



A survey of network virtualization [☆]

N.M. Mosharaf Kabir Chowdhury ^{a,1}, Raouf Boutaba ^{b,c,*}

^a *Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94702, United States*

^b *Cheriton School of Computer Science, University of Waterloo, ON, Canada N2L 3G1*

^c *Division of IT Convergence Engineering, POSTECH, Pohang KB 790-784, Republic of Korea*

ARTICLE INFO

Article history:

Received 3 March 2008

Received in revised form 21 October 2009

Accepted 25 October 2009

Available online 31 October 2009

Responsible Editor: I.F. Akyildiz

Keywords:

Network virtualization

Virtual networks

Next-generation Internet architecture

ABSTRACT

Due to the existence of multiple stakeholders with conflicting goals and policies, alterations to the existing Internet architecture are now limited to simple incremental updates; deployment of any new, radically different technology is next to impossible. To fend off this ossification, *network virtualization* has been propounded as a diversifying attribute of the future inter-networking paradigm. By introducing a plurality of heterogeneous network architectures cohabiting on a shared physical substrate, network virtualization promotes innovations and diversified applications. In this paper, we survey the existing technologies and a wide array of past and state-of-the-art projects on network virtualization followed by a discussion of major challenges in this area.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

The Internet has been stunningly successful over the course of past three decades in supporting multitude of distributed applications and a wide variety of network technologies. However, its popularity has become the biggest impediment to its further growth. Due to its multi-provider nature, adopting a new architecture or modification of the existing one requires consensus among competing stakeholders. As a result, alterations to the Internet architecture have become restricted to simple incremental updates and deployment of new network technologies have become increasingly difficult [1,2].

To fend off this ossification, *network virtualization* has been propounded as a diversifying attribute of the future inter-networking paradigm. Even though architectural purists view network virtualization as a means for evaluating new architectures, the pluralist approach considers virtualization as a fundamental attribute of the architecture itself [1]. They believe that network virtualization can eradicate the *ossifying forces* of the Internet and stimulate innovation [1,2].

1.1. What is network virtualization?

A networking environment supports network virtualization if it allows coexistence of multiple virtual networks on the same physical substrate. Each *virtual network* (VN) in a *network virtualization environment* (NVE) is a collection of virtual nodes and virtual links. Essentially, a virtual network is a subset of the underlying physical network resources.

Network virtualization proposes decoupling of functionalities in a networking environment by separating the role of the traditional Internet Service Providers (ISPs) into two: *infrastructure providers* (InPs), who manage the physical infrastructure, and *service providers* (SPs), who create

[☆] This work was jointly supported by the Natural Science and Engineering Council of Canada (NSERC) under its Discovery program, Cisco Systems, and WCU (World Class University) program through the Korea Science and Engineering Foundation funded by the Ministry of Education, Science and Technology (Project No. R31-2008-000-10100-0).

* Corresponding author. Address: Cheriton School of Computer Science, University of Waterloo, ON, Canada N2L 3G1. Tel.: +1 519 888 4820; fax: +1 519 885 1208.

E-mail addresses: mosharaf@cs.berkeley.edu (N.M.M.K. Chowdhury), rboutaba@cs.uwaterloo.ca (R. Boutaba).

¹ This work was completed when this author was a Master's student at the University of Waterloo.

virtual networks by aggregating resources from multiple infrastructure providers and offer end-to-end network services [2–4].

Specifically, network virtualization is a networking environment that allows multiple service providers to dynamically compose multiple heterogeneous virtual networks that coexist together in isolation from each other. Service providers can deploy and manage customized end-to-end services on those virtual networks for the end users by effectively sharing and utilizing underlying network resources leased from multiple infrastructure providers [4]. Such a dynamic environment will foster deployment of multiple coexisting heterogeneous network architectures without the inherent limitations found in the existing Internet.

However, as a research area network virtualization is mostly unexplored. Several technical challenges in terms of instantiation, operation, and management of virtual networks are either untouched or require further attention. This presents a wide range of theoretical as well as practical open problems and unique challenges. This paper examines the past and the state of the art in network virtualization and identifies key issues for future exploration.

1.2. Organization

The remainder of this paper is composed as follows: in Section 2, we review four existing technologies – virtual local area networks, virtual private networks, active and programmable networks, and overlay networks – that are closely related to the concept of network virtualization. Later in Section 3, we survey a number of past and present projects on network virtualization and related concepts followed by a summarization of the surveyed projects from different perspectives in Section 4. Section 5 identifies key research issues for further exploration based on a qualitative analysis of the surveyed work. We conclude in Section 6.

2. Technologies

The concept of multiple coexisting networks appeared in the networking literature in different capacities. In this section, we discuss four such incarnations: *Virtual Local Area Networks (VLAN)*, *Virtual Private Networks (VPN)*, *active and programmable networks*, and *overlay networks*.

2.1. Virtual local area network

A virtual local area network (VLAN) [5] is a group of hosts with a common interest that are logically brought together under a single broadcast domain regardless of their physical connectivity. Since VLANs are logical entities, i.e., configured in software, they are flexible in terms of network administration, management, and reconfiguration. Moreover, VLANs provide elevated levels of trust, security, and isolation, and they are cost-effective.

Classical VLANs are essentially Layer 2 constructs, even though implementations in different layers do exist. All frames in a VLAN bear a common VLAN ID in their MAC

headers, and VLAN-enabled switches use both the destination MAC address and the VLAN ID to forward frames. This process is known as *frame coloring*. Multiple VLANs on multiple switches can be connected together using *trunking*, which allows information from multiple VLANs to be carried over a single link between switches.

2.2. Virtual private network

A virtual private network (VPN) [6–8] is a dedicated communications network of one or more enterprises that are distributed over multiple sites and connected through tunnels over public communication networks (e.g., the Internet).

Each VPN site contains one or more Customer Edge (CE) devices (e.g., hosts or routers), which are attached to one or more Provider Edge (PE) routers. Normally a VPN is managed and provisioned by a VPN service provider (SP) and known as Provider-provisioned VPN (PPVPN) [9]. While VPN implementations exist in several layers of the network stack, the following three are the most prominent ones.

2.2.1. Layer 3 VPN

Layer 3 VPNs (L3VPN) [10,11] are distinguished by their use of layer 3 protocols (e.g., IP or MPLS) in the VPN backbone to carry data between the distributed CEs. L3VPNs can again be classified into two categories: CE-based and PE-based VPNs.

In the *CE-based VPN* approach, CE devices create, manage, and tear up the tunnels without the knowledge of the SP network. *Tunneling* requires three different protocols:

- (1) *Carrier protocol* (e.g., IP), used by the SP network to carry the VPN packets.
- (2) *Encapsulating protocol*, used to wrap the original data. It can range from very simple wrapper protocols (e.g., GRE [12], PPTP [13], L2TP [14]) to secure protocols (e.g., IPsec [15]).
- (3) *Passenger protocol*, which is the original data in customer networks.

Sender CE devices encapsulate the passenger packets and route them into carrier networks. When the encapsulated packets reach the receiver CE devices at the end of the tunnels, they are extracted and actual packets are injected into receiver networks.

In *PE-based L3VPNs*, the SP knows that certain traffic is VPN traffic and process them accordingly. The VPN states are stored in PE devices, and a connected CE device behaves as if it were connected to a private network.

2.2.2. Layer 2 VPN

Layer 2 VPNs (L2VPNs) [16,17] provide end-to-end layer 2 connection between distributed sites by transporting Layer 2 (typically Ethernet but also ATM and Frame Relay) frames between participating sites. The primary advantage of L2VPN is its support of heterogeneous higher-level protocols. But its lack of a control plane takes away its capability of managing reachability across the VPN.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات