layer 3 networks, subnets, IP

⁹²Nokia, CloudBand Application Manager: Get a VNFM for OpenStack - and VMware-based Virtual Machines, <https://www.nokia.com/networks/core-networks/cloudband/application-manager/>.

addresses, and virtual middleboxes.”⁹³ Furthermore, “NFV networks ... typically consist of a semi-static physical infrastructure along with a much more dynamic overlay network layer to address the needs of virtual network functions,” as seen below.⁹⁴

NFV networks will typically consist of a semi-static physical infrastructure along with a much more dynamic overlay network layer to address the needs of virtual network functions. The overlay layer needs to respond quickly to changing service demand, new service deployments and so on.

OpenStack Neutron is the networking component offering abstractions, such as Layer 2 and Layer 3 networks, subnets, IP addresses, and virtual middleboxes. Neutron has a plugin-based architecture. Networking requests to Neutron are forwarded to the Neutron plugin installed to handle the specifics of the present network. Neutron is limited to a single space of network resources typically associated with an OpenStack region. There is no capability to federate multiple network domains and manage WAN capabilities.

CloudBand delivers NFV networking abstractions that extend from the datacenter across the WAN toward multiple locations (Figure 4). CloudBand Route Domains and Network Templates support flexible aggregations of networks. CloudBand is open to interface with any networking framework using standard OpenStack Neutron APIs and plugins. CloudBand can interface with existing MPLS networks and other legacy networks using a process called VPN stitching.

For example, CloudBand supports Nuage Networks Virtualized Services Platform (VSP), a second generation SDN solution. Nuage Networks VSP supports the federation of multiple SDN network domains each with their own SDN controller. Networks are synchronized in both directions between CloudBand and Nuage Networks VSP. The SDN controller, in this case Nuage Networks VSP, has been recently identified by ETSI NFV Industry Specification Group and is now referenced as the WAN Infrastructure Manager (WIM).

279. Furthermore, Nokia supports “CloudBand Route Domains and Network Templates support flexible aggregations of networks. CloudBand is open to interface with any networking framework using standard OpenStack Neutron APIs and plugins.”⁹⁵

⁹³Alcatel-Lucent, *Cloudband with OpenStack as NFV Platform* (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>, p. 5.

⁹⁴*Id.*

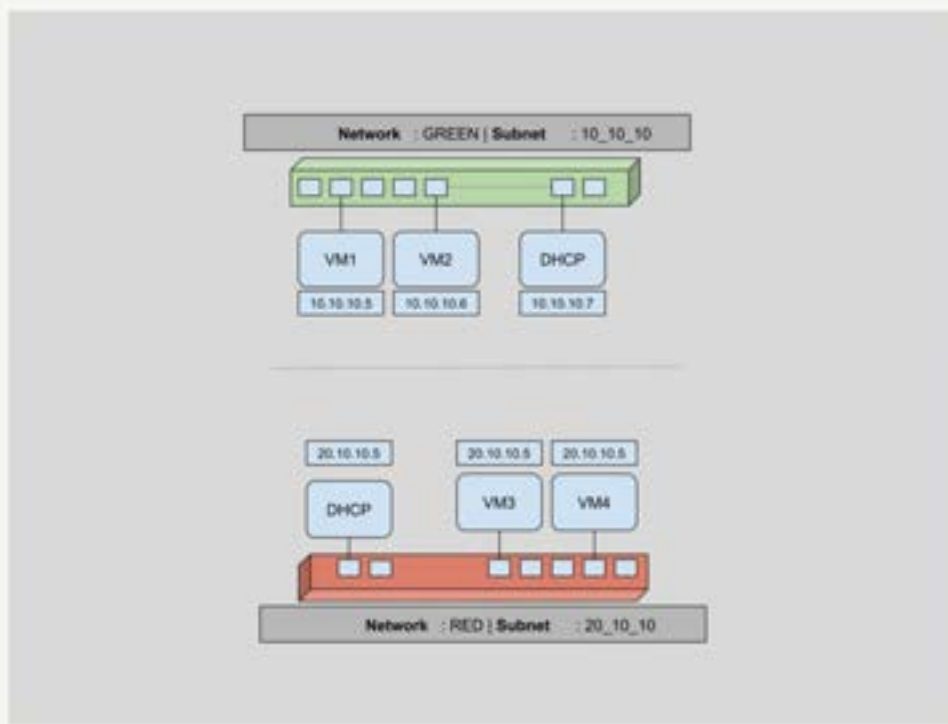
⁹⁵*Id.*

280. Nokia CloudBand uses the Neutron networking component of Openstack to receive configuration information for a virtual computer network that includes multiple computing nodes arranged via a specified network topology and configures the virtual computer network according to the configuration, as seen below.⁹⁶

Networks, Subnets & Ports

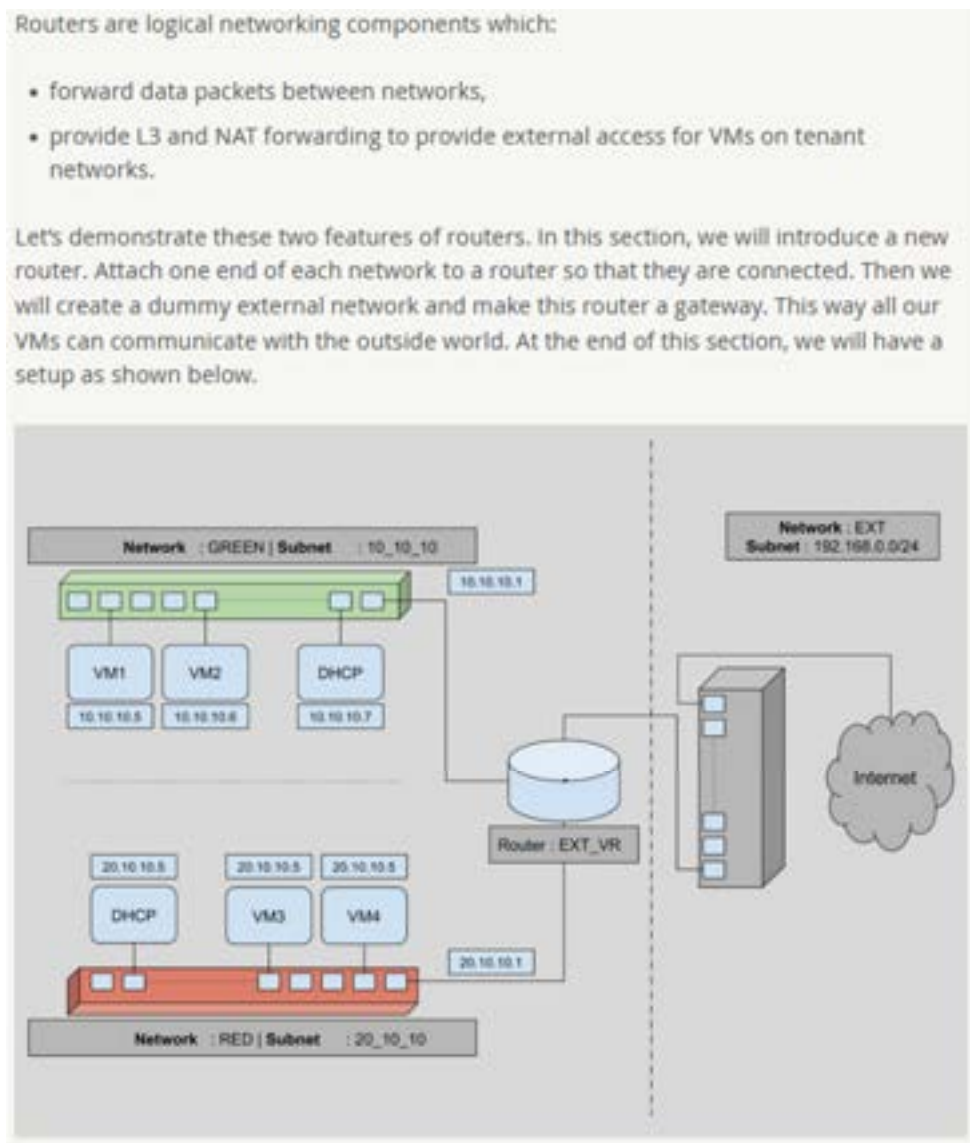
As promised, let's first create some networks. For this section our plan is to

- create a network (GREEN). Create a subnet for this network and then spawn two virtual machines (VMs) using this network.
- create another network (RED) with a different subnet and spawn another two VMs using this network
- verify ports created as part of above process and understand why OpenStack networking ports are useful.



⁹⁶Alok Kumar, Everything you need to know to get started with Neutron, Superuser (2016), <http://superuser.openstack.org/articles/everything-you-need-to-know-to-get-started-with-neutron-f90e2797-26b7-4d1c-84d8-effef03f11d2/>.

281. Nokia CloudBand uses the Neutron networking component of Openstack to receive configuration information indicating a specified virtual router device that logically interconnects groups of computing nodes, as seen below.⁹⁷



⁹⁷*Id.*

To connect any two networks we need a router. Let's create an external virtual router named EXT_VR.

```

root@controller (admin) $> neutron router-create EXT_VR
Created a new router:
+-----+-----+
| Field | Value |
+-----+-----+
| admin_state_up | True |
| distributed | False |
| external_gateway_info | |
| ha | False |
| id | <326b2b4-9076-4e12-9850-77d5f1c66c86 |
| name | EXT_VR |
| routes | |
| status | ACTIVE |
| tenant_id | 34609d0ea9ce48f98145ecc5bbac9f77 |
+-----+-----+
root@controller (admin) $> #

```

282. The Neutron networking component of Openstack allows Nokia CloudBand to create advanced virtual network topologies, as seen below.⁹⁸

Networking (neutron) concepts



UPDATED: 2019-08-23 19:39

Contents

OpenStack Networking (neutron) manages all networking facets for the Virtual Networking Infrastructure (VNI) and the access layer aspects of the Physical Networking Infrastructure (PNI) in your OpenStack environment. OpenStack Networking enables projects to create advanced virtual network topologies which may include services such as a [firewall](#), a [load balancer](#), and a [virtual private network \(VPN\)](#).

Networking provides networks, subnets, and routers as object abstractions. Each abstraction has functionality that mimics its physical counterpart: networks contain subnets, and routers route traffic between different subnets and networks.

283. Claim 4 of the '540 patent further recites “responding, by the one or more configured computing systems, to a request from one of the multiple computing nodes for

⁹⁸OpenStack, OpenStack Documentation: Networking (neutron) concepts (2019), <https://docs.openstack.org/newton/install-guide-ubuntu/neutron-concepts.html>.

information from the specified virtual router device, wherein the responding includes emulating functionality of the specified virtual router device in generating response information for the one computing node without using any physical router devices to represent the specified virtual router device.”

284. Nokia CloudBand practices this limitation. For example, Nokia CloudBand includes a Distributed Virtual Router (DVR) that implements virtual routers across the compute nodes, as seen below.⁹⁹



Neutron OVS DVR - Distributed Virtual Router

<https://blueprints.launchpad.net/neutron/+spec/neutron-ovs-dvr>

Neutron Distributed Virtual Router implements the L3 Routers across the Compute Nodes, so that tenants intra VM communication will occur without hitting the Network Node. (East-West Routing)

Also Neutron Distributed Virtual Router implements the Floating IP namespace on every Compute Node where the VMs are located. In this case the VMs with FloatingIPs can forward the traffic to the External Network without reaching the Network Node. (North-South Routing)

Neutron Distributed Virtual Router provides the legacy SNAT behavior for the default SNAT for all private VMs. SNAT service is not distributed, it is centralized and the service node will host the service.

Proposed change

The proposal is to distribute the L3 Routers across the Compute Nodes when required by the VMs.

In this case there will be Enhanced L3 Agents running on each and every compute node (This is not a new agent, this is an updated version of the existing L3 Agent). Based on the configuration in the L3 Agent.ini file, the enhanced L3 Agent will behave in legacy(centralized router) mode or as a distributed router mode.

⁹⁹OpenStack, OpenStack Documentation: Neutron OVS DVR - Distributed Virtual Router (2014), <https://specs.openstack.org/openstack/neutron-specs/specs/juno/neutron-ovs-dvr.html>.

285. Furthermore, an agent (L3 Agent) behaves as a router thereby emulating a virtual router while responding to requests from computing nodes, as seen below.¹⁰⁰

Configuration option = Default value	Description
[DEFAULT]	
<code>agent_mode = legacy</code>	(StrOpt) The working mode for the agent. Allowed modes are: 'legacy' - this preserves the existing behavior where the L3 agent is deployed on a centralized networking node to provide L3 services like DNAT, and SNAT. Use this mode if you do not want to adopt DVR. 'dvr' - this mode enables DVR functionality and must be used for an L3 agent that runs on a compute host. 'dvr_snat' - this enables centralized SNAT support in conjunction with DVR. This mode must be used for an L3 agent running on a centralized node (or in single-host deployments, e.g. devstack)

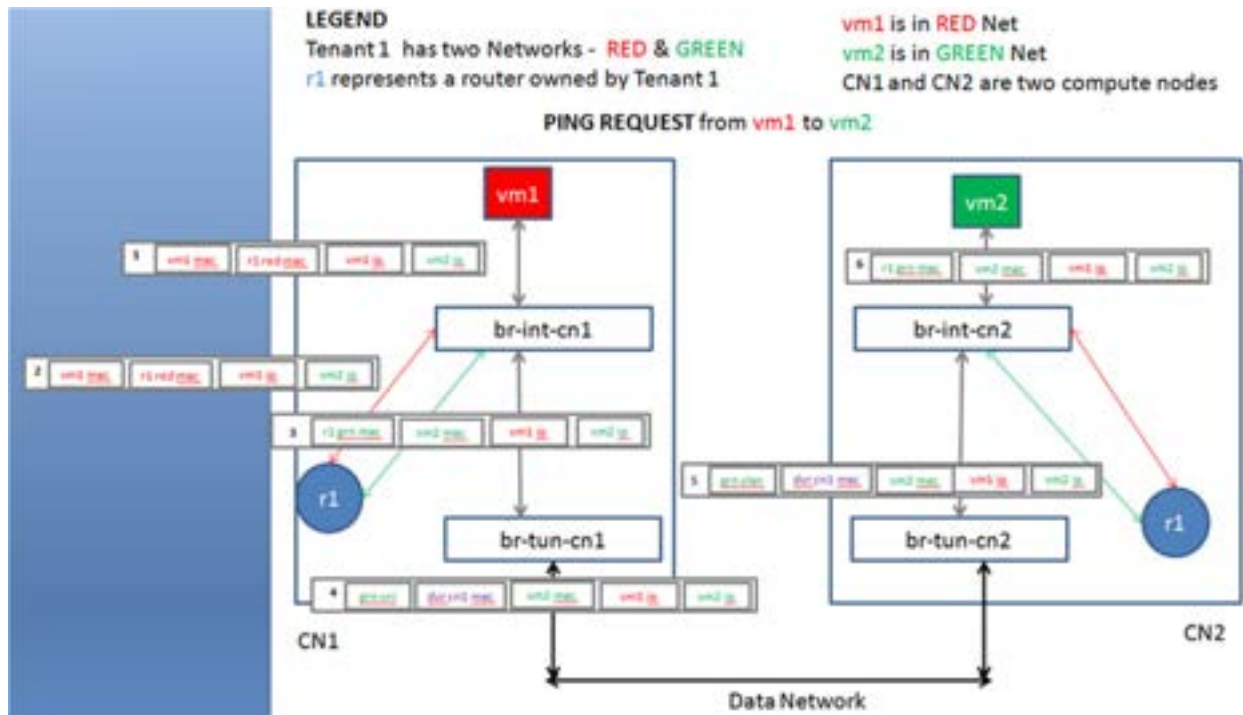
286. Claim 4 of the '540 patent further recites “forwarding, by the one or more configured computing systems, a communication between the first and second groups of computing nodes by emulating further functionality of the specified virtual router device for modifying the communication as part of the forwarding.”

287. Nokia CloudBand practices this limitation. Nokia CloudBand includes a Distributed Virtual Router (DVR) that forwards communications between groups of computing

¹⁰⁰OpenStack, OpenStack Configuration Reference: Liberty: L3 Agent, https://docs.openstack.org/liberty/config-reference/content/networking-options-l3_agent.html.

nodes by emulating the functionality of a virtual router device for modifying the communication as part of the forwarding.

288. For example, the Distributed Virtual Router (DVR) supported by Nokia CloudBand implements distributed virtual routing using L2 and L3 agents that emulate the functionality of a virtual router. “In order to accomplish distributed routing, both the L3 and L2 Agent work will need to work hand-in-hand inside the Compute Node. Today the L3 Agent runs in Network Node, but with this DVR proposal, the L3 Agent will run in Compute Nodes as well. The L2 Agent on compute nodes will continue as is today but will work in an enhanced mode called the ‘DVR Mode’ where the L2 Agents will be additionally responsible for managing (adding/removing) OVS rules in order to accomplish distributed routing,” as seen below.¹⁰¹



¹⁰¹OpenStack, Neutron/DVR L2 Agent, https://wiki.openstack.org/wiki/Neutron/DVR_L2_Agent.

289. The Distributed Virtual Router (DVR) supported by Nokia CloudBand modifies the communications as part of the forwarding. “The tunnel bridge on CN1, replaces the frame’s source mac address with a Unique DVR MAC Address of its node (one unique dvr mac address is assigned per compute node by the controller). The resulting frame is forwarded to CN2 by this tunnel bridge. Before forwarding, it also strips the local green-vlan tag and tunnels the frame with green-vni VXLAN id.” And “[t]he integration bridge on CN2, identifies the incoming frame’s source mac address is a DVR Unique MAC Address (every compute node l2-agent knows all dvr unique mac addresses used in the cloud). It then replaces the DVR Unique MAC Address with the green subnet interface MAC address and forwards the frame to vm2.”¹⁰²

290. Nokia is and has been on notice of the infringement of the ’540 patent at least as of the time Amazon filed and provided notice of this Complaint.

291. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the ’540 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, Nokia CloudBand.

292. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 1 of the ’540 patent.

293. Nokia will sell Nokia CloudBand with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the ’540 patent.

¹⁰²OpenStack, Neutron/DVR L2 Agent, [https://wiki.openstack.org/wiki/Neutron/DVR L2](https://wiki.openstack.org/wiki/Neutron/DVR_L2).

294. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts would cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

295. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '540 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia CloudBand, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '540 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT VII: PATENT INFRINGEMENT OF U.S. PATENT NO. 9,621,593

296. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

297. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 23) of the '593 patent in violation of 35 U.S.C. § 271, and will continue to do so.

298. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software, in violation of 35 U.S.C. § 271(a).

299. By way of example only, Nokia CloudBand Infrastructure Software meets all the limitations of at least independent claim 23 of the '593 patent, either literally or under the doctrine of equivalents.

300. Exemplary claim 23 of the '593 patent recites:

23. A system, comprising:

one or more computing systems having one or more processors; and;

at least one memory including instructions that, upon execution by at least one of the one or more processors, cause the system to:

receive, via an interface provided for use in configuring execution of programs by a program execution service, a request from a client that includes configuration information related to executing an indicated program; and

select one or more computing nodes of the program execution service to use for execution of the indicated program;

manage, based at least in part on the received configuration information, execution of one or more instances of the indicated program by the selected one or more computing nodes on behalf of the client.

301. Nokia infringes claims of the '593 patent, for example, claim 23 of the '593 patent since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including but not limited to Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software.

302. Claim 23 of the '593 patent begins, “A system, comprising: one or more computing systems having one or more processors; and at least one memory including instructions that, upon execution by at least one of the one or more processors.”

303. Nokia CloudBand Infrastructure Software practices this limitation. For example, Nokia AirFrame Data Center allows users to “build compact and efficient cloud computing data center” using “scalable processors,” as seen below.¹⁰³

¹⁰³Nokia, AirFrame Data Center: Adapt to Any Cloud-Based Application, <https://www.nokia.com/networks/data-center/airframe-data-center/>.

AirFrame data center solution elements

The Nokia AirFrame data center solution gives operators all the benefits and flexibility to build compact and efficient cloud computing data center that meet both Telco and IT requirements.

Nokia Data center solution is complemented with OpenNFV compatible, OpenStack and Container distribution.

Nokia data center solution provides solution from far edge to central data centers with common automation and managements system with needed workflow tools.

Nokia AirFrame Rackmount server

Nokia AirFrame Rackmount is a high-performance, energy-efficient 2-socket solution for the heavy workload data centers.

The latest generation Intel Xeon scalable processors maximize platform's speed and crucial computing performance. Utilizes the latest NVMe SSDs in a wide range of forms to minimize energy use and to deliver the fastest bandwidth, higher IOPS and lower latency.

Complete portfolio of servers, storage, networking, accelerators and commodities help communication service providers quickly move to the cloud-defined data center. With a native support for Nokia Virtual Network Functions (VNFs) and Cloud Native Network Functions (CNFs), Nokia AirFrame Rackmount is a very important building block of 5G network deployments.



Product

AirFrame Rackmount
server →

304. Claim 23 of the '593 patent further recites “receive, via an interface provided for use in configuring execution of programs by a program execution service, a request from a client that includes configuration information related to executing an indicated program.”

305. Nokia CloudBand Infrastructure Software practices this limitation. For example, Nokia CloudBand Infrastructure Software supports a “Cloud Operations Manager [that] offers single pane of glass management of hybrid, geographically distributed Kubernetes and OpenStack cloud infrastructures, including dynamic cluster creation as well as management of cluster resources,” as seen below.¹⁰⁴

¹⁰⁴Cloud Operations Manager, <https://www.nokia.com/networks/core-networks/cloudband/cloud-operations-manager/>.



What is Nokia Cloud Operations Manager?

Nokia Cloud Operations Manager (previously [Nokia CloudBand Network Director](#)) automates lifecycle operations of containerized network functions (CNFs), virtual network functions (VNFs), and network services (NSs). It optimizes and governs platform resource usage and is designed for distributed, multi-tenant, multi-vendor cloud infrastructures.

Telco applications are increasingly leveraging cloud-native principles, which will often co-exist with VNFs. Multiple integrated applications provide an inter-related service as application suites or engineered systems such as [VoLTE](#) or [core networks](#). Cloud Operations Manager automates the lifecycle management of such [hybrid networks](#) to achieve agility and cost savings.

Cloud Operations Manager provides unified lifecycle management, centralized monitoring and visualization of CNFs, VNFs and NSs. It also offers single pane of glass management of hybrid, geographically distributed Kubernetes and OpenStack cloud infrastructures, including dynamic cluster creation as well as management of cluster resources.

Today, more than 50 customers around the world rely on Cloud Operations Manager for automated cloud lifecycle operations.

On this page

- ↓ [What is Nokia Cloud Operations Manager?](#)
- ↓ [Benefits and features](#)
- ↓ [Related topics](#)
- ↓ [Resources](#)
- ↓ [Related solutions and products](#)
- ↓ [Learn more](#)

306. Nokia CloudBand Application Manager supports “virtual network functions ... referred to as VNFs” and supports use of “TOSCA template” for “VNF lifecycle management,” as seen below.¹⁰⁵

¹⁰⁵Nokia, Nokia NSP Network Services Platform, https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE15164AAAETQZZA_V1_NSP%2019.11%20Network%20Supervision%20Application%20Help.pdf, p. 28.

2.14 VNF custom actions

If there are custom actions defined in the TOSCA template for a VNF, you can trigger them using the Network Supervision application. Click on the More button on a VNF and then choose VNF > Custom Actions to view a list of any custom actions that have been defined for the VNF.

307. The TOSCA requests support “placement” of “TOSCA nodes” “by region[s],” including “[g]eographic regions (e.g., cities, municipalities, states, countries, etc.),” as seen below.¹⁰⁶

12.5.1.2 Use Case 2: Controlled placement by region

12.5.1.2.1 Description

This use case demonstrates the use of named “containers” which could represent the following:

- Datacenter regions
- Geographic regions (e.g., cities, municipalities, states, countries, etc.)
- Commercial regions (e.g., North America, Eastern Europe, Asia Pacific, etc.)

12.5.1.2.2 Features

This use case introduces the following policy features:

- Separation of resources (i.e., TOSCA nodes) by logical regions, or zones.

12.5.1.2.3 Sample YAML: Region separation amongst named set of regions

```
failover_policy_2:
  type: toasca.policy.placement
  description: My failover policy with allowed target regions (logical containers)
  properties:
    container type: region
    container_number: 3
    # If “containers” keyname is provided, they represent the allowed set
    # of target containers to use for placement for .
    containers: [ region1, region2, region3, region4 ]
```

308. Claim 23 of the ’593 patent further recites “select one or more computing nodes of the program execution service to use for execution of the indicated program.”

¹⁰⁶TOSCA Simple Profile in YAML Version 1.1, <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/csprd01/TOSCA-Simple-Profile-YAML-v1.1-csprd01.html>.

level orchestrator, which is responsible for service-specific configuration and service activation. During the lifetime of the network service, the Network Director captures network service and VNF topology, monitors the health status of the service and scales, heals and updates the service according to policies described in the network service descriptor.”¹⁰⁹ Nokia CloudBand also includes Nokia Container Services, which provides functionalities similar to those offered earlier by Nokia CloudBand Network Director. Nokia Container Services is an “offering for deploying, orchestrating, monitoring and managing containers and container-based applications” that supports “complete infrastructure life cycle management functions,” “infrastructure scaling operations which allow addition or removal of nodes from a cluster,” and “heal[ing]” to “recover a failed node.”¹¹⁰

¹⁰⁹Exhibit M, Nokia CloudBand Network Director, Product Information, <https://docplayer.net/19792521-Nokia-cloudband-network-director.html>, pp. 3-4.

¹¹⁰Nokia Container Services, https://onestore.nokia.com/asset/207821?_gl=1*_1eyhxl*_gcl_au*MTQ5MzMwNjczOC4xNzE2OTY1OTE5*_ga*MTg4NDkzNjMxMi4xNzE2OTY1OTE3*_ga_D6GE5QF247*MTcyMTE2MTkxOS4yMC4xLjE3MjExNjQ5MTIuMC4wLjA.&_ga=2.19237244.578210260.1721069959-1884936312.1716965917, p. 2.



Network service orchestrator

A network service or service chain, according to the ETSI NFV definition and the TM Forum ZOOM model, is a collection of network functions connected by a network function forwarding graph. In particular, a network service can be implemented with VNFs and a VNF forwarding graph. The Network Director supports network services as well as single VNFs. As part of network services, VNFs can be connected to other VNFs as well as physical network functions. Typical network service use cases are virtualized enterprise customer premises equipment, including service chaining, Gi-LAN, virtual evolved packet core and Voice over LTE.

The Network Director models and orchestrates network services using OASIS TOSCA descriptors and OpenStack Mistral workflows. The Network Director supports onboarding of network services and single VNFs and maintains a catalog of these. Onboarding is achieved by modeling the service as a collection of one or more VNFs connected with a forwarding graph, a set of resource requirements and a collection of workflows describing its lifecycle actions. The resulting model is documented as a cloud service archive (CSAR) and stored in the network service/VNF catalog. The archive contains a manifest file along with service templates and any workflow files and related scripts to customize the VNF or network service life cycle operations. With the Network Director domain-specific modeling language and workflow engine, virtually all network services/network functions can be onboarded without software changes to the Network Director.



Once a network service or function has been onboarded, higher level orchestrators or human operators can trigger the Network Director to deploy or update an instance of the network service as part of the service orchestration process (Figure 3). As a first step of this process, the set of VNFs and VNF forwarding graphs is determined based on a TOSCA template. Then, the Network Director identifies a generic VNF manager, such as CloudBand Application Manager, or VNF-specific VNF managers, to execute the lifecycle of the VNFs. If necessary, an appropriate VNF manager is instantiated. It then instructs the VNF manager to deploy or update and commission the VNFs and commission connected PNFs to work with the VNFs. The Network Director requests an SDN controller to establish the VNF forwarding graph between the network functions. The forwarding graph can be a collection of virtual links with or without service chaining. For Service Chaining use cases, virtual link are configured as defined in the VNF Forwarding Graph (VNFFG). When VNFs are added or removed from a service chain, the virtual links are created, modified or deleted as per the updated VNFFG. When the network service has been deployed and commissioned, control returns to the higher level orchestrator, which is responsible for service-specific configuration and service activation.

(Exhibit M, Nokia CloudBand Network Director, Product Information, <https://docplayer.net/19792521-Nokia-cloudband-network-director.html>, pp. 3-4.)

313. Nokia is and has been on notice of the infringement of the '593 patent at least as of the time Amazon filed and provided notice of this Complaint.

314. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '593 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging

others to make, use, sell, and/or offer to sell in the United States, Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software.

315. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 23 of the '593 patent.

316. Nokia will sell Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '593 patent.

317. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts will cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

318. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '593 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '593 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT VIII: PATENT INFRINGEMENT OF U.S. PATENT NO. 9,329,909

319. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

320. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 21) of the '909 patent in violation of 35 U.S.C. § 271, and will continue to do so.

321. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia AirFrame Data Center and Nokia CloudBand, in violation of 35 U.S.C. § 271(a).

322. By way of example only, Nokia AirFrame Data Center and Nokia CloudBand meets all the limitations of at least independent claim 21 of the '909 patent, either literally or under the doctrine of equivalents.

323. Exemplary claim 21 of the '909 patent recites:

21. A computing system, comprising:

one or more processors; and

one or more components of a distributed execution service that are configured to, when executed by at least one of the one or more processors, and for each of multiple programs to be executed:

select a subset of a plurality of available computing nodes to use for executing the program in a distributed manner;

initiate the executing of the program on the subset of computing nodes such that one or more jobs of the program are attempted to be executed on each of the computing nodes of the subset;

determine, while at least one of the subset of computing nodes is still executing at least one of the jobs of the program, that an actual amount of computing resources being used by the multiple computing nodes to execute the indicated program differs from an expected amount of computing resources; and

based on the determining, initiate a change in a quantity of the computing nodes of the subset being used for the executing of the program.

324. Nokia infringes claims of the '909 patent, for example, claim 21 of the '909 patent since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including but not limited to Nokia AirFrame Data Center and Nokia CloudBand.

325. For the preamble of claim 21, to the extent the preamble is determined to be limiting, Nokia AirFrame Data Center practices a “computing system, comprising: one or more processors.”

326. Nokia AirFrame Data Center practices this limitation. For example, Nokia AirFrame Data Center solution includes the necessary hardware, software and services that can adapt to any cloud-based application and supports creating scalable and distributed cloud-based architecture.¹¹¹ The mapping of this limitation of claim 21 of the '909 patent to Nokia AirFrame Data Center is exemplary, but the same infringing actions occur on configurations using other processors (or the like), as seen below.¹¹²

¹¹¹Nokia, AirFrame Data Center: Adapt to any cloud-based application, <https://networks.nokia.com/solutions/airframe-data-center-solution>.

¹¹²*Id.*

AirFrame data center

Adapt to any cloud-based application

The acceleration of telco and IT convergence and the need to support a diverse range of demanding applications requires an innovative solution that takes all the benefits from the IT and open source domains to create a scalable and distributed cloud-based architecture.

More and more services, including 5G need network functionalities and capabilities located at the most efficient point within a network. This is necessary in order to address strict latency constraints and to process huge data demands that will be critical in delivering services with real-time responsiveness.

Network architectures need a re-think, with layered and distributed network topologies containing optimized hardware to deliver unparalleled performance and the greatest flexibility.

Nokia AirFrame data center solution is designed for running demanding virtualized and cloud-native software workloads. Security of the hardware and firmware is uncompromised and fulfills increased privacy and security needs for carrier grade networks. Enhancements including advanced packet, crypto, GPU and workload-specific acceleration make AirFrame to perform better than any traditional IT servers.

327. The preamble of claim 21 further recites “one or more components of a distributed execution service that are configured to, when executed by at least one of the one or more processors, and for each of multiple programs to be executed.”

328. Nokia CloudBand practices this limitation. For example, Nokia “CloudBand supports distributed NFV [(Network Functions Virtualization)] cloud infrastructures in a variety

of ways. CloudBand provides aggregated northbound APIs that allow NFV applications and BSS/OSS to deal with the different locations as a single cloud,” as seen below.¹¹³

CloudBand supports distributed NFV cloud infrastructures in a variety of ways. CloudBand provides aggregated northbound APIs that allow NFV applications and BSS/OSS to deal with the different locations as a single cloud. CloudBand provides a policy-based placement algorithm that computes a quasi-optimal location based on server utilization at the different locations, affinity and anti-affinity rules, and other parameters. In addition, CloudBand supports Placement Zones, which span geo-distributed datacenters by “aggregating” CloudBand Node OpenStack availability zones. Virtual machines belonging to vNFs can then be instantiated across these Placement Zones according to pre-defined business policies. CloudBand provides built-in load balancing services for scalable distributed network functions. The CloudBand graphical user portal gives users an aggregated view of the infrastructure according to their different roles and responsibilities. The image management component of CloudBand manages a single catalog of images automatically and assures that images are made available where they are needed. User accounts and key pairs are also managed at the global level.

CloudBand with OpenStack as NFV Platform
COLLABORATIVE WHITE PAPER BETWEEN ALCATEL-LUCENT AND RED HAT
4

329. Claim 21 of the '909 patent further recites “select a subset of a plurality of available computing nodes to use for executing the program in a distributed manner.”

330. Nokia CloudBand practices this limitation. For example, “CloudBand provides a policy-based placement algorithm that computes a quasi-optimal location based on server utilization at the different locations, affinity and anti-affinity rules, and other parameters. In addition, CloudBand supports Placement Zones, which span geo-distributed datacenters by ‘aggregating’ CloudBand Node OpenStack availability zones. Virtual machines belonging to VNFs can then be instantiated across these Placement Zones according to pre-defined business policies. CloudBand provides built-in load balancing services for scalable distributed network functions.”¹¹⁴

¹¹³Alcatel-Lucent, Cloudband with OpenStack as NFV Platform (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>, p. 4.

¹¹⁴*Id.*

331. Claim 21 of the '909 patent further recites “initiate the executing of the program on the subset of computing nodes such that one or more jobs of the program are attempted to be executed on each of the computing nodes of the subset.”

332. Nokia CloudBand practices this limitation. For example, “CloudBand Application Manager as an ETSI NFV-compliant Virtualized Network Function Manager (VNFM) automates VNF lifecycle management and cloud resource management. It executes lifecycle management actions more easily and predictably than manual methods. It supports Nokia and third-party VNFs,” as seen below.¹¹⁵

NOKIA

CloudBand Application Manager

Get a VNFM for OpenStack- and VMware-based virtual machines

Ready-to-use VNF lifecycle management

CloudBand Application Manager as an ETSI NFV-compliant Virtualized Network Function Manager (VNFM) automates VNF lifecycle management and cloud resource management, and its standards-based APIs make it easy to work with any vendor's VNF, Element Management System (EMS), Virtualized Infrastructure Manager (VIM), and NFV Orchestrator (NFVO). For the Nokia VNF portfolio, CloudBand Application Manager offers a read-to-use, pre-integrated, one-click solution for lifecycle management.

CloudBand Application Manager automates lifecycle management by providing an open templating system, managing resources and applying associated workflows. It executes lifecycle management actions more easily and predictably than manual methods. It supports Nokia and third-party VNFs. Using OpenStack Heat orchestration templates, VMware Open Virtualization Format templates (OVF), Mistral workflows and Ansible playbooks, CloudBand Application Manager is open to a broad range of VNF onboarding options. It visualizes the structure and status of applications and performs lifecycle management, including basic functions (create, instantiate, scale, terminate, delete, operate, query and modify VNF), advanced functions (such as healing, update/patching, upgrades, backup and restore) and fault management for virtualized resources.

On this page

- ↓ VNF lifecycle management
- ↓ Resources
- ↓ Related solutions and products
- ↓ Learn more

¹¹⁵Nokia, CloudBand Application Manager, Get a VNFM for OpenStack- and VMware-based Virtual Machines, <https://www.nokia.com/networks/core-networks/cloudband/application-manager/>.

333. Claim 21 of the '909 patent further recites “determine, while at least one of the subset of computing nodes is still executing at least one of the jobs of the program, that an actual amount of computing resources being used by the multiple computing nodes to execute the indicated program differs from an expected amount of computing resources.”

334. Nokia CloudBand practices this limitation. For example, “You can assign a threshold policy to a VNF to allow the NFM-P to trigger automatic lifecycle management operations based on defined KPIs or alarms. A threshold policy allows you to monitor a set of pre-defined KPIs and create rules to define when the application indicates an overload, underload, or healing condition,” as seen below.¹¹⁶

4.1.7 VNF threshold policies

You can assign a threshold policy to a VNF to allow the NFM-P to trigger automatic lifecycle management operations based on defined KPIs or alarms. A threshold policy allows you to monitor a set of pre-defined KPIs and create rules to define when the application indicates an overload, underload, or healing condition. The policy also allows you to define an automatic triggered action to be performed when any of these conditions is met. These corrective actions include performing a scaling operation, performing a healing operation, or raising an alarm. When a lifecycle management action is triggered, the NFM-P automatically sends a lifecycle change notification to CBAM.

You can create a CMM or CMG template to define a list of conditions and specify an action to be automatically performed when those conditions are met. The template can be used to create a VNF threshold policy, but you can modify the default conditions and actions imported from the template each time you create a new policy.

Release 17.3
March 2017
Issue 1

3HE-11999-AAAA-TQZZA

19

335. Claim 21 of the '909 patent further recites “based on the determining, initiate a change in a quantity of the computing nodes of the subset being used for the executing of the program.”

¹¹⁶NSP Network Services Platform, Network Functions Manager - Packet (NFM-P) Release 17.3, NFV Solutions Guide, https://documentation.nokia.com/cgi-bin/dbaccess filename.cgi/3HE11999AAAATQZZA01_V1_NSP%20NFM-P%2017.3%20NFV%20Solutions%20Guide.pdf, p. 19.

336. Nokia CloudBand practices this limitation. For example, “The policy also allows you define an automatic triggered action to be performed when any of these conditions is met. These corrective actions include performing a scaling operation, performing a healing operation, or raising an alarm. When a lifecycle management action is triggered, the NFM-P automatically sends a lifecycle change notification to CBAM.”¹¹⁷

337. Nokia is and has been on notice of the infringement of the '909 patent at least as of the time Amazon filed and provided notice of this Complaint.

338. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '909 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, Nokia AirFrame Data Center and Nokia CloudBand.

339. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 21 of the '909 patent.

340. Nokia will sell Nokia AirFrame Data Center and Nokia CloudBand with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '909 patent.

341. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts would cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

¹¹⁷NSP Network Services Platform, Network Functions Manager – Packet (NFM-P) Release 17.3, NFV Solutions Guide, [https://documentation.nokia.com/cgi-bin/dbaccess filename.cgi/3HE11999AAAATQZZA01_V1_NSP%20NFM-P%2017.3%20NFV%20Solutions %20Guide.pdf](https://documentation.nokia.com/cgi-bin/dbaccess filename.cgi/3HE11999AAAATQZZA01_V1_NSP%20NFM-P%2017.3%20NFV%20Solutions%20Guide.pdf), p. 19.

342. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '909 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia AirFrame Data Center and Nokia CloudBand, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '909 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT IX: PATENT INFRINGEMENT OF U.S. PATENT NO. 9,766,912

343. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

344. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 10) of the '912 patent in violation of 35 U.S.C. § 271, and will continue to do so.

345. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia CloudBand Infrastructure Software in violation of 35 U.S.C. § 271(a).

346. By way of example only, Nokia CloudBand Infrastructure Software meets all the limitations of at least independent claim 10 of the '912 patent, either literally or under the doctrine of equivalents.

347. Exemplary claim 10 of the '912 patent recites:

10. One or more computer-readable storage media having instructions stored thereon for executing a method, the method comprising:

receiving a request to launch a virtual machine;

determining a virtual machine image, a virtual machine image configuration, and metadata configuration information, which are together associated with launching of the virtual machine, wherein the virtual machine image configuration identifies the virtual hardware associated with the launching of the virtual machine, wherein the virtual machine image includes software used to boot the virtual machine, and wherein the metadata configuration information is stored in a same registration record with the virtual machine image configuration, so that both the virtual machine image configuration and the metadata configuration information are identified by a single Application Programming Interface request;

launching the virtual machine using the virtual machine image and the virtual machine image configuration:

as part of the launching process, supplying the metadata configuration information to the launched virtual machine so that the virtual machine can use the metadata configuration information to customize itself after launching, the metadata configuration information identifying software packages to be loaded into the virtual machine and the virtual machine interpreting or executing the metadata configuration information to obtain the software packages to load so that the virtual machine customizes itself after launching.

348. Nokia infringes claims of the '912 patent, for example, claim 10 of the '912 patent since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including but not limited to Nokia CloudBand Infrastructure Software.

349. Claim 10 of the '912 patent begins, "One or more computer-readable storage media having instructions stored thereon for executing a method."

350. As a non-limiting example, Nokia performs this method using Nokia CloudBand Application Manager. For example, Nokia CloudBand Application Manager "makes it simple to

host, orchestrate, automate, and manage” “Virtualized Network Functions (vNFs).”¹¹⁸ Nokia CloudBand Application Manager supports “[a]utomated vNF life cycle management” that allows configuring servers that include computer-readable storage media having instructions stored thereon, as seen below.¹¹⁹

¹¹⁸CloudBand Deploy cloud infrastructure, management and orchestration for Virtual Machines and Container-based software, <https://www.nokia.com/networks/core-networks/cloudband/>.

¹¹⁹NETWORK FUNCTIONS VIRTUALIZATION – CHALLENGES AND SOLUTIONS, <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2013/9377-network-functions-virtualization-challenges-solutions.pdf>, p. 7.

AUTOMATED LIFECYCLE MANAGEMENT

In the IT world, tools such as Cloudify, Puppet and Chef² are widely used to automate server configuration. With the virtualization of the telecommunications industry, service providers have the opportunity to introduce similar tools as part of their operational arsenal.

Figure 4. Automated vNF lifecycle management

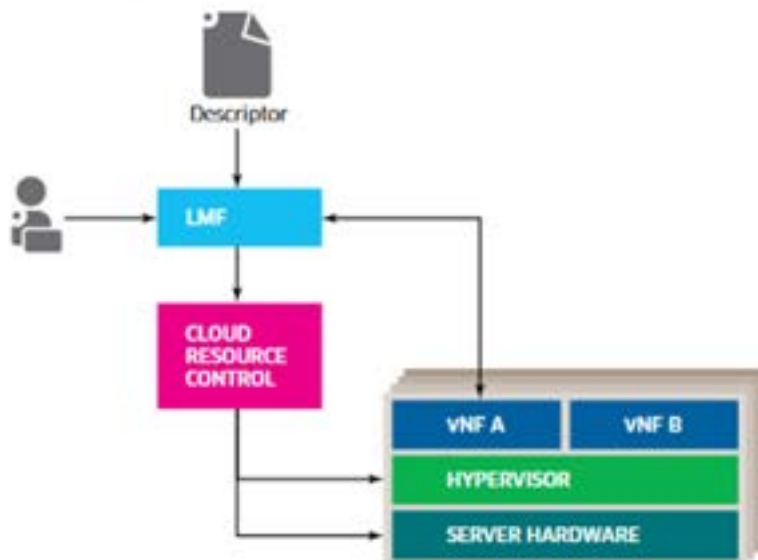


Figure 4 illustrates the role of a vNF lifecycle management function (LMF). It uses a descriptor that is generally provided by the vNF vendor. This descriptor defines the structure of the vNF (as it may consist of various sub-functions where each needs to run as an independent VM) and deployment and operational aspects, such as computation, storage and networking requirements. These descriptors are mapped to requests to the cloud resource control to create VMs and to identify software images to be downloaded and initiated on those VMs. Once the VMs are up and running, the LMF configures parameters of the vNF components based on instructions in the descriptor.

351. Claim 10 of the '912 patent further recites “receiving a request to launch a virtual machine.”

352. Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager performs “[deployment of] cloud infrastructure, management and orchestration for Virtual Machines and Container-based software.”¹²⁰

353. Claim 10 of the '912 patent further recites “determining a virtual machine image, a virtual machine image configuration, and metadata configuration information, which are together associated with launching of the virtual machine.”

354. Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager “uses a descriptor” for “virtualized network functions (vNFs)” that is “provided by the vNF vendor.”¹²¹ These “descriptors are mapped to requests to the cloud resource control to create VMs and to identify software images to be downloaded and initiated on those VMs,” as seen below.¹²²

¹²⁰Nokia, CloudBand Application Manager: Get a VNFM for OpenStack - and VMware-based Virtual Machines, <https://www.nokia.com/networks/core-networks/cloudband/application-manager/>.

¹²¹NETWORK FUNCTIONS VIRTUALIZATION – CHALLENGES AND SOLUTIONS, <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2013/9377-network-functions-virtualization-challenges-solutions.pdf>, p. 7.

¹²²*Id.*

AUTOMATED LIFECYCLE MANAGEMENT

In the IT world, tools such as Cloudify, Puppet and Chef² are widely used to automate server configuration. With the virtualization of the telecommunications industry, service providers have the opportunity to introduce similar tools as part of their operational arsenal.

Figure 4. Automated vNF lifecycle management

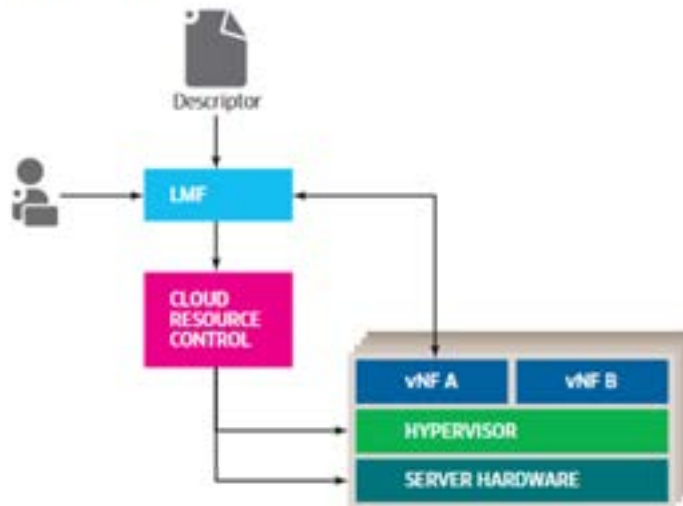


Figure 4 illustrates the role of a vNF lifecycle management function (LMF). It uses a descriptor that is generally provided by the vNF vendor. This descriptor defines the structure of the vNF (as it may consist of various sub-functions where each needs to run as an independent VM) and deployment and operational aspects, such as computation, storage and networking requirements. These descriptors are mapped to requests to the cloud resource control to create VMs and to identify software images to be downloaded and initiated on those VMs. Once the VMs are up and running, the LMF configures parameters of the vNF components based on instructions in the descriptor.

355. Claim 10 of the '912 patent further recites “wherein the virtual machine image configuration identifies the virtual hardware associated with the launching of the virtual machine.”

356. Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager supports use of “TOSCA template,” as seen below.¹²³

¹²³Nokia, Nokia NSP Network Services Platform, https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE15164AAAETQZZA_V1_NSP%2019.11%20Network%20Supervision%20Application%20Help.pdf.

2.14 VNF custom actions

If there are custom actions defined in the TOSCA template for a VNF, you can trigger them using the Network Supervision application. Click on the More button on a VNF and then choose VNF > Custom Actions to view a list of any custom actions that have been defined for the VNF.

357. The TOSCA template allows specifying “Virtual Deployment Units (VDU)” that “describe the capabilities of the virtualized containers, such as virtual CPU, RAM, disk,” as seen below.¹²⁴

3.3 VNFD: Virtualized Network Function Descriptor

A VNFD is a deployment template which describes a VNF in terms of deployment and operational behavior requirements. It also contains connectivity, interface and virtualized resource requirements [ETSI GS NFV-IFA 011]. The main parts of the VNFD are the following:

- VNF topology: it is modeled in a cloud agnostic way using virtualized containers and their connectivity. Virtual Deployment Units (VDU) describe the capabilities of the virtualized containers, such as virtual CPU, RAM, disks; their connectivity is modeled with VDU Connection Point Descriptors (VduCpd), Virtual Link Descriptors (Vld) and VNF External Connection Point Descriptors (VnfExternalCpd);
- VNF deployment aspects: they are described in one or more deployment flavours, including configurable parameters, instantiation levels, placement constraints (affinity / antiaffinity), minimum and maximum VDU instance numbers. Horizontal scaling is modeled with scaling aspects and the respective scaling levels in the deployment flavours;
- VNF lifecycle management (LCM) operations: describes the LCM operations supported per deployment flavour, and their input parameters; Note, that the actual LCM implementation resides in a different layer, namely referring to additional template artifacts.

Editor Note: VNF LCM operation modeling in TOSCA is still under discussion.

358. Claim 10 of the '912 patent further recites “wherein the virtual machine image includes software used to boot the virtual machine, and wherein the metadata configuration information is stored in a same registration record with the virtual machine image configuration.”

359. Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager allows using TOSCA for configuring virtual machine images along with metadata configuration information, as seen below.¹²⁵

¹²⁴*Id.*

¹²⁵Topology and Orchestration Specification for Cloud Applications Version 1.0, <http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.html>.

5.4.1 `tosca.artifacts.nfv.Swimage`

Shorthand Name	Swimage
Type Qualified Name	tosca:Swimage
Type URI	tosca.artifacts.nfv.Swimage
derived_from	tosca.artifacts.Deployment.Image

5.4.1.1 Properties

Name	Required	Type	Constraints	Description
name	yes	string		Name of this software image
version	yes	string		Version of this software image
checksum	yes	string		Checksum of the software image file
container_format	yes	string		The container format describes the container file format in which software image is provided.
disk_format	yes	string		The disk format of a software image is the format of the underlying disk image
min_disk	yes	scalar-unit size		The minimal disk size requirement for this software image.
min_ram	no	scalar-unit size		The minimal RAM requirement for this software image.
Size	yes	scalar-unit size		The size of this software image
sw_image	yes	string		A reference to the actual software image within VNF Package, or url.
operating_system	no	string		Identifies the operating system used in the software image.
supported_virtualization_environment	no	list		Identifies the virtualization environments (e.g. hypervisor) compatible with this software image

360. Claim 10 of the '912 patent further recites “so that both the virtual machine image configuration and the metadata configuration information are identified by a single Application Programming Interface request.”

361. Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager supports “standards-based APIs” for “VNF lifecycle management,” as seen below.¹²⁶

¹²⁶Nokia, CloudBand Application Manager: Get a VNFM for OpenStack - and VMware-based Virtual Machines, <https://www.nokia.com/networks/core-networks/cloudband/application-manager/>.



Ready-to-use VNF lifecycle management

CloudBand Application Manager as an ETSI NFV-compliant Virtualized Network Function Manager (VNFM) automates VNF lifecycle management and cloud resource management, and its standards-based APIs make it easy to work with any vendor's VNF, Element Management System (EMS), Virtualized Infrastructure Manager (VIM), and NFV Orchestrator (NFVO). For the Nokia VNF portfolio, CloudBand Application Manager offers a read-to-use, pre-integrated, one-click solution for lifecycle management.

362. Claim 10 of the '912 patent further recites “launching the virtual machine using the virtual machine image and the virtual machine image configuration.”

363. Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager allows configuring “descriptors [that] are mapped to requests to the cloud resource control to create VMs and to identify software images to be downloaded and initiated on those VMs. Once the VMs are up and running, the LMF configures parameters of the vNF components based on instructions in the descriptors,” as seen below.¹²⁷

¹²⁷NETWORK FUNCTIONS VIRTUALIZATION – CHALLENGES AND SOLUTIONS, <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2013/9377-network-functions-virtualization-challenges-solutions.pdf>.

AUTOMATED LIFECYCLE MANAGEMENT

In the IT world, tools such as Cloudify, Puppet and Chef[®] are widely used to automate server configuration. With the virtualization of the telecommunications industry, service providers have the opportunity to introduce similar tools as part of their operational arsenal.

Figure 4. Automated vNF lifecycle management

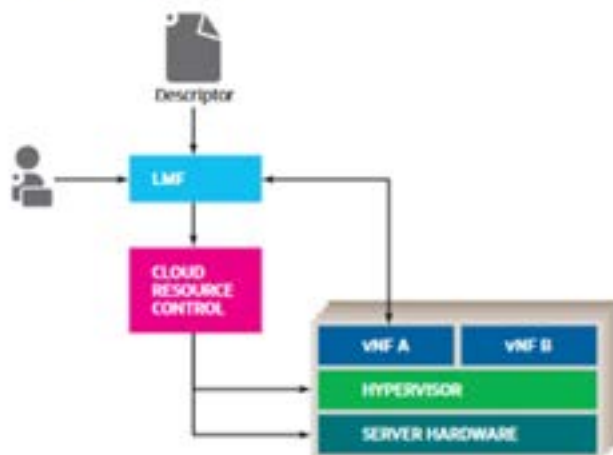


Figure 4 illustrates the role of a vNF lifecycle management function (LMF). It uses a descriptor that is generally provided by the vNF vendor. This descriptor defines the structure of the vNF (as it may consist of various sub-functions where each needs to run as an independent VM) and deployment and operational aspects, such as computation, storage and networking requirements. These descriptors are mapped to requests to the cloud resource control to create VMs and to identify software images to be downloaded and initiated on those VMs. Once the VMs are up and running, the LMF configures parameters of the vNF components based on instructions in the descriptor.

364. Claim 10 of the '912 patent further recites “as part of the launching process, supplying the metadata configuration information to the launched virtual machine so that the virtual machine can use the metadata configuration information to customize itself after launching, the metadata configuration information identifying software packages to be loaded into the virtual machine and the virtual machine interpreting or executing the metadata configuration information to obtain the software packages to load so that the virtual machine customizes itself after launching.”

365. Nokia CloudBand Application Manager practices this limitation. For example, Nokia CloudBand Application Manager uses “descriptors [that] are mapped to requests to the

cloud resource control to create VMs and to identify software images to be downloaded and initiated on those VMs. Once the VMs are up and running, the LMF configures parameters of the vNF components based on instructions in the descriptors,” as seen below.¹²⁸

AUTOMATED LIFECYCLE MANAGEMENT

In the IT world, tools such as Cloudify, Puppet and Chef² are widely used to automate server configuration. With the virtualization of the telecommunications industry, service providers have the opportunity to introduce similar tools as part of their operational arsenal.

Figure 4. Automated vNF lifecycle management

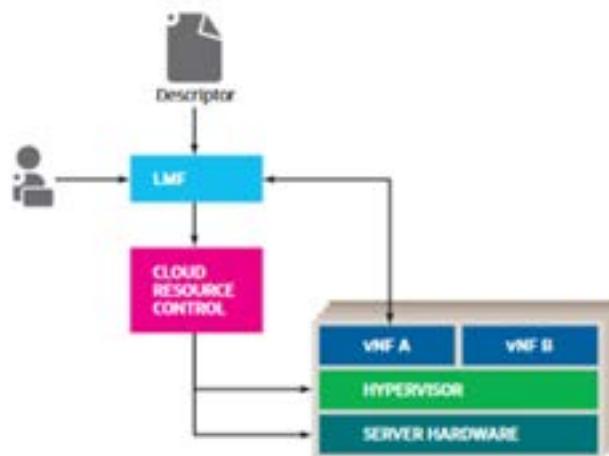


Figure 4 illustrates the role of a vNF lifecycle management function (LMF). It uses a descriptor that is generally provided by the vNF vendor. This descriptor defines the structure of the vNF (as it may consist of various sub-functions where each needs to run as an independent VM) and deployment and operational aspects, such as computation, storage and networking requirements. These descriptors are mapped to requests to the cloud resource control to create VMs and to identify software images to be downloaded and initiated on those VMs. Once the VMs are up and running, the LMF configures parameters of the vNF components based on instructions in the descriptor.

366. Nokia is and has been on notice of the infringement of the '912 patent at least as of the time Amazon filed and provided notice of this Complaint.

367. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '912 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging

¹²⁸*Id.*

others to make, use, sell, and/or offer to sell in the United States, Nokia CloudBand Application Manager.

368. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 10 of the '912 patent.

369. Nokia will sell Nokia CloudBand Application Manager with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '912 patent.

370. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts would cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

371. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '912 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia CloudBand, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '912 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT X: PATENT INFRINGEMENT OF U.S. PATENT NO. 9,756,018

372. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

373. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 18) of the '018 patent in violation of 35 U.S.C. § 271, and will continue to do so.

374. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia AirFrame Data Center and CloudBand Infrastructure Software, in violation of 35 U.S.C. § 271(a).

375. By way of example only, Nokia AirFrame Data Center and CloudBand Infrastructure Software meet all the limitations of at least independent claim 18 of the '018 patent, either literally or under the doctrine of equivalents.

376. Exemplary claim 18 of the '018 patent recites:

18. A system comprising:

one or more hardware processors of one or more computing systems; and

one or more memories with stored instructions that, when executed by at least one of the one or more hardware processors, cause the one or more computing systems to implement functionality for a network service that provides computer networks to remote clients, the implementing of the functionality including:

providing a first computer network for a first client of the network service based on information specified by the first client, wherein the providing of the first computer network for the first client includes:

selecting multiple computing nodes from a plurality of computing nodes provided by the network service for use by remote clients;

provisioning the multiple computing nodes for use in the first computer network for the client; and

configuring one or more hardware devices of the network service that provide one or more of the multiple computing nodes to route communications according to a network topology of the first computer network that is indicated in the information specified by the first client;

receiving, from the first client via an interface of the network service, a request for a secure connection between a first remote location and the first computer network provided by the network service for the first client; and

responding to the received request by providing configuration information that enables one or more devices at the first remote location to participate in the secure connection.

377. Nokia infringes claims of the '018 patent, for example, claim 18 of the '018 patent since it makes, uses, offers for sale, and/or sells in the United States and/or imports into the United States infringing products and services including but not limited to Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software.

378. Claim 18 of the '018 patent begins, "A system comprising: one or more hardware processors of one or more computing systems; and one or more memories with stored instructions that, when executed by at least one of the one or more hardware processors, cause the one or more computing systems to implement functionality for a network service that provides computer networks to remote clients, the implementing of the functionality including."

379. Nokia AirFrame Data Center practices this limitation. For example, Nokia AirFrame Data Center includes the necessary hardware, software and services that can adapt to any cloud-based application and supports creating scalable and distributed cloud-based architecture.¹²⁹ The mapping of this limitation to Nokia AirFrame Data Center is exemplary. The same infringing actions may occur on configurations using other processors or data centers, as seen below.¹³⁰

¹²⁹Nokia, AirFrame Data Center: Adapt to any cloud-based application, <https://networks.nokia.com/solutions/airframe-data-center-solution>.

¹³⁰*Id.*

AirFrame data center

Adapt to any cloud-based application

The acceleration of telco and IT convergence and the need to support a diverse range of demanding applications requires an innovative solution that takes all the benefits from the IT and open source domains to create a scalable and distributed cloud-based architecture.

More and more services, including 5G need network functionalities and capabilities located at the most efficient point within a network. This is necessary in order to address strict latency constraints and to process huge data demands that will be critical in delivering services with real-time responsiveness.

Network architectures need a re-think, with layered and distributed network topologies containing optimized hardware to deliver unparalleled performance and the greatest flexibility.

Nokia AirFrame data center solution is designed for running demanding virtualized and cloud-native software workloads. Security of the hardware and firmware is uncompromised and fulfills increased privacy and security needs for carrier grade networks. Enhancements including advanced packet, crypto, GPU and workload-specific acceleration make AirFrame to perform better than any traditional IT servers.

380. Claim 18 of the '018 patent further recites “providing a first computer network for a first client of the network service based on information specified by the first client, wherein the providing of the first computer network for the first client includes.”

381. Nokia CloudBand Infrastructure Software practices this limitation. For example, “CloudBand Infrastructure Software (CBIS) virtualizes and manages compute, storage and network resources to enable Virtualized Network Functions (VNFs) to run while meeting strict requirements for robustness, performance and security.”¹³¹ Furthermore, “CBIS performs the

¹³¹ Nokia CloudBand Infrastructure Software Data sheet, <https://onestore.nokia.com/asset/200058>, p. 1.

ETSI MANO Virtual Infrastructure Manager (VIM) function for VMs. CBIS provides the virtualization software (hypervisor, virtual switch, monitoring and Bare Metal) installed on each server,” as seen below.¹³²



Nokia CloudBand Infrastructure Software

Data sheet

CloudBand Infrastructure Software is a ready-to-use software solution geared towards customers who seek an easy way to build a serviceable, highly available, secure and protected virtual infrastructure that supports NFV, is simple to operate and is easily upgraded.

CloudBand Infrastructure Software (CBIS) virtualizes and manages compute, storage and network resources to enable Virtualized Network Functions (VNFs) to run while meeting strict requirements for robustness, performance and security. Each CBIS instance manages one NFVI node composed of VNFs distributed across one or more HW elements in one or more datacentre racks, typically in a single geographical location. CBIS performs the ETSI MANO Virtual Infrastructure Manager (VIM) function for VMs. CBIS provides the virtualization software (hypervisor, virtual switch, monitoring and Bare Metal) installed on each server.

CloudBand Infrastructure Software virtualizes storage resources across servers, supporting Ceph block storage, file storage and object storage, as well as one or more external storage arrays. It offers a rich set of features, including Enhanced Platform Awareness (EPA) for high performance VNFs; robust security solutions; flexible support of hardware and rich monitoring capabilities.

The platform is easy to install operate, and manage life cycle with automated procedures, from the CBIS Manager, an intuitive graphical user interface. In addition to its function as Infrastructure life cycle management tool, the CBIS Manager is also a launchpad to other OpenStack interfaces.

382. Furthermore, “The Network Director models and orchestrates network services using OASIS descriptors and OpenStack Mistral workflows.”¹³³ Nokia CloudBand also includes Nokia Container Services, which provides functionalities similar to those offered earlier by Nokia

¹³²*Id.*

¹³³Exhibit M, Nokia CloudBand Network Director, Product Information, <https://docplayer.net/19792521-Nokia-cloudband-network-director.html>, p. 4.

CloudBand Network Director. Nokia Container Services is an “offering for deploying, orchestrating, monitoring and managing containers and container-based applications” that supports “complete infrastructure life cycle management functions,” “infrastructure scaling operations which allow addition or removal of nodes from a cluster,” and “heal[ing]” to “recover a failed node.”¹³⁴ The Nokia Container Services “can support hundreds of worker/ edge node Virtual Machines (VMs)” and “uses Terraform providers to interface to the APIs of the virtual infrastructure.”¹³⁵

383. Claim 18 of the '018 patent further recites “selecting multiple computing nodes from a plurality of computing nodes provided by the network service for use by remote clients.”

384. Nokia CloudBand Infrastructure Software practices this limitation. For example, “CloudBand provides a policy-based placement algorithm that computes a quasi-optimal location based on server utilization at the different locations, affinity and anti-affinity rules, and other parameters. In addition, CloudBand supports Placement Zones, which span geo-distributed datacenters by ‘aggregating’ CloudBand Node OpenStack availability zones. Virtual machines belonging to vNFs can then be instantiated across these Placement Zones according to pre-defined business policies. CloudBand provides built-in load balancing services for scalable distributed network functions.”¹³⁶

¹³⁴Nokia Container Services, https://onestore.nokia.com/asset/207821?_gl=1*_eyhxlB*_gcl_au*MTQ5MzMwNjczOC4xNzE2OTY1OTE5*_ga*MTg4NDkzNjMxMi4xNzE2OTY1OTE3*_ga_D6GE5QF247*MTcyMTE2MTkxOS4yMC4xLjE3MjExNjQ5MTIuMC4wLjA.&_ga=2.19237244.578210260.1721069959-1884936312.1716965917, p. 2.

¹³⁵*Id.*

¹³⁶Alcatel-Lucent, Cloudband with Openstack as NFV Platform (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>, p. 4.

385. Claim 18 of the '018 patent further recites “provisioning the multiple computing nodes for use in the first computer network for the client.”

386. Nokia CloudBand Infrastructure Software practices this limitation. For example, “If necessary, an appropriate VNF manager is instantiated. It then instructs the VNF manager to deploy or update and commission the VNFs and commission connected PNFs to work with the VNFs. The Network Director requests an SDN controller to establish the VNF forwarding graph between the network functions. The forwarding graph can be a collection of virtual links with or without service chaining. For Service Chaining use cases, virtual links are configured as defined in the VNF Forwarding Graph (VNFFG). When VNFs are added or removed from a service chain, the virtual links are created, modified or deleted as per the updated VNFFG.”¹³⁷ Nokia CloudBand also includes Nokia Container Services, which provides functionalities similar to those offered earlier by Nokia CloudBand Network Director. Nokia Container Services is an “offering for deploying, orchestrating, monitoring and managing containers and container-based applications” that supports “complete infrastructure life cycle management functions,” “infrastructure scaling operations which allow addition or removal of nodes from a cluster,” and “heal[ing]” to “recover a failed node.”¹³⁸ Nokia Container Services “can support hundreds of worker/ edge node Virtual Machines (VMs)” and “uses Terraform providers to interface to the APIs of the virtual infrastructure.”¹³⁹

¹³⁷Exhibit M, Nokia CloudBand Network Director, Product Information, <https://docplayer.net/19792521-Nokia-cloudband-network-director.html>, p.4.

¹³⁸Nokia Container Services, https://onestore.nokia.com/asset/207821?_gl=1*_1eyhxl*_gcl_au*MTQ5MzMwNjc2OC4xNzE2OTE5*_ga*MTg4NDkzNjMxMi4xNzE2OTE5OTE3*_ga_D6GE5QF247*MTcyMTE2MTkxOS4yMC4xLjE3MjExNjQ5MTIuMC4wLjA.&_ga=2.19237244.578210260.1721069959-1884936312.1716965917, p. 2.

¹³⁹*Id.*

387. Claim 18 of the '018 patent further recites “configuring one or more hardware devices of the network service that provide one or more of the multiple computing nodes to route communications according to a network topology of the first computer network that is indicated in the information specified by the first client.”

388. Nokia CloudBand Infrastructure Software practices this limitation. For example, “As part of network services, VNFs can be connected to other VNFs as well as physical network functions,” as seen below.¹⁴⁰

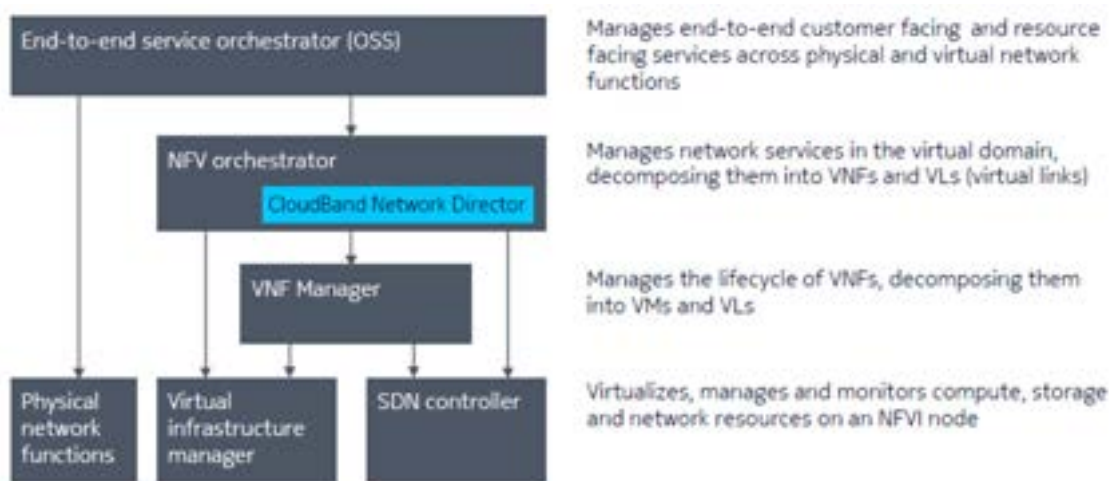


Figure 1: Hierarchy of orchestrators

389. Claim 18 of the '018 patent further recites “receiving, from the first client via an interface of the network service, a request for a secure connection between a first remote location and the first computer network provided by the network service for the first client.”

390. Nokia CloudBand Infrastructure Software practices this limitation. For example, “CloudBand is open to interface with any networking framework using standard OpenStack

¹⁴⁰*Id.*

Neutron APIs and plugins. CloudBand can interface with existing MPLS networks and other legacy networks using a process called VPN stitching,” as seen below.¹⁴¹

OpenStack Neutron is the networking component offering abstractions, such as Layer 2 and Layer 3 networks, subnets, IP addresses, and virtual middleboxes. Neutron has a plugin-based architecture. Networking requests to Neutron are forwarded to the Neutron plugin installed to handle the specifics of the present network. Neutron is limited to a single space of network resources typically associated with an OpenStack region. There is no capability to federate multiple network domains and manage WAN capabilities.

CloudBand delivers NFV networking abstractions that extend from the datacenter across the WAN toward multiple locations (Figure 4). CloudBand Route Domains and Network Templates support flexible aggregations of networks. CloudBand is open to interface with any networking framework using standard OpenStack Neutron APIs and plugins. CloudBand can interface with existing MPLS networks and other legacy networks using a process called VPN stitching.

For example, CloudBand supports Nuage Networks Virtualized Services Platform (VSP), a second generation SDN solution. Nuage Networks VSP supports the federation of multiple SDN network domains each with their own SDN controller. Networks are synchronized in both directions between CloudBand and Nuage Networks VSP. The SDN controller, in this case Nuage Networks VSP, has been recently identified by ETSI NFV Industry Specification Group and is now referenced as the WAN Infrastructure Manager (WIM).

391. Furthermore, “VPNaaS (VPN-as-a-Service) is a Neutron extension that introduces VPN feature set” and “support[s] multiple tunneling, security protocols that supports both static and dynamic routing.”¹⁴² The OpenStack Neutron APIs allow creating a secure request, as seen below.¹⁴³

¹⁴¹Alcatel-Lucent, Cloudband with OpenStack as NFV Platform (2014), <https://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10694-cloudband-with-openstack-as-nfv-platform.pdf>, p. 5.

¹⁴²Neutron/VPNaaS <https://wiki.openstack.org/wiki/Neutron/VPNaaS>.

¹⁴³*Id.*

ipsec-site-connection Create

JSON Request

```
#!highlight javascript numbers=disable
POST /v1.0/ipsec-site-connections
Accept: application/json
Content-Type: application/json
X-Auth-Token:xyz
Content-Length: abc

{
  "ipsec_site_connection" : {
    "name": "ipsec_connection_1",
    "peer-address": "192.168.2.255",
    "peer-id" : "192.168.2.255",
    "peer-cidr" : "10.30.2.0/24",
    "dpd": "action=hold,interval=20,timeout=120",
    "mtu": "1500",
    "psk": "bla_bla_bla",
    "initiator": "bi-directional",
    "vpnservice-id": "02b1fef7-16f5-4917-bf19-c40a9af805ed",
    "ikepolicy-id": "03299abc-16f5-4917-bf19-c40a9af805ed",
    "ipsecpolicy-id": "0dbc1234-16f5-4917-bf19-c40a9af805ed"
  }
}
```

392. Claim 18 of the '018 patent further recites “responding to the received request by providing configuration information that enables one or more devices at the first remote location to participate in the secure connection.”

393. Nokia CloudBand Infrastructure Software practices this limitation. For example, OpenStack Neutron APIs specifies the response format, as seen below.¹⁴⁴

¹⁴⁴ *Id.*

JSON Response

```

#highlight javascript numbers=disable
HTTP/1.1 202 Accepted
Content-Type: application/json
Content-Length: abc

{
  "ipsec_site_connection" : {
    "id": "cfc6589d-f949-4c66-99d2-c2da56ef3764",
    "tenant-id": "310df60f-2a10-4ee5-9554-98393092194c",
    "name": "ipsec_connection_1",
    "peer-address": "192.168.2.255",
    "peer-id" : "192.168.2.255",
    "peer-cidr" : "10.30.2.0/24",
    "dpd": {
      "action" : "hold"
      "interval" : 20,
      "timeout" : 120,
    }
    "mtu": "1500",
    "psk": "bla_bla_bla",
    "initiator": "bi-directional",
    "vpnservice-id": "02b1fef7-16f5-4917-bf19-c40a9af805ed",
    "ikepolicy-id": "03299abc-16f5-4917-bf19-c40a9af805ed",
    "ipsecpolicy-id": "0dbc1234-16f5-4917-bf19-c40a9af805ed",
    "admin_state_up": true,
    "status": "PENDING_CREATE"
  }
}

```

394. Nokia is and has been on notice of the infringement of the '018 patent at least as of the time Amazon filed and provided notice of this Complaint.

395. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '018 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software.

396. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 18 of the '018 patent.

397. Nokia will sell Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '018 patent.

398. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts would cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

399. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '018 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nokia AirFrame Data Center and Nokia CloudBand Infrastructure Software, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '018 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT XI: PATENT INFRINGEMENT OF U.S. PATENT NO. 11,909,586

400. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

401. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 1) of the '586 patent in violation of 35 U.S.C. § 271, and will continue to do so.

402. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing

within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia Nuage Networks Virtualized Cloud Services, in violation of 35 U.S.C. § 271(a).

403. By way of example only, Nuage Networks Virtualized Cloud Services meets all the limitations of at least independent claim 1 of the '586 patent, either literally or under the doctrine of equivalents.

404. Exemplary claim 1 of the '586 patent recites:

1. A method, comprising:

performing, by one or more computing systems of a telecommunications infrastructure provider:

managing, by a communication manager, communications to and from a computing node in a virtual computer network of computing nodes,

wherein the computing nodes are virtual machine instances hosted on physical hosts in a substrate network of the telecommunications infrastructure provider;

wherein the communication manager is implemented on a first physical host of the physical hosts and includes a switch that physically connects to the substrate network;

wherein the managing comprises:

storing configuration information about the virtual computer network;

receiving, from the substrate network, a first communication addressed to the computing node hosted on the first physical host; based at least in part on the configuration information,

modifying the first communication and forwarding the first communication to the computing node;

receiving, from the substrate network, a second communication addressed to the computing node; and

based at least in part on the configuration information, dropping the second communication without forwarding the second communication to the computing node, and

wherein the managing further comprises:

in response to an address resolution protocol (ARP) communication from the computing node regarding a second computing node in the virtual computer network, sending a spoofed response to the ARP communication indicating a virtual hardware address of the second computing node.

405. For the preamble of claim 1, to the extent the preamble is determined to be limiting, Nuage Networks Virtualized Cloud Services practices a “method, comprising: performing, by one or more computing systems of a telecommunications infrastructure provider.”

406. For example, “Nuage Networks Virtualized Cloud Services (VCS) is a Software-Defined Networking (SDN) solution that provides network virtualization and advanced automation across any Telco Cloud data center infrastructure,” as seen below.¹⁴⁵

¹⁴⁵NOKIA Accelerate Telco Clouds with Nuage Networks Virtualized Cloud Services, https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nokia_Accelerate_Telco_Clouds_with_Nuage_Networks_VCS_Application_Note_EN.pdf, p. 8.



Nuage Networks VCS

VCS for Telco Cloud in 7 points:

Provides support for all major cloud management systems, hypervisors, and network gear. VCS leverages VMs on any x86-based hardware

Uses programmable business logic and policies to fully automate and simplify network service creation

Provides support for L2-L4 services including programmable security services

Optimizes and scales Telco Cloud connectivity and is deployable on heterogeneous networks

Offers unrestricted placement of VM, container or bare metal workloads to maximize efficiency of server resources

Includes extensive data analytics and performance monitoring capabilities

Integrates public, private and hybrid cloud applications into managed VPNs

Nuage Networks Virtualized Cloud Services (VCS) is a Software-Defined Networking (SDN) solution that provides network virtualization and advanced automation across any Telco Cloud data center infrastructure and automatically establishes connectivity between virtualized compute resources whether virtual machines, containers, or legacy bare metal servers upon their creation ensuring application traffic is served by all functions reliably and efficiently. Leveraging programmable business logic and a powerful policy engine, VCS provides an open and highly responsive solution that scales to meet the stringent needs of massive multi-tenant Telco Clouds. VCS is a software solution that can be deployed over any existing datacenter network environment.

VCS network abstraction

Nuage Networks VCS allows enterprise administrators to define their networking requirements in application terms, without being burdened or slowed down by network implementation details. CSPs can express security requirements (e.g. firewall and ACL policies),

SLA requirements, load balancing, user access-rights, and more in an intuitive and abstracted way that is translated to network policies which are further translated to network configurable parameters.

407. Claim 1 of the '586 patent further recites “managing, by a communication manager, communications to and from a computing node in a virtual computer network of computing nodes.”

408. Nuage Networks Virtualized Cloud Services practices this limitation. For example, “Nuage Networks Virtualized Cloud Services (VCS) is a Software-Defined Networking (SDN) solution that provides network virtualization and advanced automation across any Telco Cloud data center infrastructure and automatically establishes connectivity between virtualized compute resources whether virtual machines, containers, or legacy bare metal servers upon their creation ensuring application traffic is served by all functions reliably and efficiently.”¹⁴⁶

¹⁴⁶*Id.*

409. Claim 1 of the '586 patent further recites “wherein the computing nodes are virtual machine instances hosted on physical hosts in a substrate network of the telecommunications infrastructure provider.”

410. Nuage Networks Virtualized Cloud Services practices this limitation. For example, Nuage Networks Virtualized Cloud Services includes “Virtual Routing and Switching (VRS)” which is “a software module that is installed in the hypervisor layer for VMs or as part of a container structure in virtualized server environments” and “creates and manages the virtual endpoints (i.e. VXLAN Tunnel Endpoint (VTEP)) that are used for the virtual ‘overlay’ tunnels between the workloads (e.g. VMs and containers) within a prescribed overlay VPN,” as seen below.¹⁴⁷

The Nokia logo is displayed in a bold, blue, sans-serif font.

Virtual Routing and Switching

Virtual Routing and Switching (VRS) is a software module that is installed in the hypervisor layer for VMs or as part of a container structure in virtualized server environments. It creates and manages the virtual endpoints (i.e. VXLAN Tunnel Endpoint (VTEP)) that are used for the virtual “overlay” tunnels between the workloads (e.g. VMs and containers) within a prescribed overlay VPN. These tunnels are created by adding a VXLAN encapsulation to the original Ethernet frame from designated traffic allowing it to be directly routable to other VMs or containers within that overlay VPN.

411. Claim 1 of the '586 patent further recites “wherein the communication manager is implemented on a first physical host of the physical hosts and includes a switch that physically connects to the substrate network.”

412. Nuage Networks Virtualized Cloud Services practices this limitation. For example, Nuage Networks Virtualized Cloud Services performs “Virtual Routing and Switching (VRS)”

¹⁴⁷*Id.* at p. 12.

based on the “based on the Open vSwitch (OVS)” which uses “SmartNIC” to “to leverage [] resident switching and packet processing capabilities for OVS offloading,” as seen below.¹⁴⁸



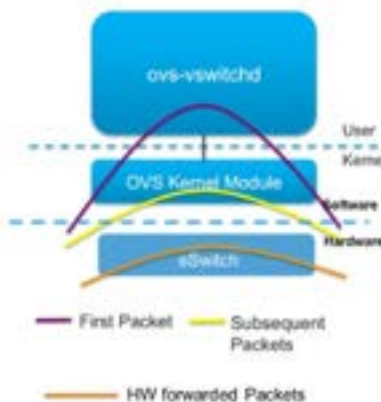
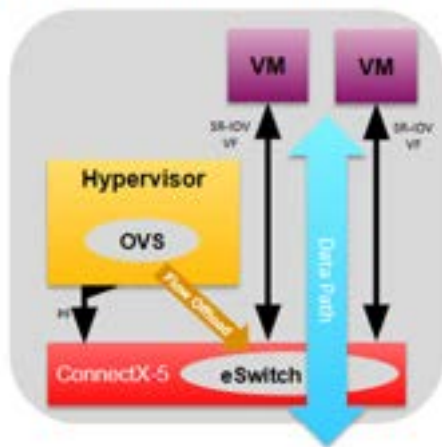
OVRs - OVS offload using SmartNIC technology

In an OVRs environment, Nuage Networks has partnered with SmartNIC vendors to leverage their resident switching and packet processing capabilities for OVS offloading. This solution combines the performance and efficiency of networking hardware on the SmartNIC with the flexibility of VRS all while leveraging the Single-Root Input/Output Virtualization (SR-IOV) standard. This offload approach scales performance linearly with the number of NICs to achieve some of the highest packet performance results. This approach is ideally suited for the most stringent and demanding applications such as the infrastructure components of the 5G packet core.

413. The “OVS Hardware” “offload[s] to NIC embedded Switch (eSwitch),” as seen below.¹⁴⁹

Full OVS Hardware Offload – NIC Architecture

- **Accelerated Switching and Packet Processing (ASAP²)**
 - Open vSwitch as Standard SDN Control
 - OVS data-plane offload to NIC embedded Switch (eSwitch) – SR-IOV Data Path
- **Best of Both Worlds: SDN Programmability and Faster Switching Performance**



¹⁴⁸Id. at p. 15.

¹⁴⁹Look Mom – No Patches in our Blazing Fast and Smart Telco Cloud (2018), <https://object-storage-ca-ymq-1.vexxhost.net/swift/v1/6e4619c416ff4bd19e1c087f27a43eea/www-assets-prod/presentation-media/No-Patches-OSS-Vancouver-v0.8.pdf>, p. 5.

414. Claim 1 of the '586 patent further recites “wherein the managing comprises: storing configuration information about the virtual computer network.”

415. Nuage Networks Virtualized Cloud Services practices this limitation. For example, the OVS “Open vSwitch” used by the Nuage Networks Virtualized Cloud Services uses “OVS using traffic control” or “OVS-TC,” an “extension of OVS” that “allows matching on a variety of predefined flow keys” that represent configuration information that allows users to “match on IP addresses, UDP/TCP ports, metadata and more,” as seen below.¹⁵⁰

OVS-TC

TC Flower is a packet classifier in the Linux kernel and part of the kernel traffic classification subsystem. OVS-TC is the extension of OVS user space code to enable it to offload flows using TC Flower and TC actions. OVS-TC allows matching on a variety of predefined flow keys. The user can match on IP addresses, UDP/TCP ports, metadata and more. Like OVS, OVS-TC includes an action side which allows packets to be modified, forwarded or dropped. The user can influence what is offloaded or not down to a per flow basis.

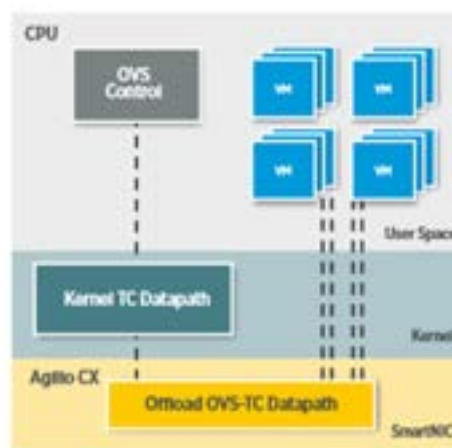


Figure 4: OVS-TC datapath

416. Claim 1 of the '586 patent further recites “receiving, from the substrate network, a first communication addressed to the computing node hosted on the first physical host; based at least in part on the configuration information.”

417. Nuage Networks Virtualized Cloud Services practices this limitation. For example, the Open vSwitch “implement[s] the OVS datapath in user space,” in which “the DPDK poll mode

¹⁵⁰Virtual Switch Acceleration with OVS-TC and Agilio 40GbE SmartNICs, https://d3ncevyc0dfnh8.cloudfront.net/media/documents/WP_OVS-TC_40G.pdf.

driver delivers packets”, *e.g.*, first communication, “directly into the dedicated user space application, bypassing the Linux kernel stack altogether,” as seen below.¹⁵¹

OVS-DPDK

OVS-DPDK is a noteworthy attempt to address the fundamental limitations of Kernel-OVS. By implementing the OVS datapath in user space, the DPDK poll mode driver delivers packets directly into the dedicated user space application, bypassing the Linux kernel stack altogether.

This eliminates unnecessary overhead in the stack and can enable additional optimizations for the vSwitch, such as loading packets directly into caches and batch processing. The DPDK community regularly provides further optimizations and tuning for OVS with upstreamed DPDK setup for network interface cards (NICs).

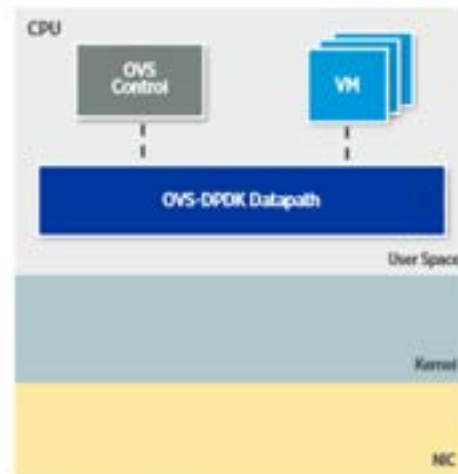


Figure 2: OVS-DPDK datapath

418. Claim 1 of the '586 patent further recites “modifying the first communication and forwarding the first communication to the computing node.”

419. Nuage Networks Virtualized Cloud Services practices this limitation. For example, “OVS-TC includes an action side which allows packets to be modified, forwarded or dropped,” as seen below.¹⁵²

¹⁵¹*Id.*

¹⁵²*Id.*

OVS-TC

TC Flower is a packet classifier in the Linux kernel and part of the kernel traffic classification subsystem. OVS-TC is the extension of OVS user space code to enable it to offload flows using TC Flower and TC actions. OVS-TC allows matching on a variety of predefined flow keys. The user can match on IP addresses, UDP/TCP ports, metadata and more. Like OVS, OVS-TC includes an action side which allows packets to be modified, forwarded or dropped. The user can influence what is offloaded or not down to a per flow basis.

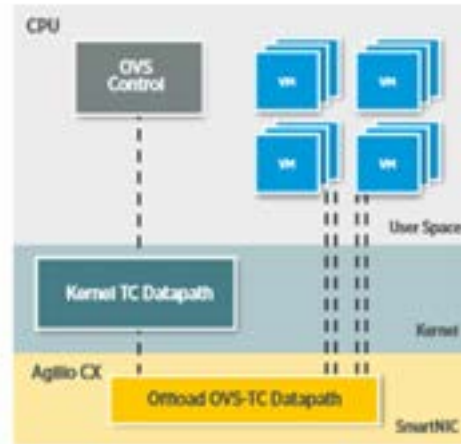


Figure 4: OVS-TC datapath

420. Claim 1 of the '586 patent further recites “receiving, from the substrate network, a second communication addressed to the computing node.”

421. Nuage Networks Virtualized Cloud Services practices this limitation. For example, the Open vSwitch “implement[s] the OVS datapath in user space,” in which “the DPDK poll mode driver delivers packets”, *e.g.*, second communication, “directly into the dedicated user space application, bypassing the Linux kernel stack altogether.”¹⁵³

422. Claim 1 of the '586 patent further recites “based at least in part on the configuration information, dropping the second communication without forwarding the second communication to the computing node.”

423. The Nuage Networks Virtualized Cloud Services practices this limitation. For example, “OVS-TC includes an action side which allows packets to be modified, forwarded or dropped,” as seen below.¹⁵⁴

¹⁵³*Id.*

¹⁵⁴*Id.*

OVS-TC

TC Flower is a packet classifier in the Linux kernel and part of the kernel traffic classification subsystem. OVS-TC is the extension of OVS user space code to enable it to offload flows using TC Flower and TC actions. OVS-TC allows matching on a variety of predefined flow keys. The user can match on IP addresses, UDP/TCP ports, metadata and more. Like OVS, OVS-TC includes an action side which allows packets to be modified, forwarded or dropped. The user can influence what is offloaded or not down to a per flow basis.

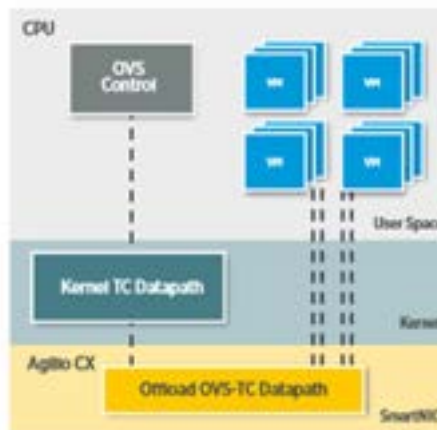


Figure 4: OVS-TC datapath

424. Claim 1 of the '586 patent further recites “wherein the managing further comprises: in response to an address resolution protocol (ARP) communication from the computing node regarding a second computing node in the virtual computer network, sending a spoofed response to the ARP communication indicating a virtual hardware address of the second computing node.”

425. Nuage Networks Virtualized Cloud Services practices this limitation. For example, the “VRS agent programs the flow tables in the kernel without VSC involvement. It also handles DHCP and ARP requests from the local VMs.”¹⁵⁵

¹⁵⁵Nuage Networks from Nokia, TECHNICAL DESCRIPTION: Nuage Networks Virtualized Services Platform: Service chaining (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Service_Chaining_Technical_Description_Document_EN-compressed.pdf, p. 8.

As VMs get added or removed, the related VRS agent(s) are fully programmed as a result with all the information required for local VMs:

- ACLs, L2 and L3 FIBs
- ARP tables
- QoS marking and policing
- Required frequency of statistics-gathering

As new flows get activated, the VRS agent programs the flow tables in the kernel without VSC involvement. It also handles DHCP and ARP requests from the local VMs.

426. Furthermore, “VRS includes” a “VRS Agent” that “replies to all ARP (no flooding),” as seen below.¹⁵⁶

3. VRS (Data Plane) - Virtual Routing and Switching plugin inside the Hypervisor. It's based on OVS, and it's responsible for L2/L3 forwarding, encapsulation.

On VRS you can define various VSC for redundancy and load balancing (one active and one standby), and each of them establishes an OpenFlow session using the Underlay network, not Management, using TCP port 6633 (SSL is optional).

VRS includes two main Nuage components:

- VRS Agent, that talks to VSC using OpenFlow. It's responsible for programming L2/L3 FIBs, and it replies to all ARP (no flooding). It also reports changes in VMs to the VSC. The forwarding table is pushed to VRS from VSC via OpenFlow. It has not only a view of all the IP and MAC addresses of the VMs being served by the local hypervisor, but also those which belong to the same domain (L2 and L3 segments), that is, all possible destinations of traffic for the VMs being served by that HV.
- Open vSwitch (OVS), provides Switching and Routing components and Tunneling to forward the traffic.

VRS supports a wide range of L2 and L3 encapsulation methods (VXLAN, VLAN, MPLSoGRE) so that it can communicate with a wide range of external network endpoints (other hypervisors, IP- or MPLS-based routers).



¹⁵⁶Mat Jovanovic, Welcome to Mat's Cloud, Hitchhikers Guide to Hybrid Cloud, <https://matscloud.blogspot.com/2017/06/nuage-networks-vsp-deep-dive.html>.

427. Nokia is and has been on notice of the infringement of the '586 patent at least as of the time Amazon filed and provided notice of this Complaint.

428. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '586 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, the Nuage Networks Virtualized Services Platform.

429. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 1 of the '586 patent.

430. Nokia will sell the Nuage Networks Virtualized Services Platform with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '586 patent.

431. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts would cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

432. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '586 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of the Nuage Networks Virtualized Services Platform, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '586 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

COUNT XII: PATENT INFRINGEMENT OF U.S. PATENT NO. 11,336,529

433. Amazon incorporates by reference the preceding paragraphs as if fully stated herein.

434. Amazon is informed and believes, and on that basis alleges, that Nokia has infringed and is currently infringing one or more claims (*e.g.*, claim 1) of the '529 patent in violation of 35 U.S.C. § 271, and will continue to do so.

435. Nokia has infringed and is currently infringing literally and/or under the doctrine of equivalents, by, among other things, making, using, offering for sale, selling, and/or importing within this judicial district and elsewhere in the United States, infringing products, including but not limited to Nokia Nuage Networks Virtualized Services Platform, in violation of 35 U.S.C. § 271(a).

436. By way of example only, Nokia Nuage Networks Virtualized Services Platform meets all the limitations of at least independent claim 1 of the '529 patent, either literally or under the doctrine of equivalents.

437. Exemplary claim 1 of the '529 patent recites:

1. A method, comprising: performing, by one or more computing systems:

providing a configurable network service accessible over one or more networks by a plurality of clients;

creating, by the configurable network service, a virtual computer network of virtual machines according to one or more client requests, including:

creating at least one logical sub-network of the virtual computer network; and

creating a first virtual machine and a second virtual machine in the logical sub-network; and

managing, by the configurable network service, communications in the virtual computer network, including:

intercepting an Address Resolution Protocol (ARP) request sent by the first virtual machine, wherein the ARP request requests a Media Access Control (MAC) address associated with an Internet Protocol (IP) address of with the second virtual machine;

responsive to the ARP request, sending an ARP response to the first virtual machine indicating the MAC address, wherein the ARP response causes the first virtual machine to send packets to the IP address of the second virtual machine using frames comprising the MAC address;

receiving a frame from the first virtual machine, the frame comprising the MAC address and a packet;

determining that the frame complies with an access control policy, wherein the access control policy allows or denies a communication in the virtual computer network based on a source, a destination, a direction, or a protocol used for the communication; and

routing the packet to the second virtual machine.

438. For the preamble of claim 1, to the extent the preamble is determined to be limiting, Nokia Nuage Networks Virtualized Services Platform practices a “method, comprising: performing, by one or more computing systems.”

439. For example, “The Nuage Networks Virtualized Services Platform (VSP) is the industry leading network automation platform enabling a complete range of SDN, SD-WAN, and cloud solutions,” as seen below.¹⁵⁷

¹⁵⁷Nuage Networks from Nokia, Nuage Networks Virtualized Services Platform, <https://www.nuagenetworks.net/platform/virtualized-services-platform/>.



Nuage Networks Virtualized Services Platform

The Nuage Networks Virtualized Services Platform (VSP) is the industry leading network automation platform enabling a complete range of SDN, SD-WAN, and cloud solutions. VSP provides advanced network automation across networks and clouds of all sizes and architectures, from datacenter private clouds to large enterprise wide area networks (WANs) and some of the largest public clouds in the world. VSP enables both Virtualized Network Services (VNS) for SD-WAN and Virtualized Cloud Services (VCS) for SDN.

440. Claim 1 of the '529 patent further recites “providing a configurable network service accessible over one or more networks by a plurality of clients.”

441. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, “VSP provides advanced network automation across networks and clouds of all sizes and architectures, from datacenter private clouds to large enterprise wide area networks (WANs) and some of the largest public clouds in the world. VSP enables both Virtualized Network Services (VNS) for SD-WAN and Virtualized Cloud Services (VCS) for SDN.”¹⁵⁸

442. Claim 1 of the '529 patent further recites “creating, by the configurable network service, a virtual computer network of virtual machines according to one or more client requests.”

443. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, Nokia Nuage Networks Virtualized Services Platform performs “virtual networks orchestration across the WAN, SD-WAN, and datacenter,” as seen below.¹⁵⁹

¹⁵⁸*Id.*

¹⁵⁹Virtual networks orchestration (VNO), <https://www.nokia.com/networks/ip-networks/virtual-networks-orchestration/>.



The slide features the Nokia logo in the top left corner. To its right is a navigation menu with the following items: "Service providers", "Enterprises and governments", "Licensing", "Partners", and "Insigh". The main content area has a dark blue-to-green gradient background. It contains the title "Virtual networks orchestration (VNO)" in large white font, followed by a subtitle in smaller white font: "Maximize total network ROI and increase average revenue per branch through virtual networks orchestration across the WAN, SD-WAN, and datacenter."

Overview

The Nokia Virtual Networks Orchestration (VNO) solution leverages Nuage Networks delivered SD-WAN services to dynamically connect existing WAN services and new hybrid services to branches faster, and to automate setup for rapid delivery of revenue-generating virtualized cloud services to branch sites.

Nokia VNO automates the setup of SD-WAN overlay services from the datacenter to branch sites across the private cloud of the service provider, or across the public cloud, with security being enabled by using IPsec, for example. Additionally, a broad range of multi-vendor virtualized security and other value-added service (VAS) applications can be enabled through establishing network connectivity through service chains in the datacenter.

444. Furthermore, Nokia Nuage Networks Virtualized Services Platform supports “deployment of a particular VM in a mixed hypervisor environment” and allows “select[ing] a mixture between KVM, VMware and Hyper-V,” as seen below.¹⁶⁰

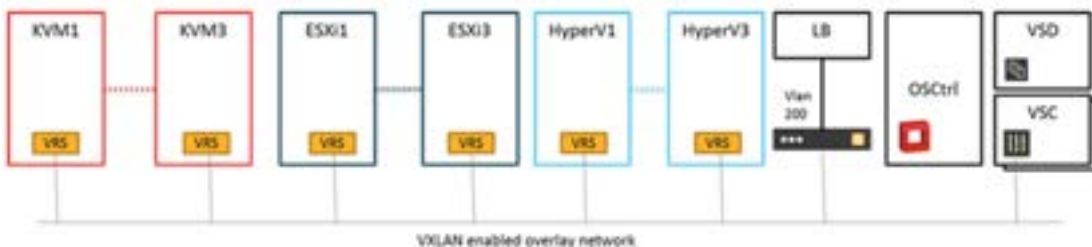
¹⁶⁰Hyper-V Integration for OpenStack (2017), <https://nuagenetworks.github.io/2017/05/04/Hyper-V-Integration.html>.

Example Deployment

Sample Architecture

As an example, we will show the deployment of a particular VM in a mixed hypervisor environment.

In this example, we selected a mixture between KVM, VMware and Hyper-V, and we also added a native VLAN port.



445. Claim 1 of the '529 patent further recites “creating at least one logical sub-network of the virtual computer network.”

446. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, a “sample architecture” based on Nokia Nuage Networks Virtualized Services Platform includes “[t]wo logical subnets [] deployed in OpenStack to boot VMs against, with a variety of virtual instances in each subnet.”¹⁶¹

447. Claim 1 of the '529 patent further recites “creating a first virtual machine and a second virtual machine in the logical sub-network.”

448. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, a “sample architecture” based on Nokia Nuage Networks Virtualized Services Platform includes “[t]wo logical subnets [] deployed in OpenStack to boot VMs against, with a variety of virtual instances in each subnet.”¹⁶²

¹⁶¹*Id.*

¹⁶²*Id.*

449. Claim 1 of the '529 patent further recites “managing, by the configurable network service, communications in the virtual computer network.”

450. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, Nokia Nuage Networks Virtualized Services Platform includes a “Virtualized Services Controller (VSC)” and a “Virtual Routing and Switching (VRS)” and “[t]he VSC runs a virtualized version of the Nokia Service Router operating system (SR OS) and acts as the SDN controller of the solution. It maintains the forwarding table for every tenant router (logical), and programs the forwarding plane elements (VRS).”¹⁶³

451. Claim 1 of the '529 patent further recites “intercepting an Address Resolution Protocol (ARP) request sent by the first virtual machine, wherein the ARP request requests a Media Access Control (MAC) address associated with an Internet Protocol (IP) address of with the second virtual machine.”

452. Nokia Nuage Networks Virtualized Services Platform supports a “mac command” that “sets the MAC address used in ARP responses when the virtual router instance is master. Routing of IP packets with mac-address as the destination MAC is also enabled. The mac setting must be the same for all virtual routers participating as a virtual router or indeterminate connectivity by the attached IP hosts will result. All VRRP advertisement messages are transmitted with mac-address as the source MAC.”¹⁶⁴

453. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, the “Virtual Routing and Switching (VRS)” of Nokia Nuage Networks Virtualized

¹⁶³*Id.*

¹⁶⁴m Commands, Nokia CLASSIC CLI COMMAND REFERENCE GUIDE RELEASE 21.7.R1, https://infocenter.nokia.com/public/7750SR217R1A/index.jsp?topic=%2Fcom.sr.classic%2Fhtml%2Fclassic_m_commands.html, p. 3666.

Services Platform “participates in the forwarding-plane of the SDN environment,” “handles all L2-L4 operations according to the network policy defined by the VSD,” and “handles ARP requests locally, implements its own DHCP server, relays OpenStack metadata, etc.”¹⁶⁵

454. Furthermore, “Virtual Routing and Switching (VRS) is a software module that is installed in the hypervisor layer for VMs or as part of a container structure in virtualized server environments. It creates and manages the virtual endpoints (i.e. VXLAN Tunnel Endpoint (VTEP)) that are used for the virtual ‘overlay’ tunnels between the workloads (e.g. VMs and containers) within a prescribed overlay VPN. These tunnels are created by adding a VXLAN encapsulation to the original Ethernet frame from designated traffic allowing it to be directly routable to other VMs or containers within that overlay VPN,” as seen below.¹⁶⁶



Virtual Routing and Switching

Virtual Routing and Switching (VRS) is a software module that is installed in the hypervisor layer for VMs or as part of a container structure in virtualized server environments. It creates and manages the virtual endpoints (i.e. VXLAN Tunnel Endpoint (VTEP)) that are used for the virtual “overlay” tunnels between the workloads (e.g. VMs and containers) within a prescribed overlay VPN. These tunnels are created by adding a VXLAN encapsulation to the original Ethernet frame from designated traffic allowing it to be directly routable to other VMs or containers within that overlay VPN.

455. Claim 1 of the ’529 patent further recites “responsive to the ARP request, sending an ARP response to the first virtual machine indicating the MAC address, wherein the ARP response causes the first virtual machine to send packets to the IP address of the second virtual machine using frames comprising the MAC address.”

¹⁶⁵Hyper-V Integration for OpenStack (2017), <https://nuagenetworks.github.io/2017/05/04/Hyper-V-Integration.html>.

¹⁶⁶NOKIA Accelerate Telco Clouds with Nuage Networks Virtualized Cloud Services, https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nokia_Accelerate_Telco_Clouds_with_Nuage_Networks_VCS_Application_Note_EN.pdf, p. 12.

456. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, the “Virtual Routing and Switching (VRS)” “handles ARP requests locally, implements its own DHCP server, relays OpenStack metadata, etc.”¹⁶⁷ Furthermore “[i]n the case of a Nuage Networks installation: The OVS user-space component was further extended with functionality such as local ARP/DHCP responder, support for redundant SDN controllers, multicast handlers, etc.”¹⁶⁸

457. Nokia Nuage Networks Virtualized Services Platform further supports a “mac command” that “sets the MAC address used in ARP responses when the virtual router instance is master. Routing of IP packets with mac-address as the destination MAC is also enabled. The mac setting must be the same for all virtual routers participating as a virtual router or indeterminate connectivity by the attached IP hosts will result. All VRRP advertisement messages are transmitted with mac-address as the source MAC,” as seen below.¹⁶⁹

¹⁶⁷Hyper-V Integration for OpenStack (2017), <https://nuagenetworks.github.io/2017/05/04/Hyper-V-Integration.html>.

¹⁶⁸*Id.*

¹⁶⁹m Commands, Nokia CLASSIC CLI COMMAND REFERENCE GUIDE RELEASE 21.7.R1, https://infocenter.nokia.com/public/7750SR217R1A/index.jsp?topic=%2Fcom.sr.classic%2Fhtml%2Fclassic_m_commands.html, p. 3666.

Syntax	<code>mac mac-address</code> <code>no mac</code>
Context	[Tree] (config>router>if>vrrp mac) [Tree] (config>router>if>ipv6>vrrp mac)
Full Context	configure router interface ipv6 vrrp mac configure router interface vrrp mac
Description	<p>This command sets an explicit MAC address used by the virtual router instance overriding the VRRP default derived from the VRID.</p> <p>Changing the default MAC address is useful when an existing HSRP or other non-VRRP default MAC is in use by the IP hosts using the virtual router IP address. Many hosts do not monitor unessential ARPs and continue to use the cached non-VRRP MAC address after the virtual router becomes master of the host's gateway address.</p> <p>The <code>mac</code> command sets the MAC address used in ARP responses when the virtual router instance is master. Routing of IP packets with <code>mac-address</code> as the destination MAC is also enabled. The <code>mac</code> setting must be the same for all virtual routers participating as a virtual router or indeterminate connectivity by the attached IP hosts will result. All VRRP advertisement messages are transmitted with <code>mac-address</code> as the source MAC.</p> <p>The command can be configured in both non-owner and owner vrrp nodal contexts.</p> <p>The <code>mac</code> command can be executed at any time and takes effect immediately. When the virtual router MAC on a master virtual router instance changes, a gratuitous ARP is immediately sent with a VRRP advertisement message. If the virtual router instance is disabled or operating as backup, the gratuitous ARP and VRRP advertisement message is not sent.</p> <p>The <code>no</code> form of the command restores the default VRRP MAC address to the virtual router instance.</p>

458. Furthermore, “Virtual Routing and Switching (VRS) is a software module that is installed in the hypervisor layer for VMs or as part of a container structure in virtualized server environments. It creates and manages the virtual endpoints (i.e. VXLAN Tunnel Endpoint (VTEP)) that are used for the virtual ‘overlay’ tunnels between the workloads (e.g. VMs and containers) within a prescribed overlay VPN. These tunnels are created by adding a VXLAN encapsulation to the original Ethernet frame from designated traffic allowing it to be directly routable to other VMs or containers within that overlay VPN,” as seen below.¹⁷⁰

¹⁷⁰NOKIA Accelerate Telco Clouds with Nuage Networks Virtualized Cloud Services, https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nokia_Accelerate_Telco_Clouds_with_Nuage_Networks_VCS_Application_Note_EN.pdf, p. 12.

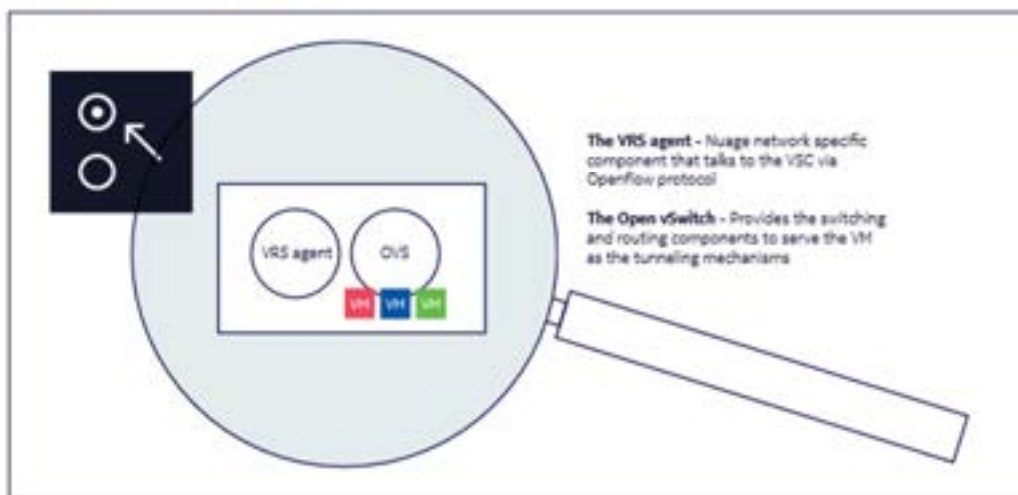


Virtual Routing and Switching

Virtual Routing and Switching (VRS) is a software module that is installed in the hypervisor layer for VMs or as part of a container structure in virtualized server environments. It creates and manages the virtual endpoints (i.e. VXLAN Tunnel Endpoint (VTEP)) that are used for the virtual "overlay" tunnels between the workloads (e.g. VMs and containers) within a prescribed overlay VPN. These tunnels are created by adding a VXLAN encapsulation to the original Ethernet frame from designated traffic allowing it to be directly routable to other VMs or containers within that overlay VPN.

As shown in Figure 6, VRS is based on the Open vSwitch (OVS) which is an open-source implementation of a distributed virtual multilayer switch which also provides the VTEP function. In addition to providing the virtual switching and forwarding plane for VMs and containers, VRS also offers L2-L4 capabilities such as offering a distributed L4 ACL-based firewall.

Figure 6. The Anatomy of the VRS



459. Claim 1 of the '529 patent further recites "receiving a frame from the first virtual machine, the frame comprising the MAC address and a packet."

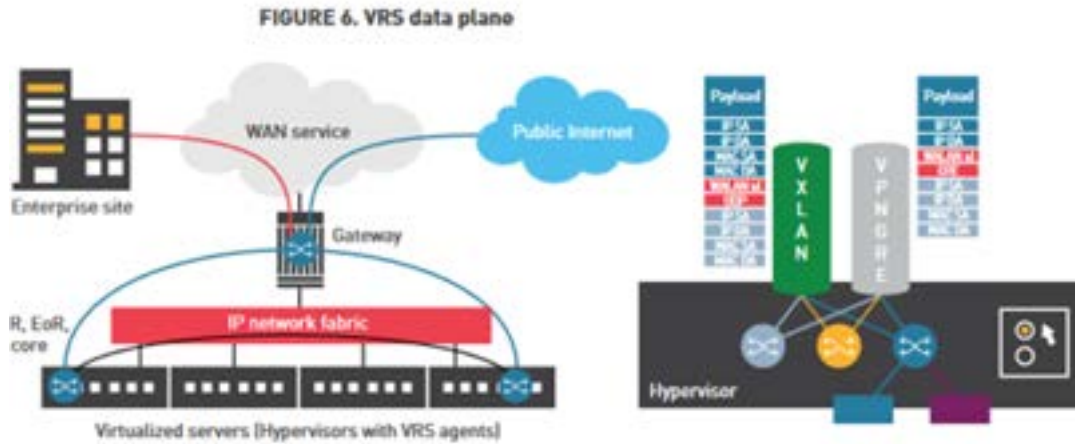
460. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, Nokia Nuage Networks Virtualized Services Platform enables Virtual Private LAN Service (VPLS) in which "[w]hen the MAC is known (populated in the VPLS FDB), all packets destined for the MAC (routed or bridged) are targeted to the specific virtual port where the MAC has been learned."¹⁷¹

¹⁷¹ARP and VPLS FDB Interactions, https://infocenter.nokia.com/public/7750SR217R1A/index.jsp?topic=%2Fcom.nokia.L2_Services_and_EVPN_Guide_21.7.R1%2Farp_and_vpls_fd-ai9enrmqvk.html, p. 261.

461. Claim 1 of the '529 patent further recites “determining that the frame complies with an access control policy, wherein the access control policy allows or denies a communication in the virtual computer network based on a source, a destination, a direction, or a protocol used for the communication.”

462. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, Nokia Nuage Networks Virtualized Services Platform implements ACLs (access control lists) such that “[a]s packets are received from the VM VNICs they are processed by the associated VRS instance: ACLs are evaluated, L2 and optionally L3 lookups are performed to determine how the rest of the packets in the flow should be treated,” as seen below.¹⁷²

¹⁷²Nuage Networks from Nokia, TECHNICAL DESCRIPTION, Nuage Networks Virtualized Services Platform: Service chaining (2016), https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nuage_Networks_Service_Chaining_Technical_Description_Document_EN-compressed.pdf, p. 6.



The VRS service on the left, used to network the green and red VMs, is implemented using the service primitives shown on the right. In each of the two hypervisors there is a blue VRS instance used to isolate the tenant quintuple flows from other tenant instances and to provide internal and external connectivity for the green and red VMs. Two other tenant instances share tunnels with the blue VRS: one or more VXLAN tunnels to other hypervisors and one more VPN tunnels to the gateway(s). Each VRS instance is instantiated as a combination of L3 and L2 forwarding entities:

- one distributed L3 Virtual Routing and Forwarding [VRF] instance
- one VXLAN virtual bridge for every VM subnet implementing the MAC FIB
- every VXLAN is represented in VRF as an IRB interface

As packets are received from the VM VNICs they are processed by the associated VRS instance: ACLs are evaluated, L2 and optionally L3 lookups are performed to determine how the rest of the packets in the flow should be treated. The resulting next-hop could be either a local VM VNIC or a tunnel:

463. Claim 1 of the '529 patent further recites “routing the packet to the second virtual machine.”

464. Nokia Nuage Networks Virtualized Services Platform practices this limitation. For example, “Virtual Routing and Switching (VRS)” of Nokia Nuage Networks Virtualized Services Platform creates “overlay VPNs or ‘tunnels’” “in the forwarding plane by adding a further layer of encapsulation to native Ethernet frames making it routable yet isolated from all other neighboring traffic.”¹⁷³

¹⁷³NOKIA Accelerate Telco Clouds with Nuage Networks Virtualized Cloud Services, https://www.nuagenetworks.net/wp-content/uploads/2020/06/Nokia_Accelerate_Telco_Clouds_with_Nuage_Networks_VCS_Application_Note_EN.pdf, p. 4.

465. Nokia is and has been on notice of the infringement of the '529 patent at least as of the time Amazon filed and provided notice of this Complaint.

466. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '529 patent by active inducement under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and/or offer to sell in the United States, Nuage Networks Virtualized Services Platform.

467. For example, Nokia will provide directions, technical support, guides, marketing materials, instruction manuals, and/or other information that encourage and facilitate infringing use by others of the method as claimed in claim 1 of the '529 patent.

468. Nokia will sell Nuage Networks Virtualized Services Platform with the knowledge and intent that customers who buy it will use it for their infringing use and therefore that customers will be directly infringing the '529 patent.

469. Nokia will intend and continue to intend to induce patent infringement by its customers and will have knowledge that the inducing acts would cause infringement or will be willfully blind to the possibility that its inducing acts will cause infringement.

470. Nokia will also infringe indirectly and continue to infringe indirectly one or more claims of the '529 patent by contributory infringement under 35 U.S.C. § 271(c) by offering to sell, selling, or importing into the United States, components of Nuage Networks Virtualized Services Platform, which include non-standard software, knowing the same to be especially made or especially adapted for use in an infringement of the '529 patent, and not a staple article or commodity of commerce suitable for substantial non-infringing use.

ATTORNEYS' FEES

471. Amazon is entitled to recover reasonable and necessary attorneys' fees under applicable law.

DEMAND FOR JURY TRIAL

Amazon hereby demands trial by jury on all claims and issues so triable.

PRAYER FOR RELIEF

WHEREFORE, Amazon respectfully requests that this Court enter judgment against NOKIA as follows:

- A. That NOKIA has infringed, actively induces infringement of, and contributorily infringes, and continues to infringe the Asserted Patents in violation of 35 U.S.C. § 271 (a), (b), and (c);
- B. An injunction against further direct or future indirect infringement of the Asserted Patents;
- C. An award of damages adequate to compensate Amazon for the patent infringement that has occurred, together with pre-judgment interest and costs;
- D. An accounting for any infringing sales not presented at trial and an award by the Court of additional damages for any such infringing sales;
- E. An award of all other damages permitted by 35 U.S.C. § 284;
- F. That this is an exceptional case and merits an award to Amazon of its costs and reasonable attorneys' fees incurred in this action as provided by 35 U.S.C. § 285; and
- G. Such other relief as this Court deems just and proper.

MORRIS, NICHOLS, ARSHT & TUNNELL LLP

/s/ Jeremy Tigan

OF COUNSEL:

J. David Hadden
FENWICK & WEST LLP
801 California Street
Mountain View, CA 94041
(650) 988-8500

July 30, 2024

Jack B. Blumenfeld (#1014)
Jeremy A. Tigan (#5239)
1201 North Market Street
P.O. Box 1347
Wilmington, DE 19899
(302) 658-9200
jblumenfeld@morrisnichols.com
jtigan@morrisnichols.com

Attorneys for Plaintiffs