

Network Virtualization – a View from the Bottom

Jorge Carapinha
PT Inovação
Rua Eng. Ferreira Pinto Basto
3810-106 Aveiro Portugal
+351 234 403200
jorgec@ptinovacao.pt

Javier Jiménez
Telefónica I+D
C/ Emilio Vargas, 6
28043 Madrid Spain
+34 913 374 000
fjjc@tid.es

ABSTRACT

The interest in network virtualization has been growing steadily among the networking community in the last few years. Network virtualization opens up new possibilities for the evolution path to the Future Internet by enabling the deployment of different architectures and protocols over a shared physical infrastructure. The deployment of network virtualization imposes new requirements and raises new issues in relation to how networks are provisioned, managed and controlled today. The starting point for this paper is the network virtualization reference model conceived in the framework of the EU funded 4WARD project. In this paper we look at network virtualization mainly from the perspective of the network infrastructure provider, following the 4WARD network virtualization architecture and evaluate the main issues and challenges to be faced in commercial operator environments.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Architecture and Design

General Terms

Design, Management

Keywords

Network Virtualization, Resource Management

1. INTRODUCTION

In the last few years, the concept of network virtualization has attracted a great deal of attention from industry and research fora. Although it is not a strictly new concept, a network virtualization “renaissance” has been originated mainly from the realization that it can provide a platform upon which novel network architectures can be built, experimented and evaluated, freed from legacy technology constraints. Furthermore, network virtualization has been heralded as the keystone of a new global Internet architecture, a new component enabling the coexistence of multiple networking approaches, thus overcoming the present

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
VISA '09, August 17, 2009, Barcelona, Spain.
Copyright 2009 ACM 978-1-60558-595-6/09/08...\$10.00.

Internet “ossification” problem and stimulating the development and deployment of new network technologies and advanced applications. In addition, virtualization is expected to provide a clean separation of services and infrastructure and facilitate new ways of doing business by allowing the trading of network resources among multiple providers and customers.

Up to now the concept of network virtualization has been explored essentially in testbeds, on a limited scale. Several proposals for network virtualization architecture or for the utilization of network virtualization in several contexts have been put forward [2, 3, 4, 5, 12, 13]. However, so far, the impact of network virtualization on the underlying physical infrastructure has been relatively overlooked. In particular, the feasibility of the network virtualization concept in large scale operator networks has not been adequately assessed.

There are multiple challenges associated with the deployment of network virtualization in operator infrastructures. A few examples are how to enforce isolation between different virtual networks, how to conciliate a wide range of requirements of different virtual networks, how to fulfill carrier-class level reliability, how to guarantee scalability. In this paper, we are mainly interested in evaluating the impact of virtualization on the underlying network infrastructure. We use the 4WARD network virtualization architecture as the main reference [1].

2. OVERVIEW AND STATE OF THE ART

2.1 From VPNs to VNs

In general, virtualization provides an abstraction between user and physical resources, so that the user gets the illusion of direct interaction with those physical resources. In the last few years, significant advances in operating system virtualization technologies have been achieved, with contributions from major players [6, 7, 8, 9], which has enabled the use of virtualization in a growing number of contexts and applications. Major router vendors are also following this trend and it seems likely that network virtualization will become a reality in operator networks in a not too distant future [10, 11].

The concept of network virtualization is not new – it was materialized in the past (albeit in a limited way) with network-based VPNs¹, which have been a highly successful approach to provide separate virtual networks over a common physical infrastructure. One may wonder why full-blown network

¹ In the context of this paper, by VPN we mean a network-based provider provisioned layer 3 Virtual Private Network, of which the BGP/MPLS VPN model [19] is a common example.

virtualization is in fact required if the VPN model has been so successful to deploy the concept of virtual network (VN).

VPNs fulfill the basic goal of providing different logical networks over a shared infrastructure, but suffer from a few limitations, among others:

- All virtual networks are based on the same technology and protocol stack, which precludes the coexistence of different networking solutions;
- A real isolation of virtual network resources is not possible, by default;
- A clean separation of the roles of infrastructure provider and VPN service provider is not possible and in practice they are played by the same entity.

Network virtualization goes a step further by enabling independent programmability of virtual networks. This means that a virtual network is no longer necessarily based on IP, or any other specific technology, and any kind of network architecture can in principle be built over a virtual network.

Another important strength of virtualization is the native capability to handle multi-provider scenarios and hide the specificities of the network infrastructure, including the existence of multiple administrative domains. Although some workaround solutions exist, this has traditionally been a complicated issue for VPNs.

Finally, network virtualization provides a real isolation of virtual networks sharing the same infrastructure, for example by means of different operating system instances, and not just an illusory isolation, as provided by VPNs.

2.2 Basic components

Network virtualization is composed of two main components – link virtualization and node virtualization.

Link virtualization allows the transport of multiple separate virtual links over a shared physical link. A virtual link is often identified explicitly by a tag, but can also be identified implicitly, by a time slot or a wavelength, for example. A wide variety of the standard link virtualization techniques available in today’s Internet (e.g. ATM, Ethernet 802.1q, MPLS) may in principle be used for this purpose.

The basic elements of network virtualization are shown in Figure 1. At the substrate level, the substrate node is network equipment capable of supporting virtual nodes by means of any virtualization technology. A single substrate node typically contains a number of virtual nodes.

The virtual node is another central element in virtual network architecture. The concept of node virtualization has been materialized, to some extent, by virtual routers². Although the practical realization of the concept has had a few variations, in broad terms, a virtual router appears for all purposes (e.g. configuration, management, monitoring, troubleshooting) like a

dedicated physical router. However, in most cases, a virtual router provides an illusion of isolation, rather than a real isolation, as it lacks dedicated memory, processing and I/O resources.




Figure 1 - Network virtualization model

Node virtualization is based on isolation and partitioning of hardware resources. Physical resources of a substrate node (e.g. CPU, memory, storage capacity, link bandwidth) are partitioned into slices and each slice is allocated to a virtual node according to a set of requirements. Recent developments in operating system virtualization have permitted significant advances in terms of fairness and performance.

Virtualization of substrate nodes, in combination with virtualization of links interconnecting those substrate nodes, enables the creation of virtual networks, functionally equivalent to a physical network.

A substrate node terminates one or multiple physical links, which may in turn carry multiple virtual links. Correspondingly, a virtual node terminates one or multiple virtual links. For scalability reasons, virtual links may be handled per aggregate, rather than per virtual link, as will be discussed later on. Some substrate nodes, in addition to terminate virtual links, are also able to forward traffic transported in virtual links in a transparent way (this is the case of node b in Figure 2).




Figure 2 - Basic network virtualization elements

From a functional point of view, substrate nodes can be classified in four major groups, according to the type of functionality that they provide:

² In the context of this paper, we prefer the technology-agnostic term *virtual node*, as any network protocol other than IP should in principle be deployable on a virtual network.

- Edge nodes: located at the edge of the substrate network, these nodes are mainly points of presence (PoP) of the network infrastructure provider. Typically, these edge nodes are connected to end users, either directly or through an access network provider, which may be itself virtualized, or not; geographical location is usually a key attribute of this kind of nodes.
- Core VN-capable nodes: support virtualization, i.e. can host virtual nodes and can also support non-VN traffic (again, node b in Figure 2). A core VN-capable node may be connected to edge nodes, or to other core nodes, either VN-capable or not.
- Core transport nodes: do not support node virtualization, are typically core nodes used to transport traffic.
- Border nodes: are connected to neighbor border nodes located in different network domains.

It should be noted that these functions are not mutually exclusive – for example, node b in Figure 2 supports both core VN-capable node and core transport node functions. Figure 3 shows a scenario with two network domains, A and B, and the four types of network nodes identified above.




Figure 3 - Types of substrate nodes

2.3 Business roles

In general, the 4WARD network virtualization model includes three different business roles:

- The Infrastructure Provider (InP) deploys and runs the network physical resources, and partitions them into isolated virtual slices, using some virtualization technology. Typically, these resources are offered to another service provider, and not to end users (but the InP customer might as well be a corporation using the virtual network for its internal use, rather than to build commercial end user services). The InP has a vision of the resources allocated to each VN, but not of the protocols running inside.

- The virtual network provider (VNP) is responsible for finding and composing the adequate set of virtual resources from one or more infrastructure providers, in order to fulfill the virtual network operator request. The VNP leases slices of the virtualized infrastructure from one or more InPs and puts them together. Strictly speaking, what the VNP provides is not a network yet, but just an empty container where the virtual network operator builds the protocols that make up the VN.
- The virtual network operator (VNO) deploys any protocol stack and network architecture over a VN, independently of the underlying physical network technologies. The VNO operates, maintains, controls and manages the virtual network. Ideally, the fact that resources are virtual, rather than physical, should not imply any major impact from an operational point of view. Thus, the role of the VNO should be indistinguishable from that of any operator running a native network infrastructure. VNOs have a unified view of the network, regardless of the multiple infrastructure domains on which it is built.

It should be noted that this model does not preclude the possibility of more than one role being played by a single entity. In a vertically integrated scenario, the three roles are typically played by the same operator. Yet, even in this case, a functional separation of roles should make sense.

3. INCENTIVES TO DEPLOY NETWORK VIRTUALIZATION

No matter how interesting the concept of network virtualization may be from a technical point of view, it will only become a reality in commercial environments if there are enough incentives for network providers to build it. Some scenarios in which an infrastructure provider could benefit from implementing network virtualization could be the following:

- Network operators reselling infrastructure to third parties: the service could be conceived as an enhancement of the VPN portfolio, in which the provision of router capacity is added. Nevertheless, networks may be considered as key assets from an operator's perspective. There might be different pricing expectations from different stakeholders. The development of novel compensation mechanisms will be necessary to ensure that each role finds a place in the value chain.
- Network operators using virtualization to diversify their own infrastructure for private purposes, or to trusted or corporate third parties (e.g. mobile operator renting a virtual network from the fixed branch): the aim would be to optimize the delivery of multiple isolated services and network architectures over a common, cost-effective infrastructure.
- Network sharing: along with a cost-reduction strategy, network operators are steadily exploring the deployment of common infrastructures (e.g. in emerging or non-strategic markets) to share capital investments. Network virtualization could be used in this scenario to assure a

proper isolation of the different networks while minimizing the total cost of ownership [17].

- **Managed services:** Some telecom operators are increasingly focusing their growth strategies towards services delivery, customer orientation and product marketing. In this context, a potential approach would be the externalization of infrastructure to better focus on their core service oriented business. A third party (e.g. a system integrator) could in this context become an infrastructure provider, and could therefore benefit from virtualization techniques to better capitalize its investments in new network deployments. The same approach could be followed by governments or public entities aiming at deploying neutral telecom infrastructure, to promote the development of the digital society. However, publicly managed networks have proven to be a difficult task, mainly for political reasons.
- In a research environment, virtual infrastructure providers could be promoted to permit the validation of emerging Internet architectures. In the long term, this approach could be followed by network operators aiming at providing enhanced peering models (i.e. more than best-effort), with the goal to increase the value for global connectivity services.

4. DESIGNING VN-ENABLED NETWORKS

4.1 Infrastructure Provider requirements

In a network virtualization scenario, an InP must fulfill generic network infrastructure provider requirements, as well as new requirements specifically related with virtualization of network resources, as detailed in the following paragraphs.

- **Robustness:** The network should continue to operate in the event of node or link failure. A virtual network should provide carrier-class reliability (typically in the order of 99,99% to 99,999%).
- **Manageability:** The InP must have a view of the underlying physical topology, operational state, provisioning status, and other parameters associated with the equipment providing the virtual network.
- **Traffic and network resource control:** Traffic engineering and management techniques performed by the InP must not constrain the basic operation of a VN in any way.
- **Isolation:** Mechanisms to guarantee deterministic degrees of isolation between virtual networks must be provided. It must be guaranteed that misbehaving VNs (either accidentally or purposefully) do not affect the performance of other VNs sharing the same resources (e.g. memory space, storage, bandwidth, processing power), as well as non-VN traffic.
- **Scalability:** The number of VNs on a specific network administrative domain may grow to the order of thousands, or more. Any technical solution must be scalable to cope with any number of virtual networks. In

particular, the isolation of different virtual networks should be guaranteed in a scalable way.

- **Inter-domain interoperability:** In the case of VNs spanning multiple infrastructure provider domains, seamless interoperability must be guaranteed.
- **On-demand provisioning:** It must be possible to create, modify and remove VNs dynamically at request, with a minimal impact on the underlying infrastructure. The characteristics of a VN may change dynamically at runtime and in many cases a VN will have a limited lifespan.
- **Network technology agnosticism:** Network virtualization should not rely or depend on any specific network technology. In particular, multi-domain scenarios across dissimilar infrastructure domains should be possible.

4.2 Virtual Network ID, Virtual Link ID, virtual link aggregation and scalability

As stated before, a VN may span multiple network infrastructure domains. In order to provision and control VNs spanning multiple infrastructure domains, there must be a way to uniquely identify a specific VN on a global scale through a Virtual Network ID (VNID). This can be achieved by concatenating two pieces of information – the identification of the organization responsible for the VN (in principle, the VNO), which is supposed to be globally unique, and the ID allocated by the VNO to the VN, which must be unique within that specific administrative domain [18].

Additionally, in the data plane, identification of virtual links by means of a virtual link ID (VLID) is also required. The most straightforward solution would consist of simply mapping the VLID to whatever data link layer specific tag is in use. A number of available link virtualization techniques, such as ATM, Ethernet VLAN or MPLS, could perfectly be used for this purpose. For example, in the case of MPLS, a different LSP would be allocated to each virtual link. This means that the VLID would be used both as a virtual link identifier and a tag for data forwarding. This apparent simplicity does not come without a cost. On the one hand, for every new virtual link, a new VLID would have to be learned by all network nodes in the data path. Since the number of virtual links has a square law relationship to the number of virtual nodes in the network, this would represent a clear scalability concern. It is clear that large scale scenarios call for a more scalable approach.

As a matter of fact, scalability is a major requirement for virtual networks. One of the reasons why VPNs, namely those based on the BGP/MPLS model, have been so successful lies in its good scalability properties, which basically result from the clear separation of edge and core functions. Only edge nodes deal with VPNs and support virtualization (or an illusion of it, to be more precise), whereas core nodes are only concerned with forwarding traffic trunks and have no awareness of VPNs whatsoever³. Thus, the growth of the number of VPNs does not have any impact on the network core.

³ With the possible exception of BGP route reflectors used to disseminate VPN routes.

For this reason, virtual link aggregation will certainly constitute an important network virtualization requirement in operator networks. A virtual link aggregate can be defined as a set of virtual links that follow a common path and are treated similarly between a pair of virtual nodes. In practice, virtual link aggregation can be performed by carrying at least two types of identifiers in the data plane – one for endpoint VN identification and another for hop-by-hop forwarding, identifying the virtual link aggregate (nodes in the core need to be aware of the latter only). Virtual link aggregation simplifies the job of core substrate nodes (as they only have to deal with a limited number of virtual link aggregates, rather than a potentially high number of individual virtual links), enhances scalability (as the growth of number of VNs does not necessarily imply additional load in the network core nodes) and improves performance.

However, there are a few issues about virtual link aggregation in a network virtualization environment to be noted, so the straight replication of the VPN scenario may not be entirely possible. Figure 4 illustrates link aggregation and its relationship with node virtualization in four different scenarios. Cases 1 and 2 correspond to the legacy VPN scenario. In both cases, only the edge nodes (a and e) support node virtualization in the data path. Virtual link aggregation, used in case 2, allows nodes b, c and d to handle a single virtual link aggregate and ignore individual virtual links, thus executing a much simpler job, compared to case 1.

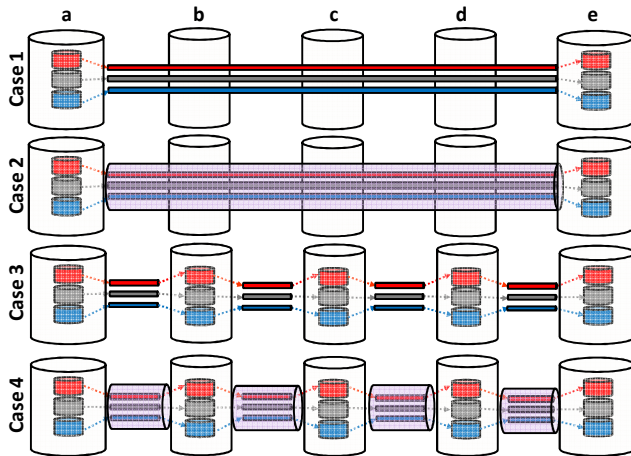


Figure 4 - Virtual link aggregation scenarios

By contrast, in cases 3 and 4, all substrate nodes in the data path support node virtualization. Contrary to the previous situation, the value of virtual link aggregation is in this case much smaller, if not null. In fact, to take advantage of the economies provided by link aggregation, the number of substrate nodes providing virtualization must be minimized. This is illustrated in cases 3 and 4 – link aggregation in case 4 does not provide any advantage compared to case 3. Link aggregation is only effective if a significant number of substrate nodes along the data path perform transport functions only, and not virtualization. This suggests a trade-off between data plane efficiency offered by virtual link aggregation and flexibility offered by node virtualization in the network core.

4.3 QoS and isolation

One of the limitations of legacy VPN approaches is that, because the matrix of incoming and outgoing traffic between every possible source-destination pair is impossible to predict, a guaranteed allocation of resources per VPN is usually not possible. On the contrary, VNs allow a fine grained allocation of resources, including bandwidth per virtual link on every segment of the substrate network. However, in practice, a fine grained per-VN resource control is not compatible with scalable aggregation-based approaches, described before. Hence, there is a trade-off between scalability (which favors aggregation) and strict VN isolation (which requires control of resources per virtual link).

Virtualization raises new QoS issues. More than QoS differentiation, virtualization requires QoS isolation. In all respects, including deterministic performance, a VN should replicate a network supported by physical resources. In a virtualization-based environment, QoS can be handled at two levels:

- Inside a VN, any QoS mechanism and tool can in principle be defined and deployed by the VNO (for example, different classes of service for voice and data). Intra-VN QoS mechanisms must not be constrained by the characteristics of the substrate (including traffic engineering performed by the InP), or by the traffic running in other VNs sharing the same resources.
- At the substrate level, QoS isolation between virtual networks is a key requirement, as any QoS policy and mechanism inside the VN cannot be built without adequate guarantees of performance. However, strict QoS isolation of virtual networks requires substrate nodes to apply a per-VN treatment to the traffic. As said before, this may not be possible, or desirable, in all cases, mainly for scalability reasons.

In the cases where virtual link aggregation is used, all virtual links inside a given virtual link aggregate get a similar QoS treatment by the substrate nodes. Isolation between virtual links in a common aggregate must be then guaranteed by policing the traffic admitted into each virtual link at the source end point. The substrate nodes in the middle have to make sure that bandwidth is shared between aggregates in a fair way.

5. BUILDING VIRTUAL NETWORKS

5.1 The VNP – InP interface

The VNP/InP interface is a key component of the network virtualization architecture. Through this interface, the VNP is able to request the establishment of a virtual network, the modification of the characteristics of a virtual network, or the removal of a virtual network. In turn, the InP acknowledges the VNP requests and notifies any relevant event (e.g. network error).

The VNP is expected to build the virtual network topology and define resource capacity requirements, namely link bandwidth and node computational capability. As discussed before, other characteristics such as geographical location of the edge nodes will be needed in many cases. The information provided by the VNP to the InP must contain a model of the virtual network topology as a graph, with a set of virtual nodes and virtual links and the applicable constraints. Each virtual node and virtual link

