# Virtualized Networks based on System Virtualization

Andreas Fischer, Andreas Berl, and Hermann de Meer

University of Passau, Innstr. 43, 94032 Passau, Germany
{andreas.fischer,andreas.berl,hermann.demeer}@uni-passau.de

## 1 Introduction

The virtualization of networks is not a new idea in network research. *Virtual Local Area Networks (VLANs)* [2] and *Virtual Private Networks (VPNs)* like IPSec [10], are widely used to virtualize links. Also *Overlays* and *Peer-to-Peer (P2P)* networks [11] are a widely used approach to get an abstraction of the physical topology of networks. In projects like PlanetLab [7] or GENI [1] end-hosts which are located all over the world are virtualized. The approach of programmable networks [6] tries to achieve network virtualization by deploying programmable network elements into the core network.

As an alternative, the method of system virtualization can be used to virtualize networks. Currently system virtualization is highly popular to virtualize servers in data centers to consolidate servers. But when this virtualization method is applied to a core network (consisting of routers and links) a new network model emerges from the combination of these technologies.

## 2 Virtualized Networks in the Context of System Virtualization

### 2.1 System Virtualization Background

*System virtualization* is used to virtualize physical hardware, called *Real Machine (RM)* in this context. A *Virtual Machine Monitor* (VMM) runs on top of the RM and virtualizes its resources by providing *Virtual Machines (VMs)*. A VM consists of virtual CPUs, virtual memory, virtual hard disks, virtual network interface cards, etc. A VM is a perfect recreation of a RM, therefore an *Operating System* (OS) can be installed within it (called *Guest OS*). Several VMs can be run in parallel on the same RM without being aware of each other. There are two popular ways of implementing a system virtualization VMM: either directly on the RM (called *full virtualization*) or on top of an OS (called *hosted virtualization*). Full virtualization VMMs are also called *hypervisors*. This kind of virtualization can be found in popular products like XEN [4] or VMWare ESX Server [12].

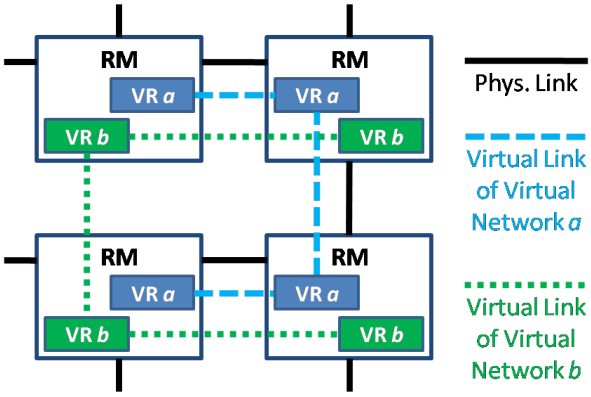


**Fig. 1.** Two independent virtual networks

### 2.2 Virtualized Networks

An existing physical network, consisting of routers and links, can be virtualized by using the method of system virtualization. Routers (that are the RMs in this case) are able to provide several VMs, each containing a Router OS. Such VMs with routing functions are called *Virtual Routers (VRs)* and are interconnected via *Virtual Links (VLs)*. VLs are logical interconnections between two VRs that have dynamically changeable properties. A single VL can be an aggregation of physical links and it can also span over several hops in the real network (tunneling). Together, VRs and VLs form *Virtual Networks* (VNs). A VN is defined as the transitive closure of interconnected VRs (i.e. two VRs that are either directly or indirectly connected belong to the same VN). Several VNs can exist in parallel on top of a physical network.

Figure 1 depicts an example of a virtualized network. Two independent VNs ($a$ and $b$) are driven on top of the physical topology, e.g. an IPv4 and an IPv6 network. It can be seen that not all of the VLs have a direct physical representation. This system virtualization based network model allows to apply management functions on VRs and VLs that were developed for VMs. VRs can easily be created, started, paused, resumed, stopped, or destroyed. Even movement of VRs is possible (both as cold migration by stopping the VR and starting it elsewhere and as live migration by transferring system state during execution of the VR [13]). Finally, VLs are manageable through the VMM [5] - i.e. link parameters like bandwidth can be modified in real time.

## 3 Conclusions and Future Work

When the method of system virtualization is applied to core networks a predetermined kind of network model emerges. This model provides a high level of flexibility in line with expected Future Internet requirements. The basic feasibility of this approach has been examined, however it still has to be worked

out whether current technology is already able to support this scenario under realistic workload conditions. Issues that have been found [8] have to be worked out. Finally, a more formal analysis of the resulting network model has to be performed in order to fully understand all implications.

**Acknowledgments.** Parts of the work in this paper were done in the context of the EU-funded project AutoI [3] and the Network of Excellence EuroNF [9]. This paper was also partly funded by the German Research Foundation (Deutsche Forschungsgemeinschaft - DFG), contract number ME 1703/4-2

## References

1. GENI - Global Environment for Networking Innovations. http://www.geni.net.
2. Virtual bridged local area networks. http://standards.ieee.org/getieee802/download/802.1Q-2003.pdf, May 2003. IEEE Standard 802.1Q.
3. AUTOI. Autonomic Internet Project, STREP, FP7, grant no. ICT-2007-1-216404. http://www.ist-autoi.eu.
4. Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. Xen and the art of virtualization. *SIGOPS Oper. Syst. Rev.*, 37(5):164–177, 2003.
5. Andreas Berl, Andreas Fischer, Hermann De Meer, Alex Galis, and Javier Rubio-Loyola. Management of virtual networks. In *4th IEEE/IFIP International Workshop on End-to-end Virtualization and Grid Management - EVGM2008*, Samos Island, Greece, September 2008.
6. Andrew T. Campbell, Herman G. De Meer, Michael E. Kounavis, Kazuho Miki, John B. Vicente, and Daniel Villela. A survey of programmable networks. *SIGCOMM Comput. Commun. Rev.*, 29(2):7–23, 1999.
7. Brent Chun, David Culler, Timothy Roscoe, Andy Bavier, Larry Peterson, Mike Wawrzoniak, and Mic Bowman. Planetlab: an overlay testbed for broad-coverage services. *SIGCOMM Comput. Commun. Rev.*, 33(3):3–12, 2003.
8. Norbert Egi, Adam Greenhalgh, Mark Handley, Mickael Hoerdt, Laurent Mathy, and Tim Schooley. Evaluating xen for router virtualization. In *16th Int. Conf. on Comp. Commun. and Networks - ICCCN 2007*, pages 1256–1261, Aug. 2007.
9. EuroNF. European Network of the Future, NoE, FP7, grant no. IST-216366, 2007. http://euronf.enst.fr/en_accueil.html.
10. S. Kent and K. Seo. Security architecture for the internet protocol. http://tools.ietf.org/html/rfc4301, Dec. 2005. IETF RFC 4301.
11. Ralf Steinmetz and Klaus Wehrle. *Peer-to-Peer Systems and Applications (Lecture Notes in Computer Science)*. Springer-Verlag New York, Secaucus, NJ, USA, 2005.
12. VMWAre. VMWare ESX - Bare-Metal Hypervisor for Virtual Machines. http://www.vmware.com/products/vi/esx.
13. Yi Wang, Eric Keller, Brian Biskeborn, Jacobus van der Merwe, and Jennifer Rexford. Virtual routers on the move: live router migration as a network-management primitive. *SIGCOMM Comput. Commun. Rev.*, 38(4):231–242, 2008.