VMM-Level Distributed Transparency Provisioning Using Cloud Infrastructure Technology

Mahsa Najafzadeh, Hadi Salimi, Mohsen Sharifi and Ali Hamidi

Click to edit Master subtitle style

Distributed Systems Laboratory,
School of Computer Engineering,
Iran University of Science and Technology,
Presented in SVM 2010,
Presenter: Hadi Salimi
hsalimi@iust.ac.ir

Agenda

- Introduction to Virtualization Technology (VT)
- Virtualized Distributed Environments
- Cloud Computing
- Cloud Infrastructure Technology (CIT)
- CIT Research Challenges
- Available CIT Solutions
  - OpenNebula
  - Nimbus
  - Eucalyptus
  - Amazon Elastic Cloud Computing (EC2)
  - Xenoserver

Introduction of Virtualization

- Displaying a real system as a different virtual system or even as multiple virtual systems(multiplex)[SMI05]
- Three main advantages of VT [UHL05]:
  - Isolation
  - Consolidation
  - Migration

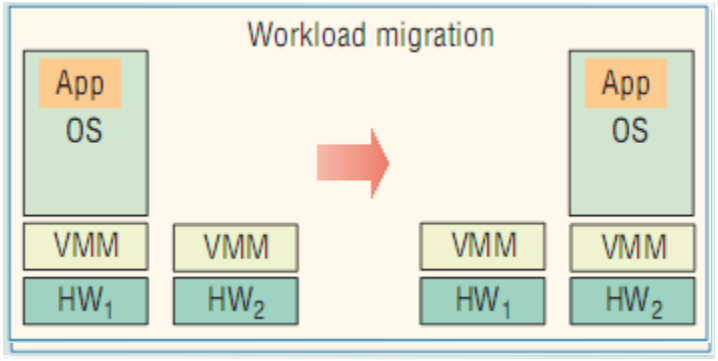

Virtual Distributed Environments

- **A distributed environment:** A collection of independent computers that appears to users as a single coherent system [TAN02]
- **Virtual distributed environment:** Using virtualization technology in distributed environment such as Cloud Computing
- The need of the research and industry communities to complex computing systems in large scales

Cloud Computing

- No common definition exists yet [VOA09]
- Our definition of Cloud Computing is:

"A distributed virtual environment in which a large set of virtualized computing resources, different infrastructures, various development platforms and useful software are delivered as a service to customers as a pay-as-you-go and transparent manner usually over the Internet"

Cloud Infrastructure Technology (CIT)

- A software layer that can provide system services in support of virtual machines such as scheduling, management and so on

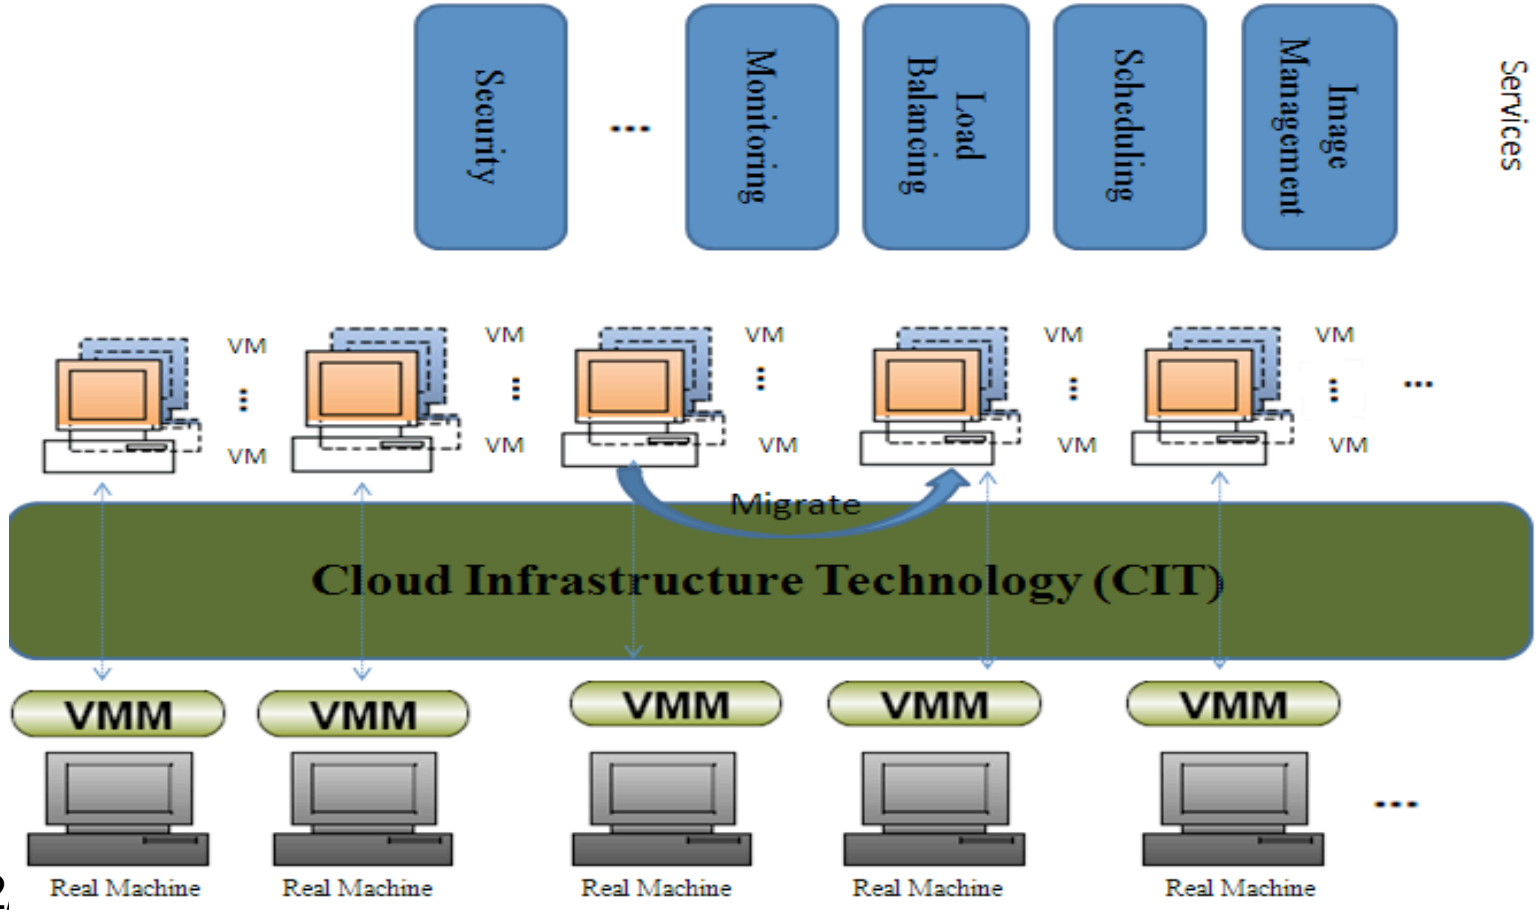

CIT Research Challenges

- Load Balancing [ZHA09]

The need to a specific CIT to be able to schedule virtual machines on physical workstations to balance the workload:

- Image Management [CAI09]

The support mechanisms in CIT for managing the virtual machine images effectively and transmitting them in a virtualized distributed environment efficiently

- Security [YUN10]

The CIT has to manage user demands and the related security issues

CIT Research Challenges (Cont'd)

- Service-Level Agreements [SAK08]

A true commercial CIT must support SLAs in order to make a tradeoff between customer objectives and low level system issues such as computing costs

- Autonomic Scalability [CHI09]

Implementation of proper mechanisms in CIT layer for making infrastructures with the ability of scaling up or down automatically in face of the change of demands

- Energy Efficiency [ABD09]

Tackling energy-awareness issues in virtualized environments by implementation of effective mechanisms in the CIT layer

Available Cloud Infrastructure Technology Solutions

- OpenNebula [BAL09]
- Nimbus [KEA08]
- Eucalyptus [NUR09]
- Amazon Elastic Compute Cloud (EC2) [VAR08]
- Xenoserver[KOT05]

OpenNebula

- An advanced software framework for Cloud computing that supports Xen, KVM and VMware virtualization platforms
- Enable the dynamic placement of groups of interconnected virtual machines on distributed infrastructures

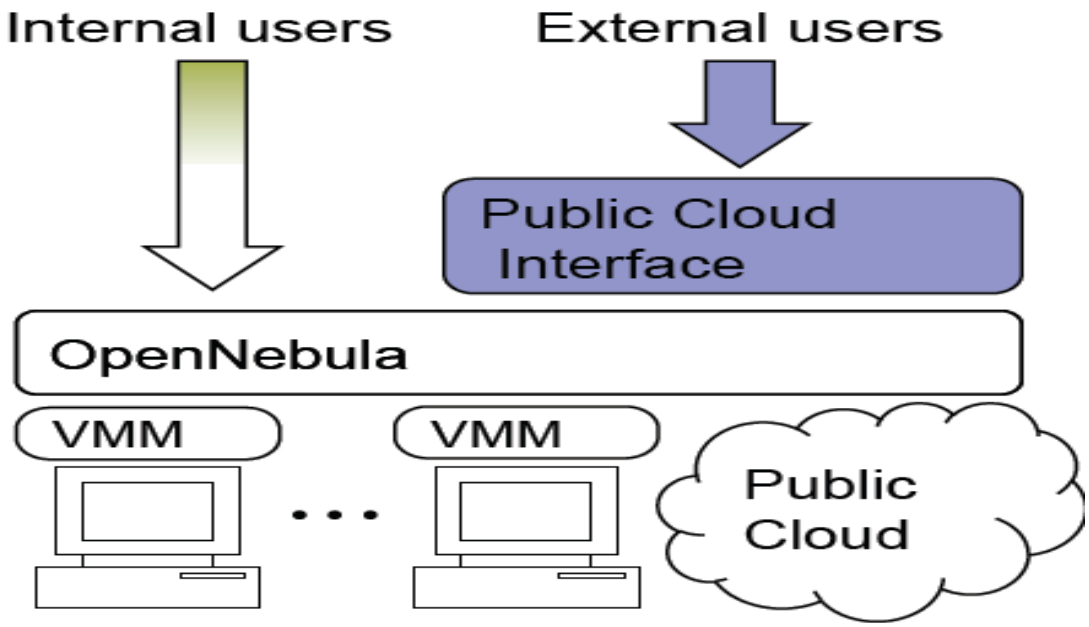

dsa-research.org

Nimbus

- A CIT that is placed on the virtual infrastructure and provides the ability of dynamic management of virtual machines in the virtualized distributed environments
- Nimbus supports Xen and KVM virtualization technologies and acts similar to OpenNebula in many aspects such as creating and managing of VMs
- Handling the related security issues

Eucalyptus

- An open source software layer that implements the infrastructure as a Cloud
- Eucalyptus's architecture is modular, simple, and hierarchical

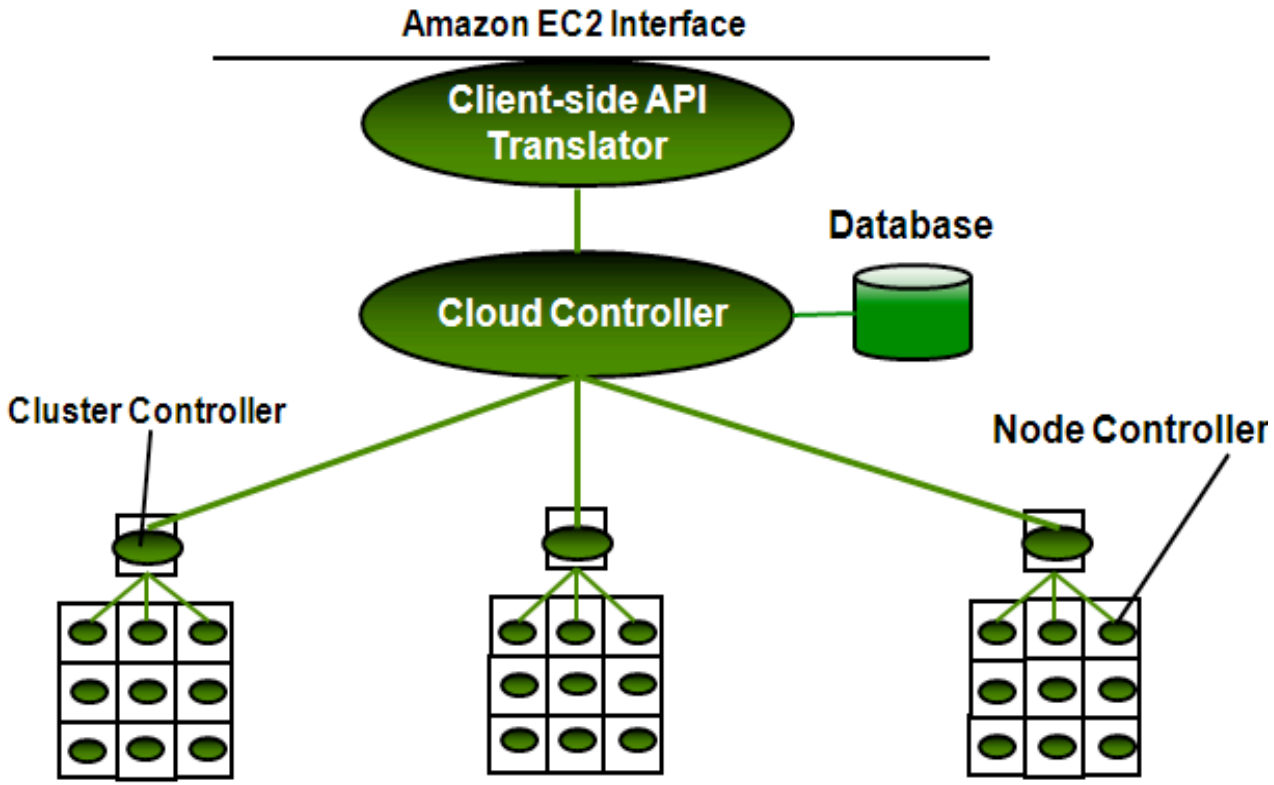

Amazon Elastic Compute Cloud (EC2)

- One of the Amazon web services that provides a resizable capacity in Cloud computing environment
- EC2 makes a business CIT for Cloud computing environments
- Its main objective is to provide the needed capacity of customers on demand

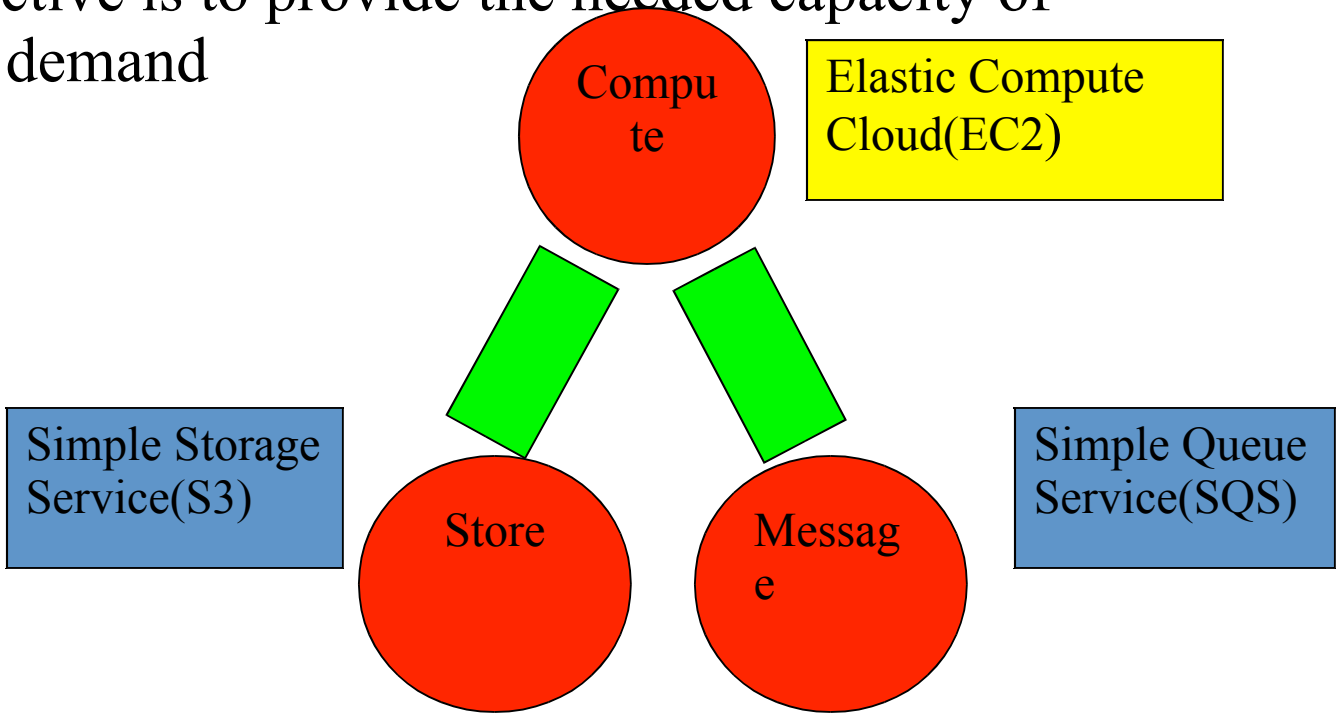

Xenoserver

ξένος

- A CIT that provides a form of transparency in Global Public Computing model (any user, any code, anywhere)
- Global Public Computing is the same as Cloud Computing

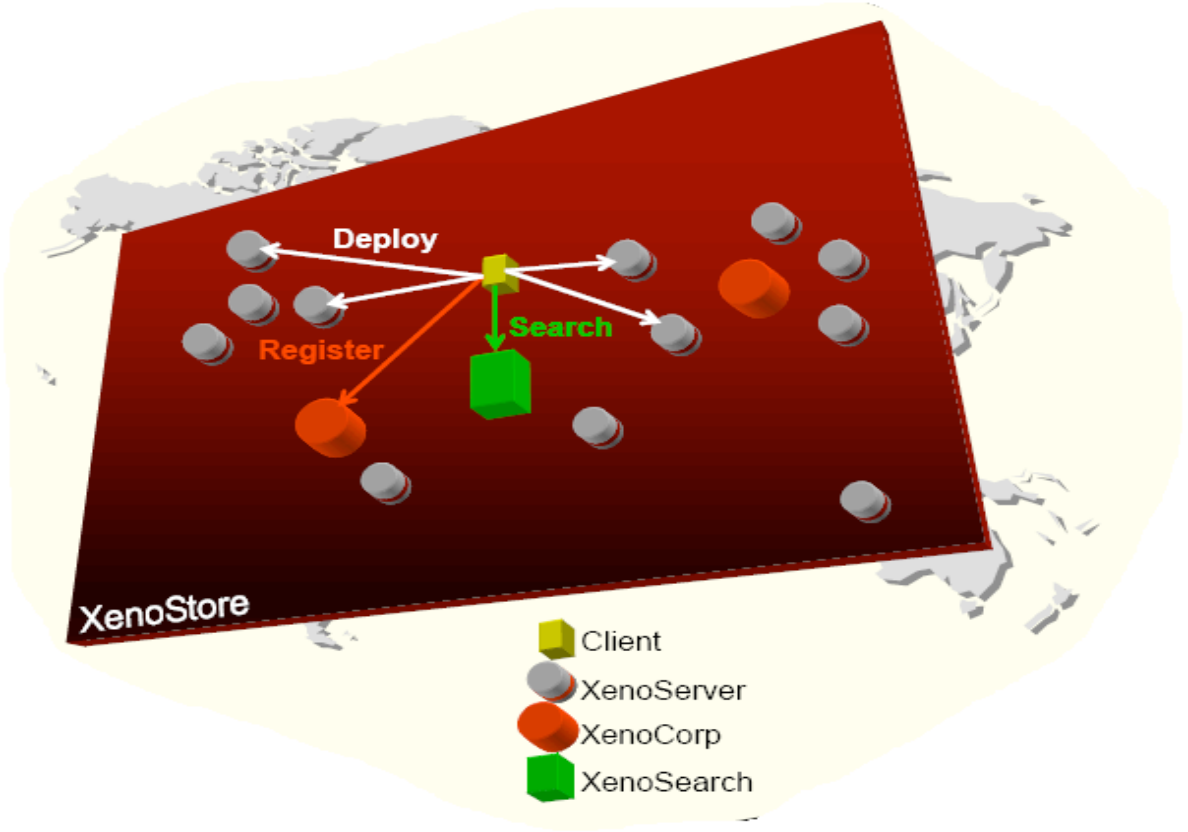

Case Study

- Provisioning Packet Compression Network Service for Virtualized Execution Environments:
  - Transparent
    - No changes in applications
    - No compilation or configuration of operating system
  - Dynamically configurable
    - Administrator can (de)activate it easily
    - Can be configured to serve selected virtual machines

Virtualization Technology

- Virtual Machine Monitor (VMM)
  - Has direct access to physical resources
  - Multiplex physical resources to virtual machines
- Implementation of a Service to change the way virtual machines access to resources inside VMM

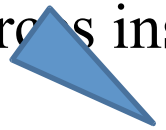

Transparent Service

- Transparent services need no modifications to applications or operating systems.

Implementation

- Xen
  - Open Source
  - High Performance
- In a Xen-based virtualized environment Dom0 has direct access to physical resources.

Implementation (Cont')

- Typical networking path of an application running inside DomU in a Xen-based environment

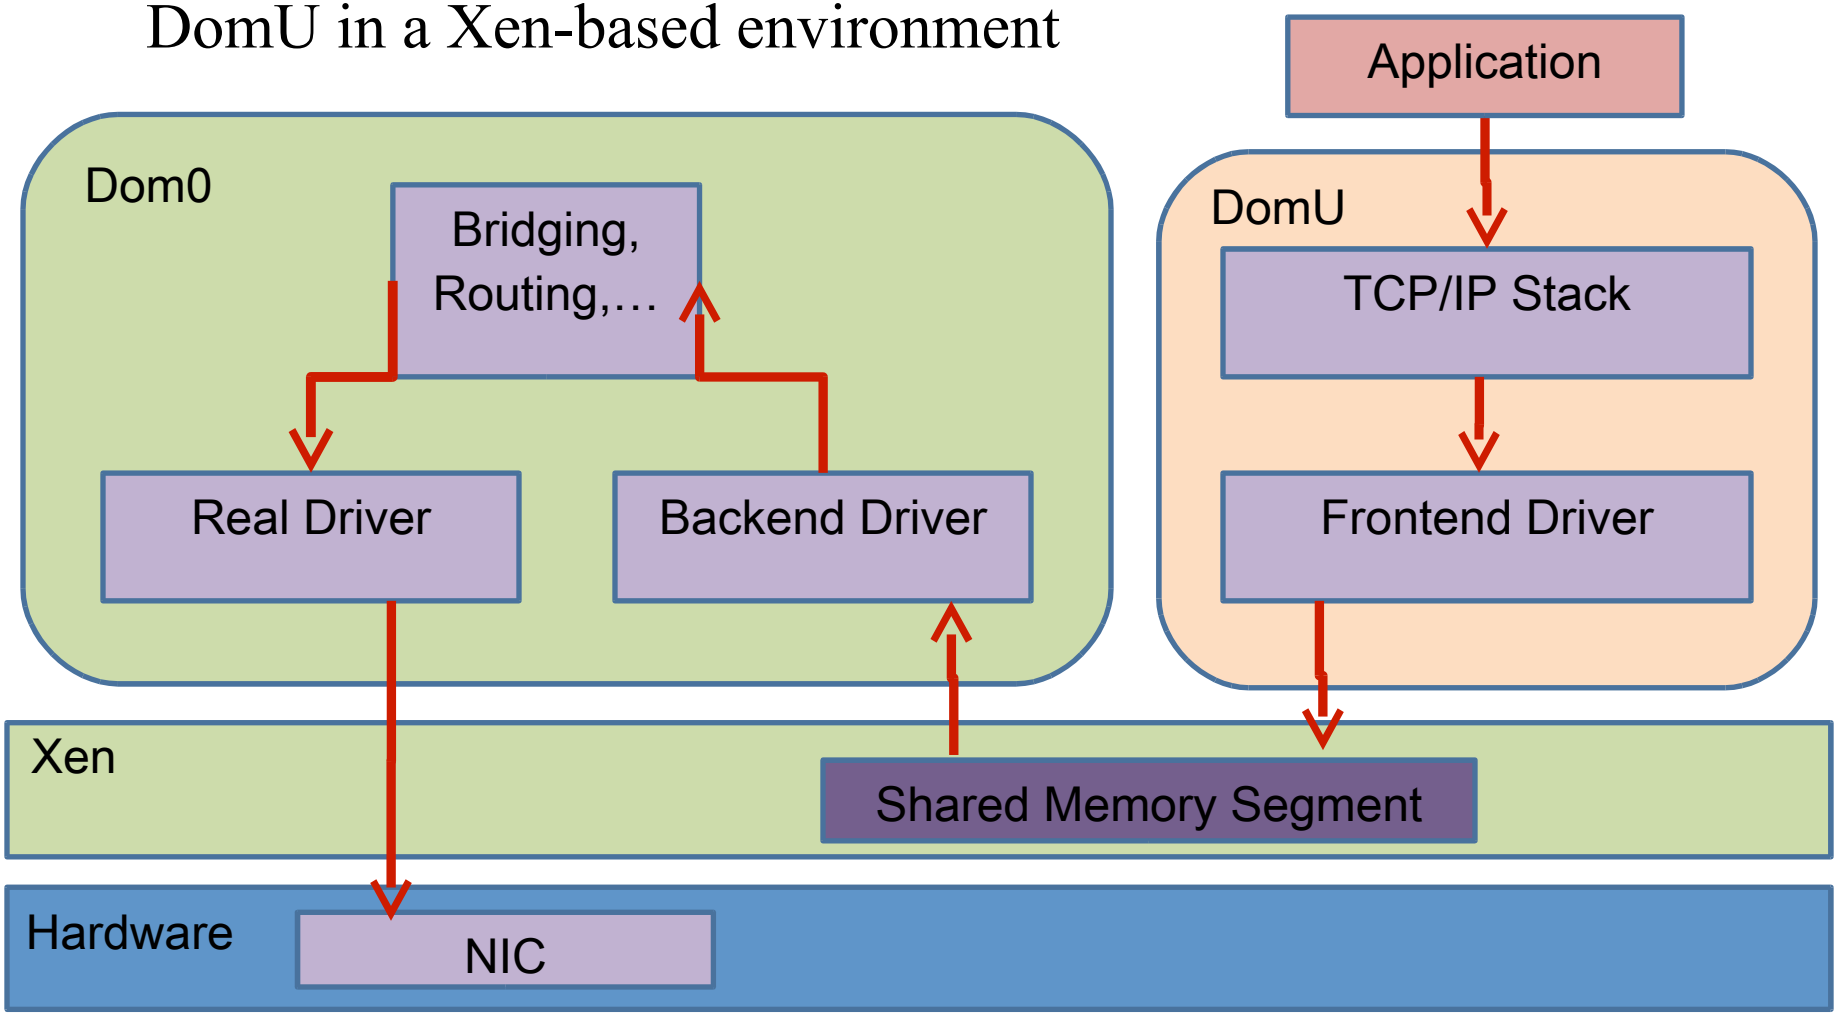

Implementation (Cont’)

- Packet Compression Network Service has been Implemented as a Netfilter Kernel Module inside Dom0.
  - Transparent = Needs no change in applications and guest operating systems
  - Can be dynamically inserted inside Dom0 kernel.
  - Provides a file-based configuration system using Proc filesystem.
  - Administrators can (de)activate it for specific virtual machines.

Implementation (Cont')

- Packet Compression Network Service inserted inside Dom0 kernel

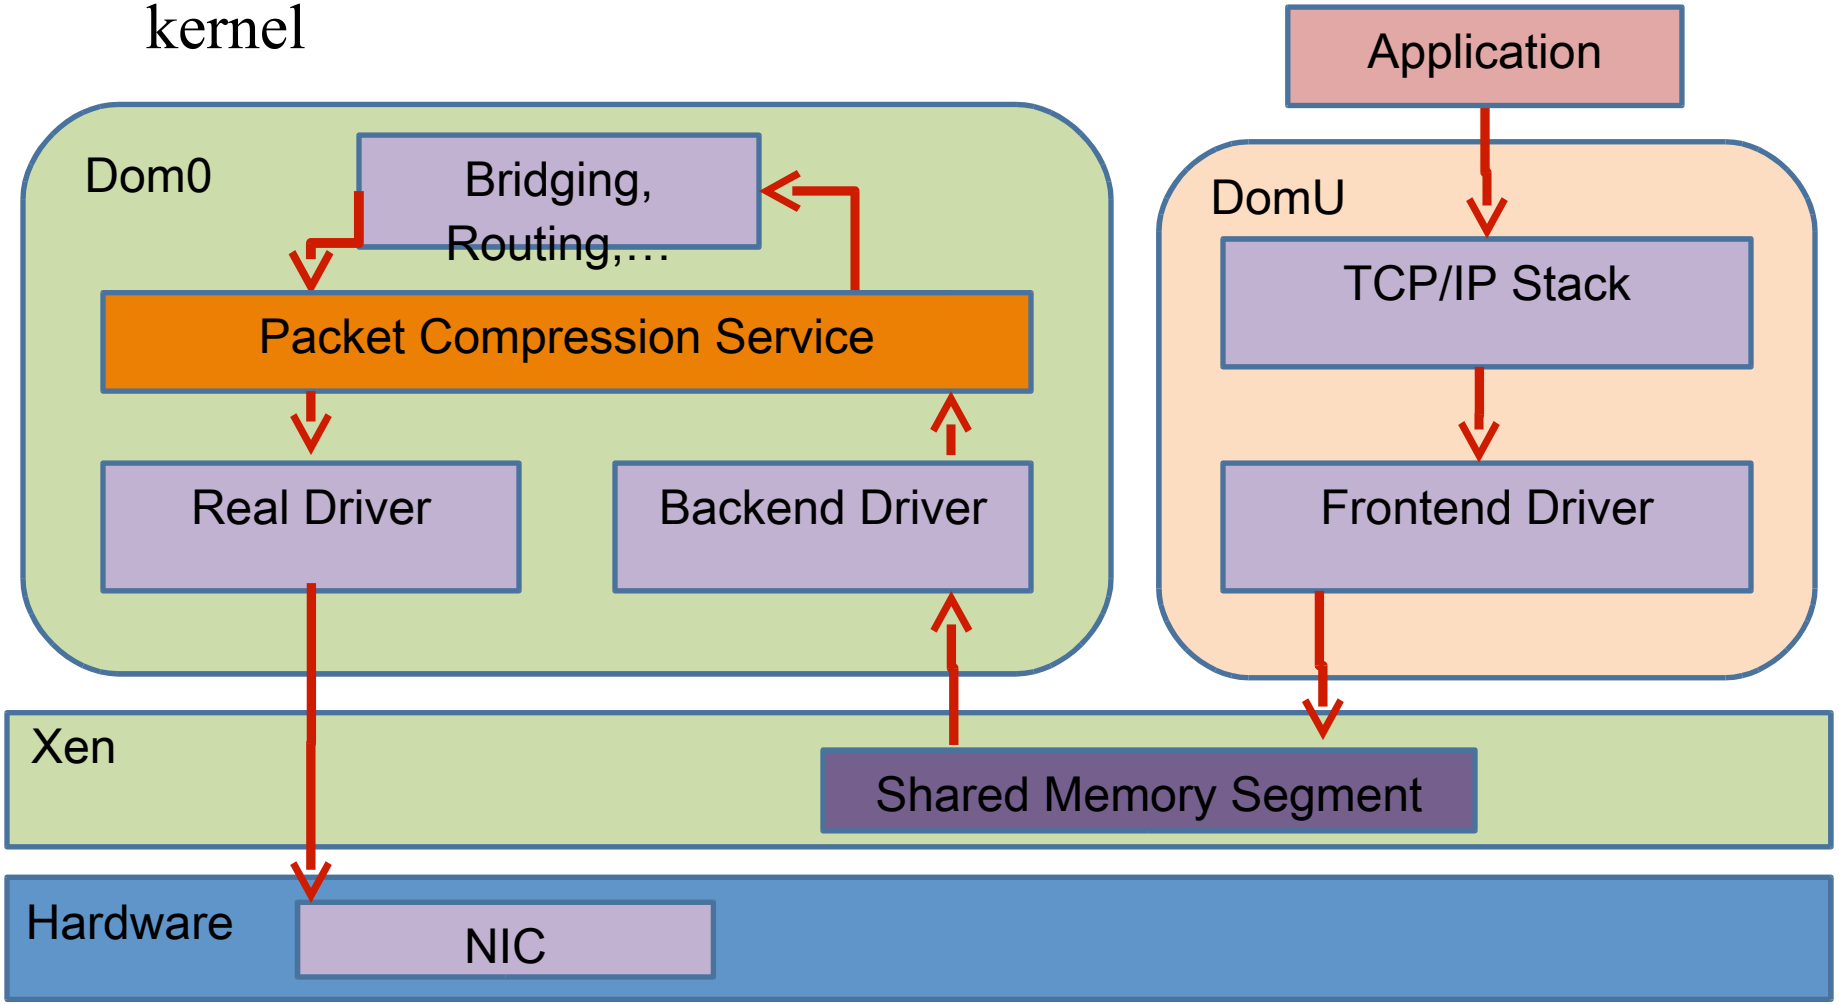

Implementation (Cont')

- Packet Compression Network Service inserted inside Dom0 kernel

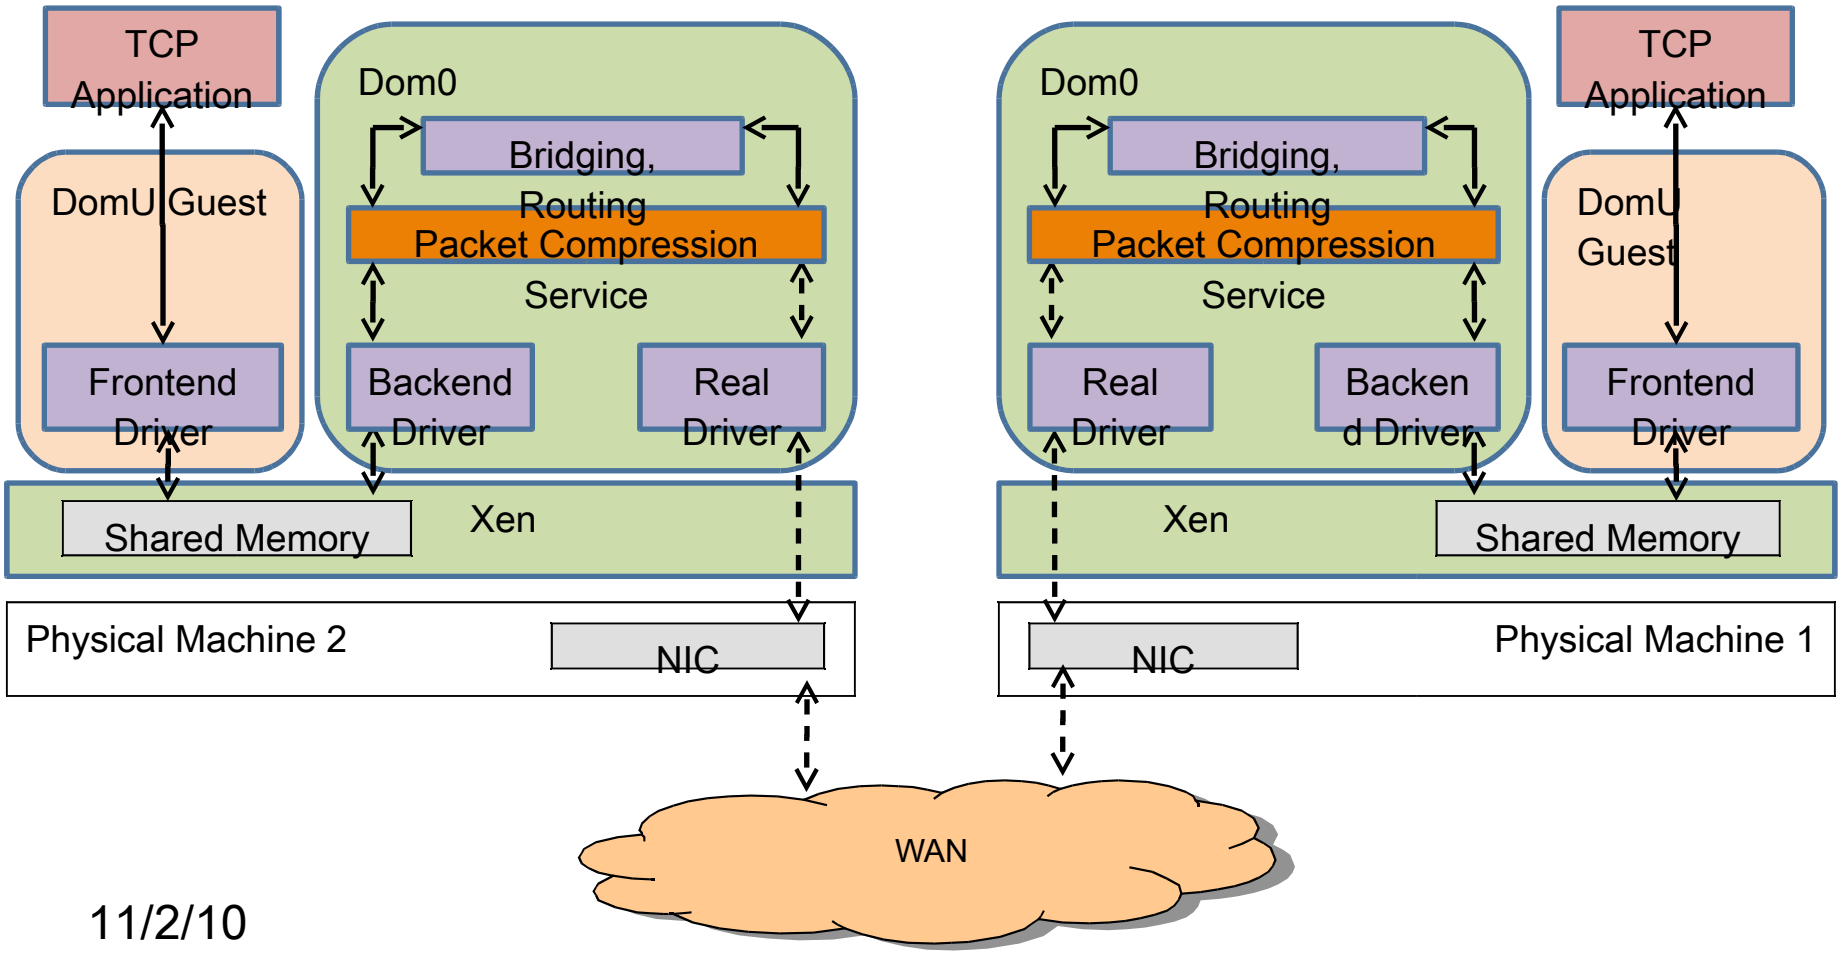

Evaluation

- Network performance evaluation results using Iperf benchmark

Conclusion

- Cloud Computing as a distributed virtual environment.
- Clouds needs a software layer to provide virtual machines with different transparency types.
- As a case study, was presented the design and implementation of a transparent network service on Xen.
- This service can be operable without any

References

- [SMI05] J. E. Smith and R. Nair, "The Architecture of Virtual Machines", IEEE Computer Magazine, May 2005.
- [UHL05] R. Uhlig, G. Neiger, D. Rodgers *et al.*, “Intel virtualization technology,” COMPUTER, vol. 38, no. 5, pp. 48-56, 2005.
- [TAN02] A.S Tanenbaum, M.V Steen and P. Hall, “Distributed Systems: Principles and Paradigms”, ISBN 0130888931,2002.
- [VOA09] J. Voas and J. Zhang, “Cloud Computing: New Wine or Just a New Bottle?”, Published by the IEEE Computer Society, 2009.
- [BAL09] C. V. Blanco, "The OpenNebula Virtual Infrastructure Manager", Xen Directions Europe 2009 joint with LinuxTag, Berlin, Germany, June, 2009.
- [KEA08] K. Keahey, T. Freeman, “Contextualization: Providing One-Click Virtual Clusters”, eScience 2008, Indianapolis, IN. 2008.

References (Cont'd)

- [NUR09] D. Nurmi, R. Wolski, C. Grzegorczyk, G. Obertelli, S. Soman, L. Youseff, D. Zagorodnov, "The Eucalyptus Open-Source Cloud-Computing System", International Symposium on Cluster Computing and the Grid, China, 2009.
- [VAR08] J. Varia,"Cloud Architecture", J. Varia, "Cloud architectures", Technical report, available at http://aws.typepad.com/aws/2008/07/white-paper-on.html
- [KOT05] E. Kotsovinos, "Global Public Computing", PhD dissertation, available as Computer Laboratory Technical Report UCAM-CL-TR-615, ISSN 1476-2986,2005.
- [ZHA09] Y. Zhao, and W. Huang, "Adaptive Distributed Load Balancing Algorithm based on Live Migration of Virtual Machines in Cloud," 5th International Joint Conference on INC, IMS and IDC, NCM, Seoul, Korea,2009.
- [CAI09] G. Cai-dong, Z. Yu, and W. Jian-ping, "Image Knowledge Management and Rapid Precise Image-mining Technology Investigation in

References (Cont'd)

- [YUN10] Y. Jung, and M. Chung, "Adaptive Security Management Model in the Cloud Computing Environment," 12th International Conference on Advanced Communication Technology, ICACT, Gangwon-Do, South Korea, 2010.
- [SAK08]R. Sakellariou, and V. Yarmolenko, "Job Scheduling on the Grid: Towards SLA-Based Scheduling," In L. Grandinetti, editor, High Performance Computing and Grids in Action pages 207–222, 2008.
- [ABD09] H. S. Abdelsalam, K. Maly, R. Mukkamala, and D. Kaminsky, "Analysis of Energy Efficiency in Cloud," Computation-World, 416-421, 2009.
- [CHI09]T. C. Chieu, A. Mohindra, A. A. Karve, and A. Segal, "Dynamic Scaling of Web Applications in a Virtualized Cloud Computing Environment,"International Conference on e-business, ICEBE, Macau, China, 2009.

Any Questions?