A Case Study of Network Virtualization: Implement Multistage Software Routers into Virtual Environment

Nanfang Li, Andrea Bianco, Luca Giraudo
Dipartimento di Elettronica, Politecnico di Torino, Italy
Email: {firstname.lastname}@polito.it

***Abstract*—Network virtualization is a promising technology that offers attractive properties to both academe and industry, like flexibility, isolation, prototyping new architecture under same infrastructure and cost-effectiveness. In this research report, we focus on the key element of virtualized network: router. To be precise, a virtual multistage software router (VMSR) has been introduced. The benefit for this architecture are twofold: multistage software routers overcome open source routers (OSR) running on commodity personal computers' (PC) limitation like low performance and limited interface; while virtualization technologies may permit to make a further step towards improved flexibility by making easier the integration of new components as they are required. Furthermore, new features based on VM management (e.g. consolidation for energy saving) may be easily enabled. In this report we analyze the proposed multistage architecture and consider its implementation when using virtual machines as internal components. Our experiments demonstrate the feasibility of the studied virtual architecture, and discuss some issues related to performance and architecture control.**

I. Virtual Multistage Software Router

Network virtualization has been recently promoted as a solution to the main issues of current Internet like security, flexibility etc. The appealing cost effective multiple virtual networks compose of virtual links and routers together. This report deals with the virtual router, we try to build a multistage software router (MSR) from VMs. The MSR is intend to overcome single PC based router limitation by exploiting a multistage switching architecture. The most benefit of this architecture is improved performance and scalability issues. Furthermore, implementation of recovery mechanisms into the management plane can increase router resilience to close the gap with carrier-grade routers.

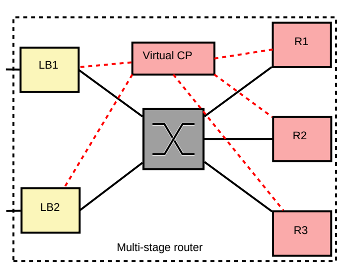

Fig. 1. Example of the multi-stage router composed by two load balancers and three back-end routers: all internal elements run on a different PC to improve performance and reliability.

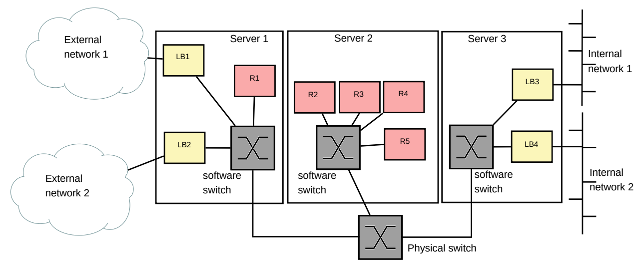

Fig. 2. Use case for introduction of virtualization in multistage router architecture: three physical servers hosts different virtual machines which are used to build up one or more virtualized multistage routers.

The multistage router defined in [1] is organized in three stages (Fig. 1), characterized by three different internal elements: first stage load balancers (LB), second stage interconnecting switches and third stage back-end routers. Interested people are invited to read [1] for a detailed explanation of the functionalities of each stage. This report work only consider to use virtual machine (VM) as the internal element when building MSR since three main advantages can be highlighted when introducing virtualization:

- **larger scalability**: new internal elements can be deployed in a seamless way when traffic increases or more interfaces are needed. This enables renting of resources from data center servers, for example when new VMs are needed to add forwarding capacity.
- **easier management and reliability**: migration of VMs during maintenance periods can be implemented and faster reaction to failures should be expected by booting new VMs on general purpose servers.
- **slicing**: sharing of the same physical infrastructure among different multistage routers possibly dedicated to different types of traffic (e.g., logical separation of the operational and of the experimental networks).

As an example, we report in Fig. 2 a use case referring to an enterprise router based on the multistage architecture under study. External and internal network connections are terminated into the computing server farm, where virtual machines act as load balancers, switches or back-end routers. This solution permits i) to locate virtual machines on different physical servers to upgrade the overall routing capacity; ii) to share the same physical server among several virtual machines to increase resource utilization; iii) to build the

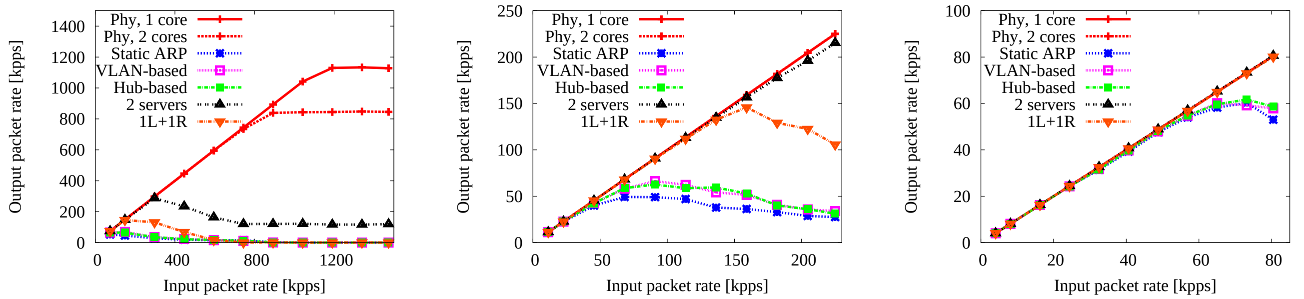

Fig. 3. Performance evaluation of multistage router (2 LB + 2 OSR) into VMware ESXi 4.0 environment with different internal configurations: from left to right, 64, 512 and 1500 bytes packets.

multistage architecture in a mixed approach exploiting both virtual and physical elements; iv) to deploy consolidation mechanisms, e.g. to move all virtual routers to a physical server and turn off unused servers during low traffic periods.

Due to the space limitation, the implementation details has been omitted and the experimental results related to the separate elements like load balancer and router have also been excluded, I will present the test results when we run the whole architecture below, interested readers can refer to [2] for details.

II. Implementation and Forwarding Performance Analysis

The goal of the work is to test the feasibility of building the multistage architecture in a virtual environment, where all internal elements are running on VMs instead of physical PCs. We consider Click Modular router as the LB, with our customer element we configure the frond-end LBs send traffic according to pre-defined rules like round-robin- or table based. The router is a VM installed with XORP routing platform. Our own DIST protocol inside XORP coordinates among the back-end routers to ensure a unique routing table. The central stage switch is a software switch implemented by different virtual infrastructure. We considered XEN and VMware before, but due to the lack of space, only VMware is reported now.

We test a full multistage router architecture considering three simple scenarios:

1) one physical server, one LB and one R
2) one physical server, two LBs and two Rs
3) two physical servers, one LB and one R per server

Results are reported in Fig. 3. Throughput is rather poor for small packet size. This is expected due to the performance limitations of single elements induced by virtualization and resource sharing (e.g. frequent interruptions to run different VMs, context switching and execution status restoration). Indeed in the first case (one LB and one R) we obtain good performance for large packets only (almost wire speed when approaching the maximum packet size). In the second case, where more elements (two LBs and two Rs) share the same resources, the CPU becomes the bottleneck and the upper performance limit is around 70 kpps, not enough to reach wire speed even for large packet size.

Running the multistage in two physical servers, the third case, leads to performance improvements, which are still unsatisfactory for 64 bytes packets, but allows to reach wire-speed from approximately 512 bytes packet size onwards. This is due to the larger amount of deployed resources and to reduced contention, as in the first case.

Finally, no major performance differences can be observed among the various internal networking configurations: thus, the utilization of hubs or VLAN tagging functionalities does not influence performance in the studied scenario, where the bottleneck is CPU overload.

The measurements reported assess the feasibility of the virtualized multi-stage router architecture, but our tests highlight some performance issues due to software limitations and resource sharing. In the next future, the natural evolution of hardware and software solutions will help to solve the inefficiency issues by adding more capacity and improved virtualization infrastructures; instead the resource sharing is more interesting from an architectural point of view.

In the recent work on this topic, we are trying to understand the network bottlenecks and resource sharing issues of this architecture thus we propose new measuring scenarios with different mapping relations between VMs and physical servers. The aim for this work is to find the best allocation of VMs into physical server in terms of performance and suggest guidelines for following works. Due to the space limitation, no details are shown here.

We demonstrate the feasibility of a multistage router architecture in a virtualized environment and highlighted strong performance limitations that makes this approach rather difficult to pursue today. To improve performance, different mapping techniques have been tested and we showed that by carefully assigning VMs to servers better performance can be achieved by isolating the resource-hungry VMs. Research work is needed in many areas to improve further this kind of approach.

References

[1] A. Bianco, J. Finochietto, M. Mellia, F. Neri, and G. Galante, “Multistage switching architectures for software routers,” *IEEE Network*, vol. 21, no. 4, pp. 15–21, Jul.-Aug. 2007.

[2] A. Bianco, R. Birke, L. Giraudo, and N. Li, “Multistage software routers in a virtual environment,” in *IEEE Globecom 2010 (Next Generation Networking Symposium)*, Dec. 2010.