UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

Virtual Network Management with XEN

GI/ITG Fachgruppen Betriebssysteme

Herbsttreffen 2008, Garching

Andreas Fischer, Andreas Berl, and Hermann de Meer

Computer Networks & Communications
Prof. Hermann de Meer

Overview

- Motivation
- Classification of Virtualization Techniques
- Virtualization of Networks
  - Definition
  - Benefits
- Virtual Network Management
  - Testbed
  - Usability
  - Security

Motivation

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Today's network layer is too inflexible
  - Slow adoption of new techniques (e.g. DiffServ/IntServ, IPv6)
  - Leads to makeshift solutions (e.g. Network Address Translation)
  - Stopgap measures becoming permanent solutions
  - New services are restricted by current limitations
- We need to overcome ossification of today's Internet
  - Networks need to cater to new services
  - Networks should be dynamically adaptable
- Virtualization of networks can help to overcome these problems

Virtualization Techniques

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Process virtualization
  - Virtualization of resources for a process
  - Process runs on virtual CPU, uses virtual memory, etc.
  - Space sharing, time sharing / multitasking
  - Example: Java VM
- System virtualization
  - Virtualization of full systems (not only userland)
  - OS runs on virtual hardware
    - Virtual CPU, memory, disk, graphic, sound, network interface card…

Virtualization Techniques

> Different approaches of system virtualization

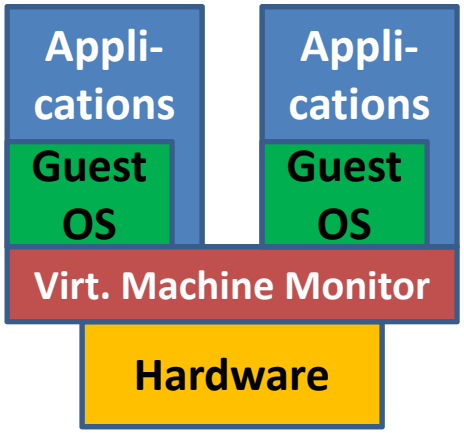

Full Virtualization
(e.g. XEN, VMWare ESX-Server)

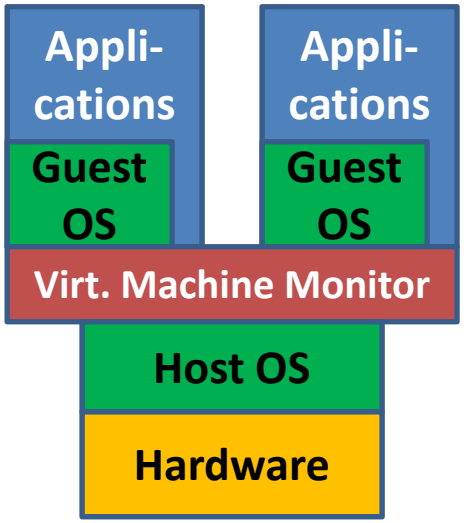

Hosted Virtualization
(e.g. VMWare Workstation)

Virtualization of Networks

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

> Virtual networks (in our view) consist of
  - ***Virtual Routers***
  - ***Virtual Topologies***

> Virtual Routers (VR)
  - Encapsulated in virtual machines
  - Have features of virtual machines

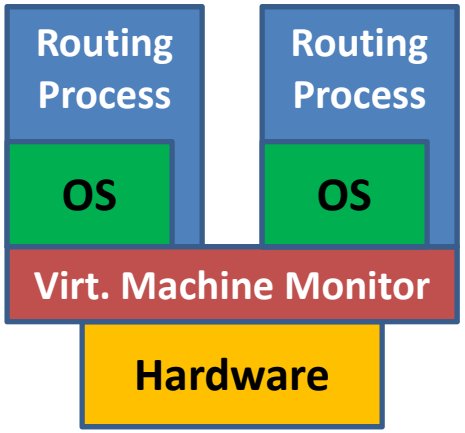

Virtual Routers

Virtualization of Networks

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Virtual Routers based on virtual machines
  - Isolated from each other (sandboxed)
  - Stored in files (images)
  - Easy to handle (Start/Stop, Create/Delete, Copy/Move)
    - Easy Backup / Restore
    - "Roll back" in time possible
  - Live migration (used in data centers)
  - Provision of different software configurations simultaneously
    - E.g. different protocol stacks

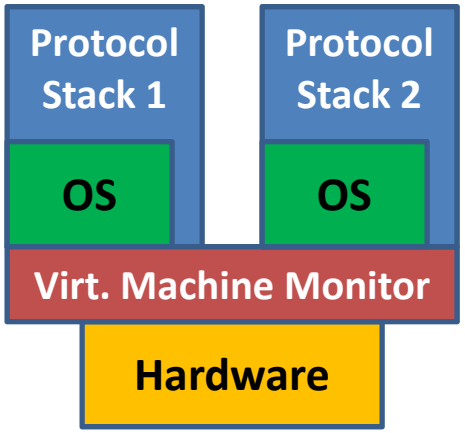

Virtual Routers

Virtualization of Networks

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

> Virtual Topologies
  - Can be different from physical topologies (subset)
  - Multiple different topologies are possible
  - Dynamic change of topology is possible
    - Powering up / Shutting down Virtual Routers
    - Changing link properties

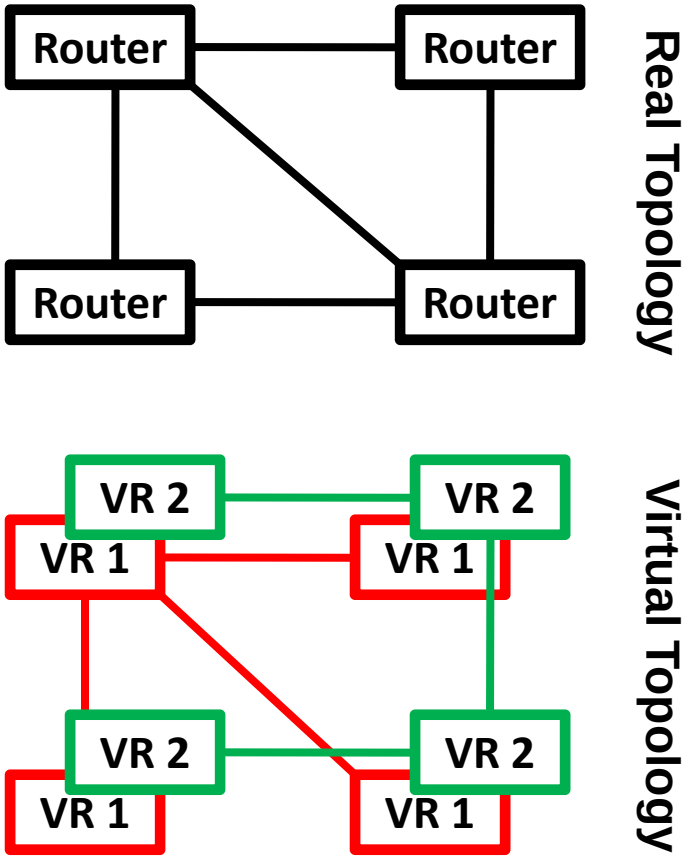

Benefits of Virtual Networks

- Create networks with different characteristics
- Adapt to service demands
  - Optimized topology
  - Adjustable link properties (e.g. bandwidth)
- Dynamic reconfiguration
  - Within hours
  - Network can adapt to changing business rules

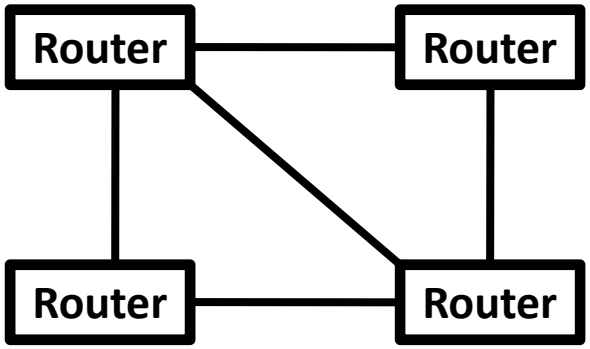

Real Topology

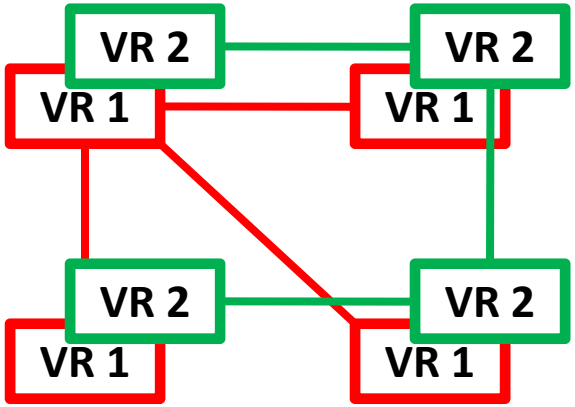

Virtual Topology

Benefits of Virtual Networks

- Encapsulation: Different networks don't interfere with each other
- Use different techniques in parallel
  - E.g. IPv4/IPv6
  - Smooth transition possible
- Add new functionality (IPv8?) without disturbing legacy network

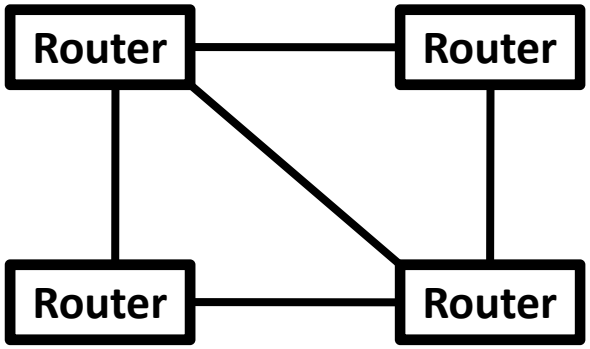

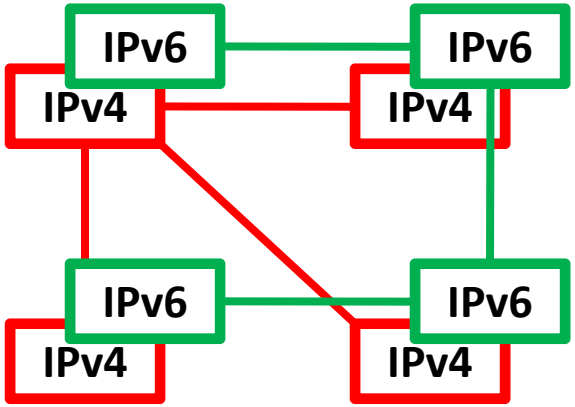

Testbed Implementation

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Selection of virtualization techniques
  - XEN seems to be the appropriate choice to start with
  - XEN implements the concept of paravirtualization
- Paravirtualization supports high performance
  - Guest OS is aware of the virtualization
  - Guest OS performs hypercalls instead of system calls
- Exclusive hardware allocation to virtual machines is possible
  - E.g. to network interface card

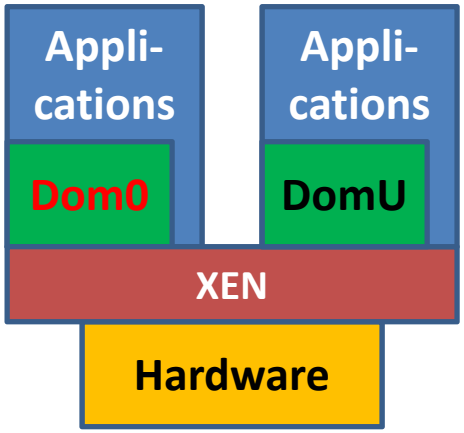

Xen

Testbed Implementation

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- XEN Testbed implemented
  - Try to limit bandwidth of a virtual router
  - Try to give bandwidth guarantees (in face of contention)
- Test results are promising
  - Bandwidth distribution stabilizes within seconds
  - Dynamic reconfiguration is possible

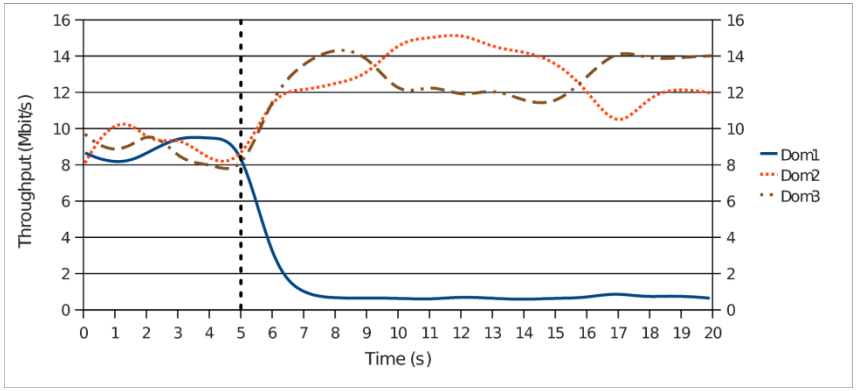

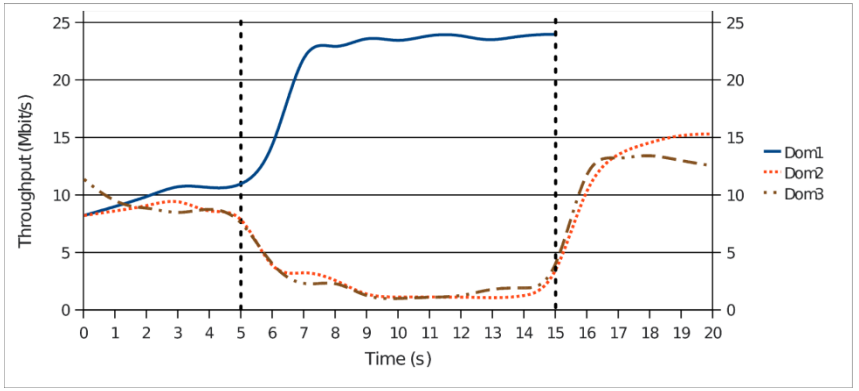

Usability of Virtual Networks

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Virtualization adds complexity
  - Not only real resources to handle but also virtual resources
  - How to manage the additional complexity?
- Our goal: Separate virtualization related problems from other network management problems
  - Provide a “pool of virtual resources”
  - Relieve clients (administrators, network management software) from dealing with real resources
- Virtualization interface is needed for
  - Monitoring the virtual network
  - Managing the virtual network

Usability of Virtual Networks

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Monitoring the virtual network – Available resources
  - Enable reasonable management
    - Decide whether to start a new virtual router or not
    - Perform load-balancing
  - Allow network management to dynamically react
    - To bottlenecks (e.g. by moving the virtual router)
    - To unexpected new user requirements (e.g. by increasing the bandwidth)
- Provide an appropriate abstraction
  - Abstract from specific hardware issues
  - Abstract from hypervisor resource overhead

Usability of Virtual Networks

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Designing management functions
  - Providing reasonable service primitives
    - Modify virtual routers (e.g. start, stop, move…)
    - Modify virtual links (e.g. change bandwidth)
    - Take into account work done by the DMTF
  - Grouping high level methods oriented at specific tasks
    - Example – group into Performance-, Fault-, Topology-Management
    - Allows clients to concentrate on a specific aspect
- Determine how to identify a virtual router
  - Identificator/Locator problem – which router is where?

Security in Virtual Networks

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Attacks using virtualization techniques have been published
- Virtual machine based rootkits may become a relevant threat
  - (Nearly) unlimited power for the attacker
  - Really hard to detect if everything is virtualized
- Clear access definitions/restrictions needed
  - Who is allowed to manage a virtual router? Its creator? Its host? Its users?
  - Who is allowed (and under what circumstances) to create new virtual routers?
  - Who is allowed to read monitoring data?
  - Too often security is an afterthought - don't repeat that mistake

Open Issues

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

- Examination of more use cases
- Verification of the applicability of XEN under realistic network conditions (high load, spiked traffic)
- Definition of an appropriate interface
  - Finding the right granularity of management and monitoring functions – low complexity, high functionality
- Determination of security requirements
  - Access rights to management functions
  - Privacy issues with monitoring values

Questions & Answers

UNIVERSITÄT PASSAU
Fakultät für Informatik und Mathematik

Thanks for your attention.

Any Questions?

Contact:

andreas.fischer@uni-passau.de

andreas.berl@uni-passau.de

hermann.demeer@uni-passau.de

BTW: According to Wikipedia, „Pils“ seems to be a special type of „Helles“

References

1. Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I.,Warfield, A.: Xen and the art of virtualization. In: SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles. (October 2003)
2. Popek, G.J., Goldberg, R.P.: Formal requirements for virtualizable third generation architectures. Communications of the ACM 17(7) (July 1974)
3. Berl, A., Fischer, A., Meer, H.D., Galis, A., Rubio-Loyola, J.: Management of virtual networks. In: 4th IEEE/IFIP International Workshop on End-to-end Virtualization and Grid Management - EVGM2008, Samos Island, Greece (September 2008)
4. Davy, S., Fahy, C., Griffin, L., Boudjemil, Z., Berl, A., Fischer, A., Meer, H.D., Strassner, J.: Towards a policy-based autonomic virtual network to support differentiated security services. In: TEMU 2008 - International Conference on Telecommunications & Multimedia, Ierapetra, Crete, Greece (July 2008)
5. Fahy, C., Davy, S., Boudjemil, Z., Van der Meer, S., Rubio-Loyola, J., Serrat, J., Strassner, J., Berl, A., De Meer, H., Macedo, D.: An information model that supports service-aware self-managing virtual resources. In: 3rd IEEE International Workshop on Modelling Autonomic Communications Environments - MACE2008, Samos Island, Greece (September 2008)
6. A. T. Campbell, H. G. De Meer, M. E. Kounavis, K. Miki, J. Vicente and D. A. Villela, A Survey of Programmable Networks, *ACM Computer Communications Review*, Vol. 29, No. 2, pp. 7-24, April 1999.
7. A. T. Campbell, H G. De Meer, M. E. Kounavis , K. Miki, J. Vicente, and D. A. Villela, "The Genesis Kernel: A Virtual Network Operating System for Spawning Network Architectures", *2nd IEEE International Conference on Open Architectures and Network Programming (OPENARCH'99)*, New York March 26-27 1999.
8. System virtualization profile (August 2007) DMTF profile DSP1042, version 1.0.0a, http://www.dmtf.org/standards/published documents/DSP1042.pdf.
9. Virtual system profile (May 2007) DMTF profile DSP1057, version 1.0.0a, http://www.dmtf.org/standards/published documents/DSP1057.pdf.

References

10. Bassi, A., Denazis, S., Galis, A., Fahy, C., Serrano, M., Serrat, J.: Autonomic Internet: A Perspective for Future Internet Services Based on Autonomic Principles. In: IEEE 3rd International Week on Management of Networks and Services End-to-End Virtualization of Networks and Services - Manweek 2007 / MACE 2007 2nd IEEE International Workshop on Modelling Autonomic Communications Environments, San José, California, USA (October 2007)
11. Autonomic Internet (AutoI) (2007) EU FP7 IST Project, http://ist-autoi.eu/.
12. Future Generation Internet (EuroFGI) (2006) Network of Excellence, FP6, grant no. 028022, http://eurongi.enst.fr/.
13. European Network of the Future (EuroNF) (2007) Network of Excellence, FP7, grant no. 216366, http://euronf.enst.fr/.
14. Rutkowska, J.: Subverting vista kernel for fun and profit. In: Black Hat 2006, Las Vegas, Nevada, USA (August 2006)
15. King, S.T., Chen, P.M., min Wang, Y., Verbowski, C., Wang, H.J., Lorch, J.R.: Subvirt: Implementing malware with virtual machines. In: IEEE Symposium on Security and Privacy. (May 2006)
16. Egi, N., Greenhalgh, A., Handley, M., Hoerdt, M., Mathy, L., Schooley, T.: Evaluating Xen for Router Virtualization. In: 16th International Conference on Computer Communications and Networks - ICCCN 2007. (August 2007) 1256–1261