Virtualization for Future Internet

2010.02.23

Korea University

Chuck Yoo

(hxy@os.korea.ac.kr)

Why Virtualization

- Internet today
  - Pro and con
- Your "wonderful" research results
  - Mostly with simulation
  - Deployment impossible
- Virtualization
  - Aim to set you free
  - Thus, center piece to future internet

Life before Virtualization

- Testbeds
  - Production: Internet2
  - Research: DETER
    - Mainly incremental changes
- Overlay
  - PlanetLab
    - Centered around specific problem to solve
    - No interaction between overlays

- NO major architectural innovation

Contents

- Virtualization introduction
- Related work in Virtualization
  - Xen
  - Embedded System Virtualization
- Network Virtualization
  - Router Virtualization
  - XEN for Router Virtualization
- Summary & Reference

What is Virtualization?

- General meaning
  - An abstraction of resources that provides a logical rather than an actual physical incarnation of those resources
  - Such resources are: CPU, storage, network, etc.
- Introduced to fully utilize mainframes in 1960s
  - Could run several projects on a single mainframe
- Has been extended to several areas
  - System Virtualization → Consumer Electronics
  - Storage Virtualization → Cloud Computing
  - Network Virtualization → Future Internet

Virtualization Features

- **Decoupling**: minimizes dependency between software and hardware
  → Provides S/W portability

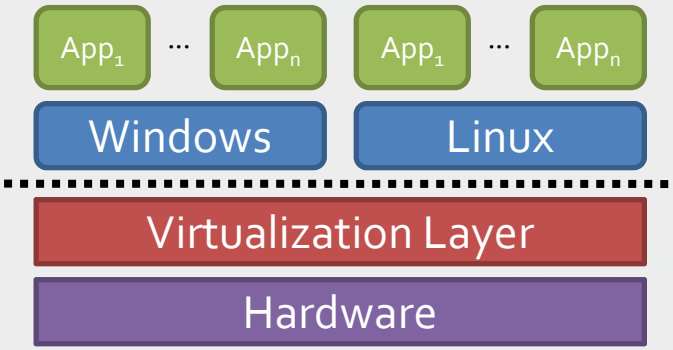

- **Isolation**: Each OS uses own virtual resources independently
  → Provides reliability and fault-tolerance

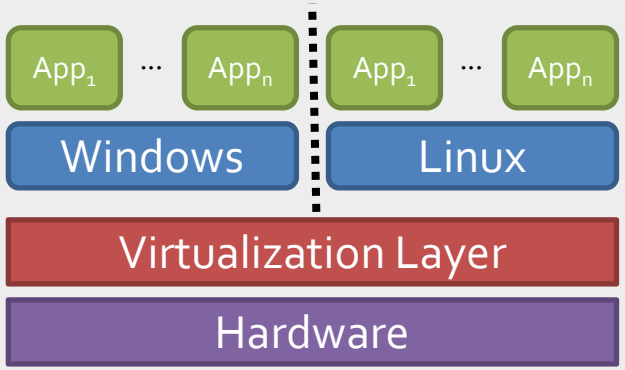

Virtualization Features

- **Consolidation**: consolidates multiple similar machines into one machine → Reduces H/W cost

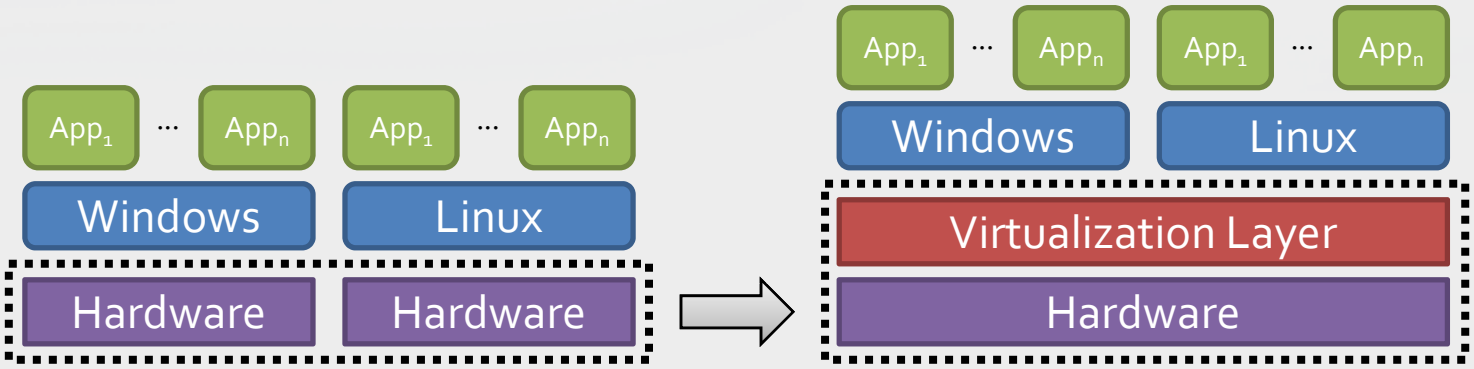

Virtualization Architecture

- Virtual Machine (VM)
  - Self-contained software execution environment
  - Guest OS runs over a virtual machine
- Virtual Machine Monitor (VMM)
  - Supervises virtual machines
  - Interacts with physical hardware

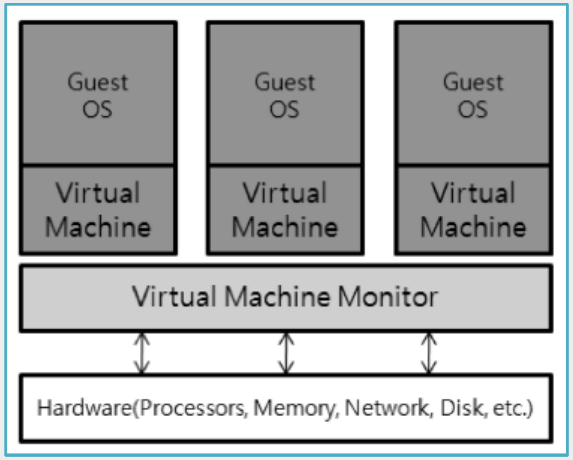

Related work in Virtualization

- Xen
- Embedded System Virtualization

Recent Virtualization(Xen)

- Xen – Hypervisor
  - University of Cambridge
  - Para-virtualization
    - modifies partially guest OS to improve Xen performance
  - **Hypercalls** are defined as a virtual machine interface
  - Split-driver model
    - device drivers are at a special domain (Domain0) to improve system reliability

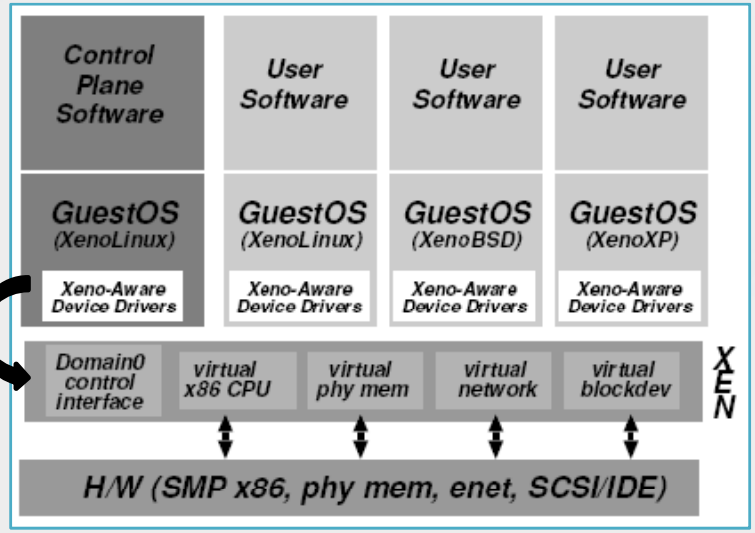

Xen: Interrupt Handling (Network, I/O)

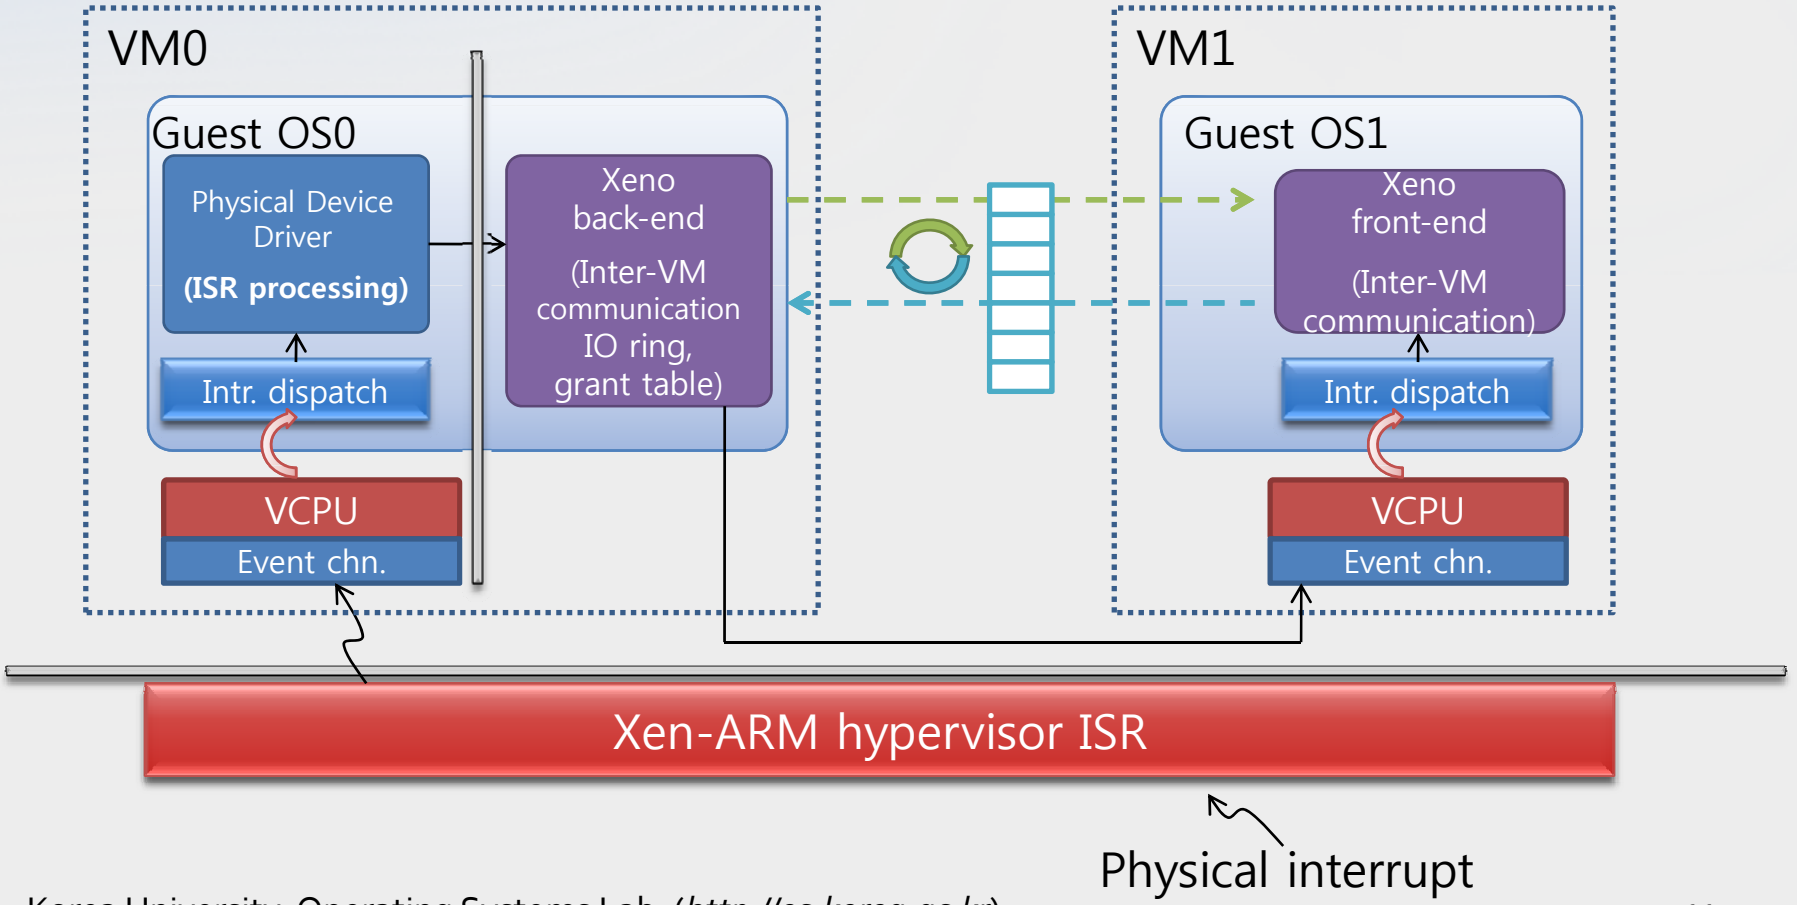

Xen: TCP performance for GigE today

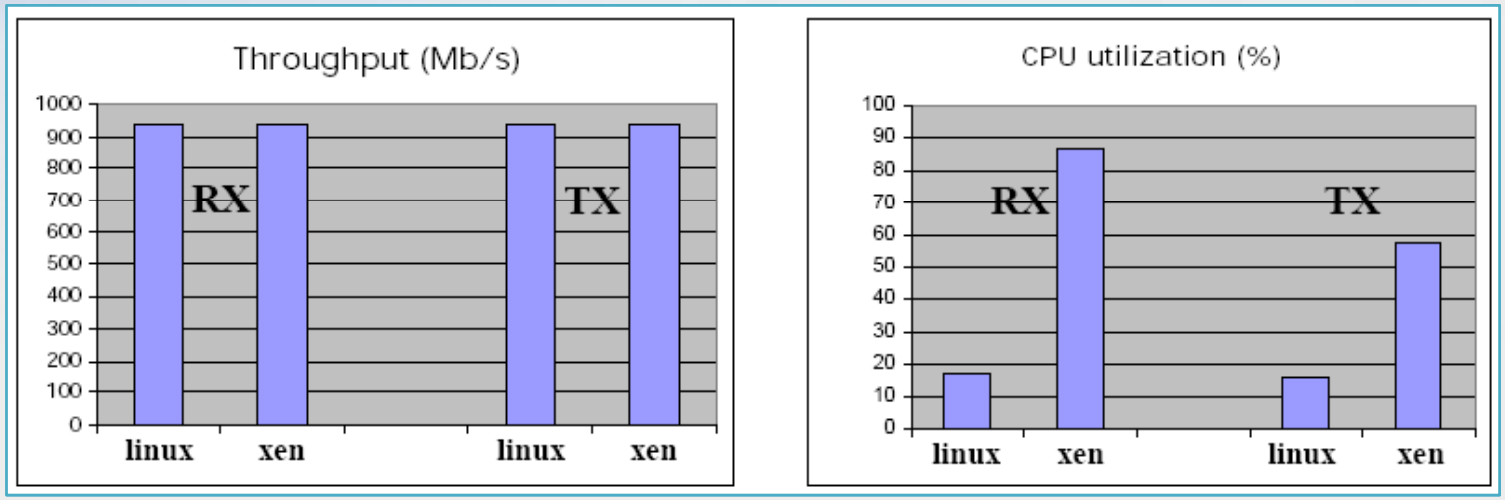

* Nov. 2007, Xen summit (HP)

- Xen split driver can sustain peak throughput on GigE
- But Xen uses significantly more CPU cycles than Linux
  - Less available cycles for application
  - 10 Gig networks: CPU saturation prevents achieving line rate
- Need to reduce I/O virtualization overhead in Xen networking

Embedded System Virtualization

- Currently hot research topic in virtualization
  - New requirements
    - Real-time support, Multi-core support, Power efficiency, etc.

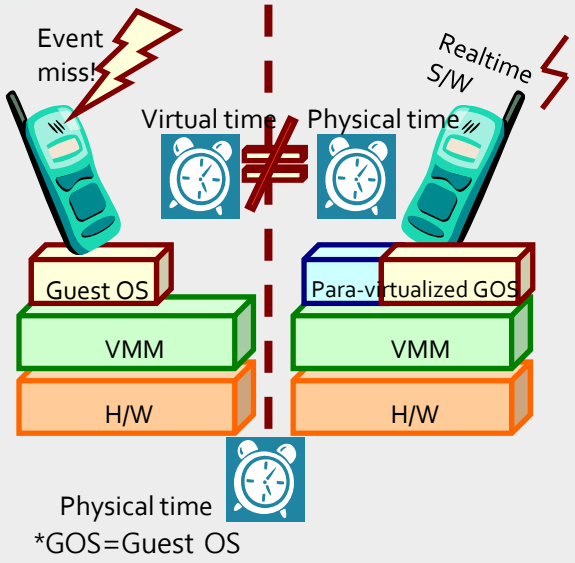

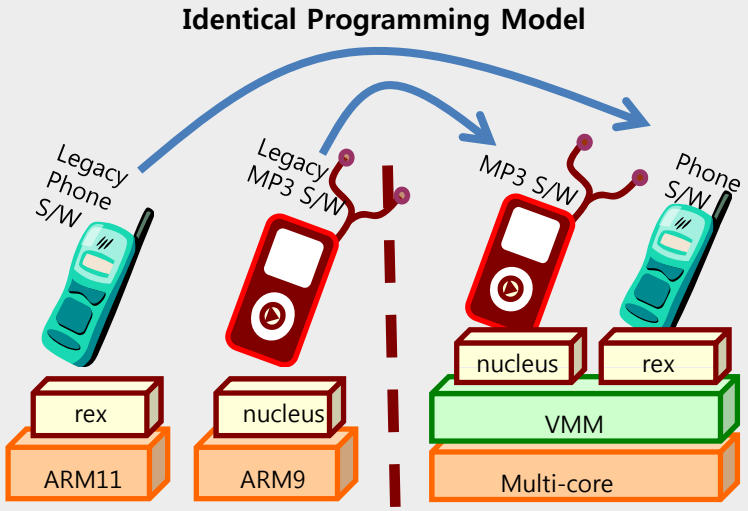

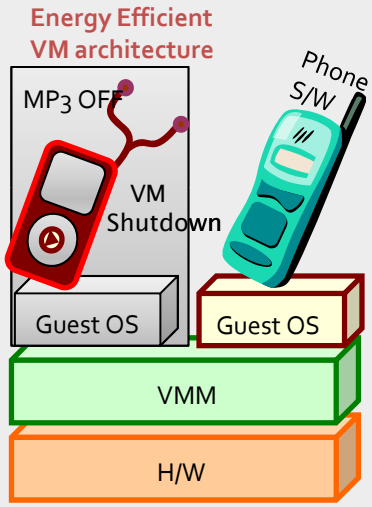

Embedded System Virtualization Topics

- Xen on ARM
  - Samsung Electronics S/W Lab.
  - Embedded S/W Platform Virtualization
- Real-time Xen
  - Supports real-time for embedded mobile devices based on Xen on ARM
  - Provides real-time scheduling between guest OSes
  - Supports user RTOS
  - Considers power efficiency

Embedded System Virtualization Example

- MobiVMM (MobiVirt'08)
  - New VMM for mobile phones
  - Supports mobile multi-core processors such as ARM11 MPCore
  - Provides multi-core load balancing by modified credit scheduler

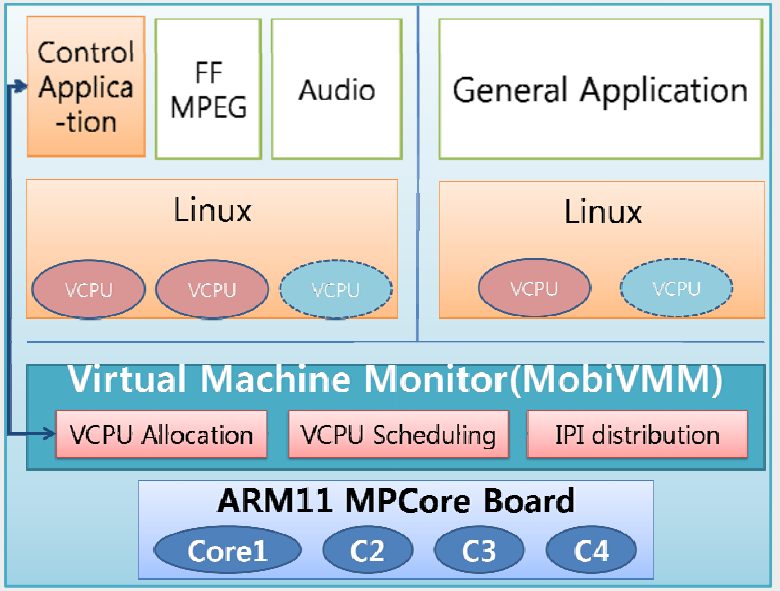

Embedded System Virtualization Korea University status

- MobiVMM
  - Run two guest OSes with MobiVMM over multi-core environments - Demo is available at http://www.youtube.com/watch?v=SABYIon9vOg
  - Feb. 2009 Xen Summit presentation
    - http://www.xen.org/xensummit/xensummit_winter_2009.html
- Real-time
  - Para-virtualized uC/OS-II (RTOS) for Xen on ARM contributed to Xen community
    - Source code is available at http://wiki.xensource.com/xenwiki/XenARM

Network Virtualization

- Router Virtualization

Router Virtualization

- Network virtualization environment

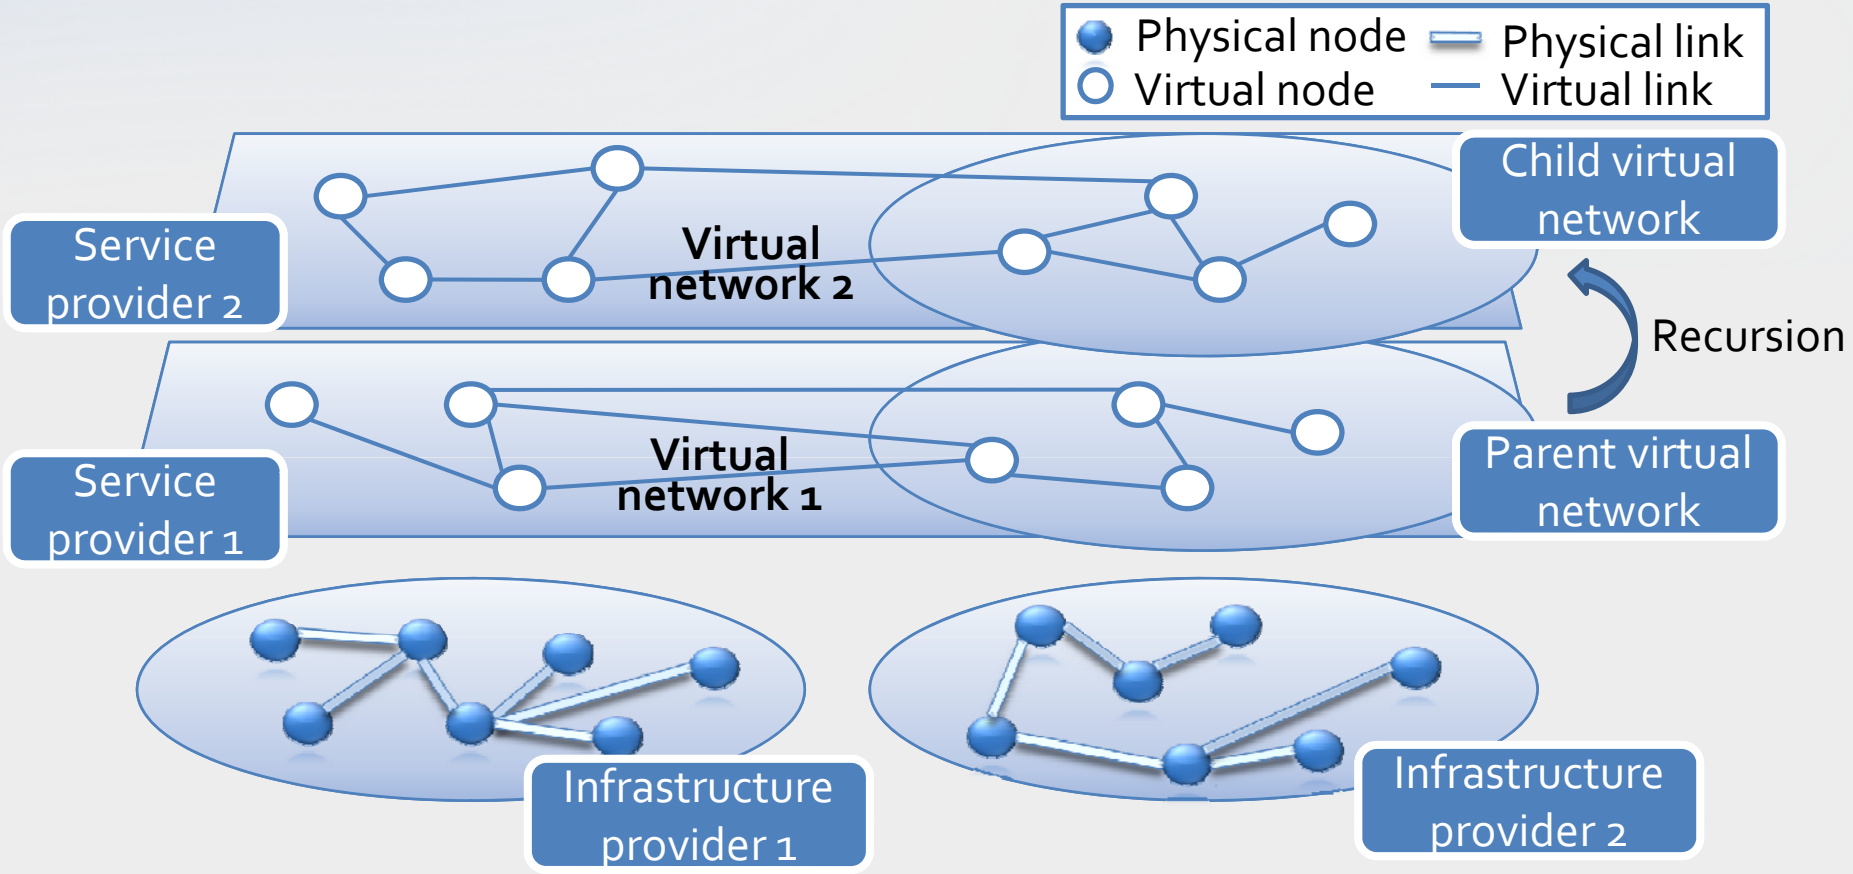

Router Virtualization

- Architecture

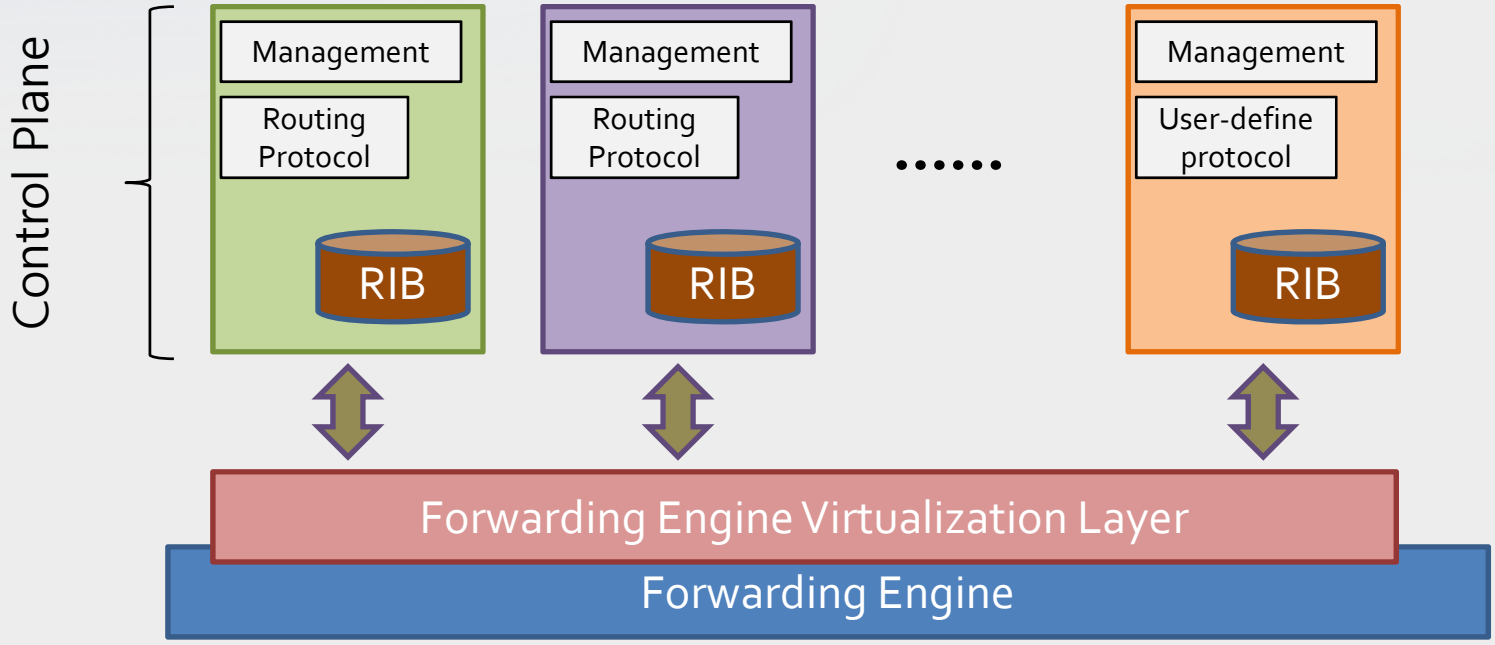

*RIB : Routing Information Base

Router Virtualization

- Advantages
  - Multiple concurrent network architectures
    - Innovation in network architecture design
    - Decoupling the providers from the physical network
  - Excellent platform for experimentation
    - Rolling out new and unstable solutions without risk
  - Independent and flexible management
  - Resource sharing
  - Lower hardware and support cost

Router Virtualization

- Requirements
  - High performance
  - Flexibility, Modularity, Scalability and Programmability
  - Isolation provided by the each virtual machine
  - Efficient resource allocation
  - Minimized virtualization overhead

Router Virtualization

- H/W based router virtualization(NetFPGA, OpenFlow)
  - Feature
    - Limited service deployment & number of virtual machine
    - Almost no effect of performance degradation by virtualization
    - Limited network flexibility due to H/W
      - Difficulty of application to different protocol policy
      - Limitations of scalability

Router Virtualization

- S/W based router virtualization(Linux-VServer, Xen, Open Vz)
  - Feature
    - Low forwarding performance
    - Flexibility, Scalability better than H/W based virtual router
      - Can be applied to various protocol policy
    - Isolation issues between virtual machines
    - Fairness issues in software virtual routers
      - Static, Configurable, Customizable virtual forwarding plane

Router Virtualization

- S/W based router virtualization

1. Container based

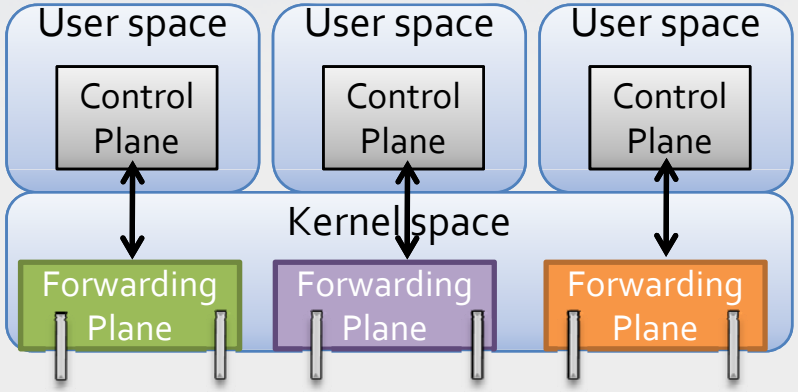

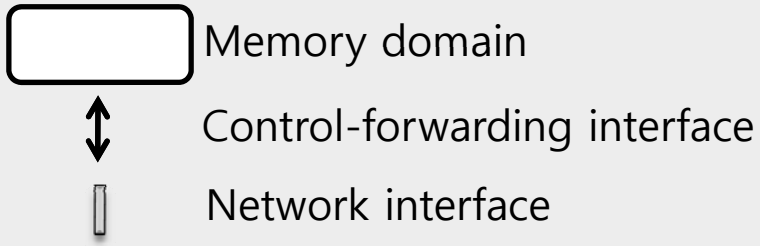

2. Hypervisor based

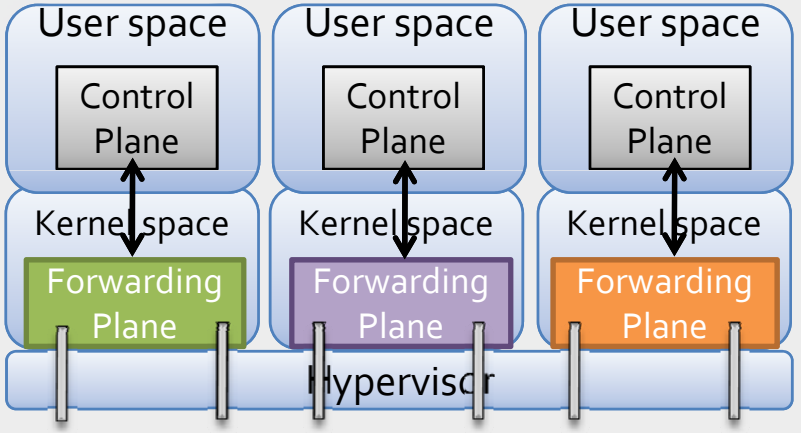

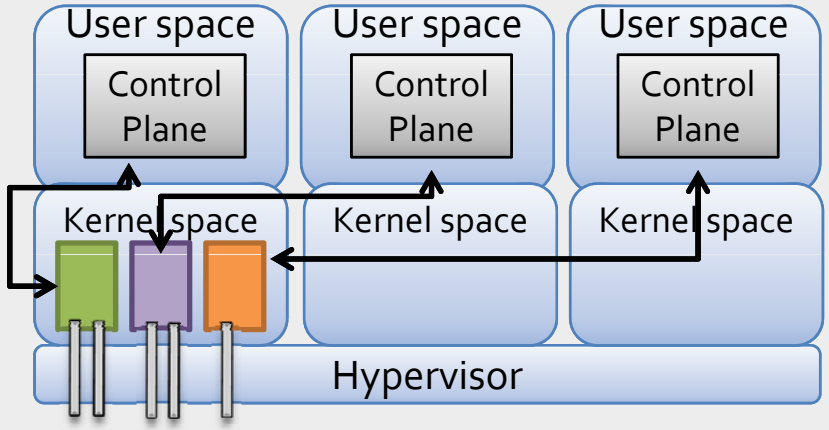

Router Virtualization

- S/W based router virtualization
  - Container based Vs. Hypervisor based(Split forwarding plane)

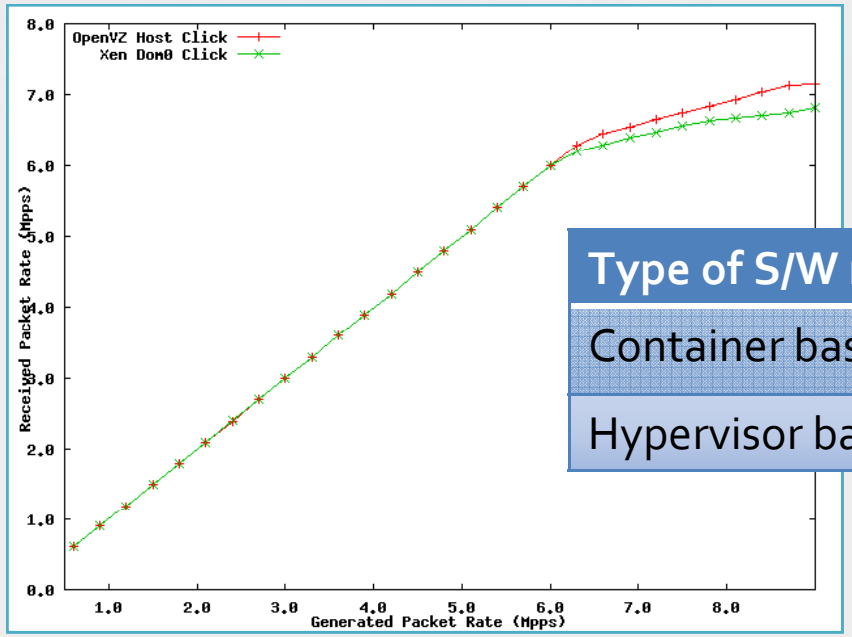

Type of S/W router	Performance
Container based(Open Vz)	7.1 Mpps
Hypervisor based(Xen)	6.8Mpps(4% overhead)

Reference: LANCASTER Univ. Computing department

Router Virtualization

- Summary

Type		Merit	Demerit
H/W		▶ Performance close to native Linux	▶ Difficulty of flexibility, scalability
S/W	Container	▶ Excellent flexibility, scalability	▶ Difficulty of isolation between VMs ▶ Low routing performance (Better than hypervisor type)
	Hypervisor	▶ Isolation, reliability	▶ Low routing performance

Summary

- Virtualization technology
  - Starts from system virtualization
  - Has been extended to embedded system, storage, network virtualization, etc.
- Virtualization for router
  - Hardware and software approaches
  - need a better approach
    - Toward dynamic and reconfigurable virtual network

Reference

[1] J. Smith and R. Nair, "Virtual Machines: Versatile Platforms for Systems and Processes (The Morgan Kaufmann Series in Computer Architecture and Design)", USA: Morgan Kaufmann Publishers inc., 2005.

[2] G. J. Popek and R. P. Goldberg, "Formal requirements for virtualizable third generation architectures", Commun. ACM, vol. 17, no. 7. pp. 412-421, 1974.

[3] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the Art of Virtualization", Proc. ACM SOSP, 2003.

[4] J. Sugerman, G. Venkitachalam, and B.-H. Lim, "Virtualizing I/O devices on vmware workstation's hosted virtual machine monitor", Proc. USENIX ATC, 2002.

[5] N. M. Mosharaf Kabir Chowdhury, and R. Boutaba, "Network Virtualization: State of the Art and Research Challenges", IEEE Communications Magazine, vol. 47, no. 7, pp. 20-26, 2009.

[6] X. Chen, Z. M. Mao, and J. V. Merwe, "ShadowNet: A Platform for Rapid and Safe Network Evolution", Proc. USENIX ATC, 2009.

[7] Norbert Egi, Adam Greenhalgh, Mark Handley, "Towards high performance virtual routers on commodity hardware", ACM CoNEXT, 2008.12

[8] Norbert Egi, Adam Greenhalgh, Mickael Hoerdt, "A Platform for High Performance and Flexible Virtual Routers on Commodity Hardware", SIGCOMM 2009, 2009. 8

감사합니다.

Operating Systems Lab.

Korea University

(http://os.korea.ac.kr)