

Chapter 2

Network Virtualization: Overview

Before presenting different network and wireless virtualization architectures and techniques, this chapter summarizes what other existing surveys on similar topics have already covered. However, the readers might wonder *why network virtualization is discussed in great details in a book on wireless virtualization?* In fact, the so-called *wireless virtualization* borrows many concepts and benefits from network virtualization. Additionally, wireless virtualization must *coexist* within the *context* of network virtualization in most applications. One of the goals of the first half of this book, (this chapter and [Chaps. 3, 4](#)) is to highlight different concepts and implementation technologies used in network virtualization that can be relevant, applied, integrated or extended to wireless virtualization. In this chapter, the readers are introduced to *virtualization*, a popular but often misused term. First, the evolution from computer and server virtualization to network virtualization is briefly discussed. Then, a classification tree to organize different virtualization architectures and technologies is constructed based on existing surveys. At last, an interpretation of the definition of *network virtualization* is given. The clear link between the network virtualization concepts presented in this chapter and wireless virtualization is emphasized.

2.1 From Server Virtualization to Network Virtualization

As mentioned in the previous chapter, virtualization technologies have recently moved from server virtualization to network virtualization. *Server virtualization*, also known as *host* or *computer virtualization*, allows multiple users to *share* the same server through *virtual machines* (VMs) by abstracting and decoupling the computing functionalities from the underlying hardware. Server virtualization has played a pivotal role in cloud computing as one of its main enablers. Through server virtualization, *on-demand provisioning* and *flexible management* of computing resources are made possible. Strictly speaking, server virtualization also includes the virtualization of network interfaces from the operating system point of view. However, it does not involve any virtualization of the network fabric, such




Fig. 2.1 Server virtualization versus network virtualization

as switches and routers. On the other hand, *network virtualization* enables multiple isolated virtual networks to share the same physical network infrastructure, as shown in Fig. 2.1. Thus, it decouples the network infrastructure from the services that it provides. This paradigm shift allows virtual networks with truly differentiated services to coexist on the same infrastructure, maximizing its reusability. These virtual networks can be deployed on demand and dynamically allocated, just as VMs in server virtualization would. The *functionalities* that each virtual network can obtain from the virtualized infrastructure range from simple *connectivity* and *performance guarantees* to advanced support for *new network protocols*.

Along with the recent development in software-defined networking (SDN) technologies, network virtualization has also become involved in cloud computing applications. For instance, network virtualization enables a more flexible management of the network interconnection between physical servers in a large data-center. New application of network virtualization can eventually allow the cloud services to extend beyond the datacenter into the network infrastructure. Using similar logic, the interesting question that arises here is whether virtualization can now be carried beyond the core network infrastructure and into the *wireless access network*. In other words, the feasibility and applicability of virtualization concepts in wireless access technologies must be considered. As explained in the Chap. 1, there are compelling motivations to do so.

2.2 Network Virtualization: Diverse Perspectives

Some surveys on network virtualization can be found in literature such as [1] and [2]. The survey [1] summarizes basic network virtualization concepts and covers selected virtualization testbeds developed by different research communities.

The survey [2] is an extension of [1] by presenting additional commercial technologies and formulating its own unified definition for network virtualization. According to survey [2], the term *network virtualization* is a very general term that can refer to *different concepts* under *different contexts*. It is often associated with resource allocation, virtual private network (VPN), virtual local area network (VLAN), multi-user multiple access techniques, SDN and a number of different concepts and technologies. In simplified scenarios, network virtualization is thought to be a technique of *network resource sharing*. However, depending on the context, it could be a combination of *resource abstraction*, *resource partitioning*, *resource aggregation*, *centralization of control*, *separation of data plane and control plane* and much more. To fully understand network virtualization, different *aspects* of it must be considered. An attempt to reorganize the different architectures has been made in [2]. Figure 2.2 represents a reconstruction of that classification. Interestingly, *wireless virtualization* also encompasses such complexities and will be approached in a similar fashion.

In Fig. 2.2, a separation between the industrial perspective and the academic perspective is made as in [2]. The term *industrial* can be contested because many of these commercial technologies were initially developed through academic research. It is more appropriate to view the component on the left-hand side of Fig. 2.2 as more matured technologies and implementation techniques [1]. The network virtualization techniques have been classified based on the network component where the virtualization is implemented. This format of classification is very helpful at providing some insight into what specific resources are being virtualized in each case. The right branch of the tree represents testbed architectures and management protocols that can incorporate a combination of specific technologies. This “academic” perspective mainly consists of emerging technologies and active research testbeds. However, these technologies will eventually mature




Fig. 2.2 Different aspects of network virtualization (based on [1], [2] and [4])

and merge with the commercial perspective. For instance, OpenFlow [3] will join *datapath virtualization* as a powerful alternative to label switching and tunneling. Many of these testbeds and technologies, including OpenFlow, will be respectively presented in [Chaps. 3 and 4](#).

Once again, an important question that arises is: *where does wireless network virtualization fit in the context of network virtualization?* The survey [2] presents wireless virtualization as one of the *frontier* research areas. Many of the wireless network virtualization techniques covered in this book can fall into the category of *device virtualization*. This is because the wireless network, formed by devices such as wireless access points and cellular basestations, lies at the *edge* of the core network. In addition, *wireless network management* is an important aspect that often requires collaboration with core network management functions in a virtualized wireless infrastructure. Nevertheless, the classification for wireless network virtualization is *multi-dimensional*. Some aspects of it are unique to wireless technologies and remain independent from core network concepts. The different aspects of wireless virtualization will be discussed in [Chap. 5](#). The different *wireless virtualization perspectives* will be derived in [Chap. 6](#).

2.3 Definition of Network Virtualization

Discussion in [2] pointed out that the usage of the term *network virtualization* is highly overloaded. Thus, an attempt to unify the various definitions is made in [2]. In this book, the definition of *virtualization* in the context of *computer virtualization* is the *abstraction* and *sharing* of *computing resources* to give multiple parties the illusion of sole ownership [2]. In computing, the basic resources are processing power, active memory, data storage and driver interfaces. Nonetheless, even within computer virtualization, this definition can be ambiguous depending on the *depth* of virtualization. For instance, before the introduction of modern virtual machines, the operating system itself can be considered as a basic form of computer virtualization at the program level. Before the advent of operating systems, mainframe computers can only execute one program at a time. The operating system allows multiple programs to *time-share* the same processing and memory resources, effectively ‘virtualizing’ the computer hardware. Thus, *virtualization* itself is clearly not a new concept. However, the usage of the term *computer or server virtualization* in this book does not refer to the operating system but generally implies a platform to offer simultaneous support for *multiple* operating systems.

In computer networks, the basic resources are *nodes* and *links*. A node can be any network equipment, most often a switch or a router. A link is a physical or logical (soft) connection between two nodes in the network. Some *interfacing resources*, such as network interfaces, can be considered a part of both host/server and network virtualization. *Identification resources* such as address space and labeling tags can also be treated as one type of network resources. Most often, a mechanism for the management of the identification for devices (IP address and MAC

address), links (flow ID), and networks (VLAN ID, VPN ID) is needed for the management and monitoring of a virtualized infrastructure. According to [2], identification resources are *soft resources* used to *assist* the virtualization of nodes and links. They are not considered to be *physical resources* on their own. High-level organizational structure such as a *network topology* can be created by assembling both soft and physical resources. Therefore, network virtualization is a *multi-dimensional* and *multi-layered* concept. This explains why the term *virtualization* is employed very loosely and often confusingly. After all, traditional VPN and VLAN can be considered a form of (primitive) network virtualization. However, many of the latest network virtualization projects can support more flexible functions beyond what VPN or VLAN can provide, such as the dynamic on-demand provisioning of resources and the support for entirely new network architectures.

The definition presented in the survey [2] covers very well the diverse spectrum of network virtualization by combining two basic concepts. Firstly, [2] clarifies the purpose of network virtualization, which is to provide a user or an application with an *illusion* of the *full ownership* of a network resource. Then, [2] describes the methods used to achieve this objective. In short summary, network virtualization can be defined as the *sharing of network resources* through the *abstraction* and *isolation* of *network functionalities* of the physical hardware. It is important not to confuse network virtualization with software-defined networking (SDN). SDN can be one of the *enabling technologies* for virtualization because by definition it provides abstraction and decoupling of functionalities from the hardware. However, the subtlety is that SDN by itself does not *necessarily* involve sharing and isolation.

This definition for network virtualization is broad enough to be applied to wireless virtualization to some extent. In the case of wireless virtualization, the purpose is to provide the mobile virtual network operators (MVNOs) with an abstracted ownership of the wireless access network. One can argue that wireless virtualization is simply a *subset* of network virtualization. One possible interpretation is that the wireless access points are equivalent to the network nodes and that the wireless channels are equivalent to the network links. However, the wireless medium is fundamentally different from the wired medium. This is due to the time-varying nature of the wireless channel and the high mobility of the end users. These wireless-specific issues will be addressed by different aspects of wireless virtualization and will be discussed in [Chap. 5](#). Finally, the term *hypervisor* refers to the layer responsible for the isolation among slices of the virtualized infrastructure. It is often referred to as virtual machine monitor (VMM) in the context of computer virtualization.

References

1. N.M. Mosharaf, K. Chowdhury, R. Boutaba, A survey on network virtualization. *J. Comp. Telecommun. Networking* **54**(5), 862–876 (2010)
2. A. Wang, M. Iyer, R. Dutta, G. Rouskas, I. Baldine, Network virtualization: technologies, perspectives, and frontiers. *J. Lightwave Technol.* **31**(4), 523–537 (2012)
3. N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Perterson, J. Rexford, S. Shenker, J. Turner, OpenFlow: Enabling innovation in campus networks, in *ACM SIGCOMM Computer Communication Review*, Apr 2008
4. A. Leon-Garcia, NSERC strategic network on smart applications on virtual infrastructure. University of Toronto, Sep 2011. Available: <http://www.savinetwork.ca/wp-content/uploads/AI-Leon-Garcia-SAVI-Introduction.pdf>



<http://www.springer.com/978-3-319-01290-2>

Wireless Virtualization

Wen, H.; Tiwary, P.K.; Le-Ngoc, T.

2013, XIV, 112 p. 42 illus., Softcover

ISBN: 978-3-319-01290-2