Migrating Xen Virtual Routers with No Packet Loss

Pedro S. Pisa, Marcelo D. D. Moreira, Hugo E. T. Carvalho, Lyno H. G. Ferraz, Otto C. M. B. Duarte
Universidade Federal do Rio de Janeiro – GTA/COPPE - Rio de Janeiro, Brazil

***Abstract*—This paper addresses the issue of migrating Xen virtual machines acting as routers. We present an implementation and an evaluation of a new live migration tool for Xen. Our mechanism migrates without any data packet loss.**

I. Introduction

The virtualization technology creates an abstraction layer that decouples virtual machines from the underlying hardware. Because of this, virtual machines can freely migrate over the physical infrastructure. Live migration, or hot migation, transfers a virtual machine from a physical machine to another, keeping alive all available services and applications that were running. Therefore, live migration allows the network manager to reallocate the network resources in order to save power, load balance, or schedule maintenance.

II. Server and Router Live Migration

Xen has already a built-in live migration mechanism, which consists of an iterative copy of virtual machine memory pages. In the first iteration, all memory pages are transferred to the destination physical machine. But the applications running in the virtual machine keep writing in memory. Hence, in each following iteration, all memory pages written by running applications during the previous iteration are transferred. In the last iteration, the virtual machine is suspended before the last transfer is started. Finally, the virtual machine is resumed on the destination physical machine [1].

The Xen live migration mechanism performs well for virtual servers, because applications often use TCP, which implements a retransmission mechanism that recovers the packet loss caused by the suspension of the virtual machine. Nevertheless, if the virtual machine acts as a virtual router, its main role is to forward packets, and there is no retransmission mechanism at IP level. As a consequence, this type of live migration incurs in packet losses.

In order to solve this problem, we propose a modification to Xen packet forwarding mechanism to separate the control and data planes of Xen-based virtual routers. The control plane implements routing protocols, such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), and the data plane is responsible for packet forwarding [2]. We propose to maintain the control plane running in the virtual machine, but move the data plane to privileged domain of Xen, called Domain 0. Using this architecture, it is possible to use the Xen built-in live migration tool without any data packet loss.

For each virtual machine, our implementation gets the calculated routes by routing protocols running in the virtual machine and creates in Domain0 a new route table associated with this virtual machine. That way, the packets that would be forwarded by the virtual machine are immediately forwarded by Domain 0. Therefore, the data packets continue to be forwarded in the meanwhile the virtual machine is suspended. The communication between the Domain 0 and the virtual machine is made by creating a bridge in the Domain 0 and a network interface in the virtual machine that are both connected in the same dedicated network. Every Domain 0 has a bridge with the same dedicated network, and the virtual machine is able to communicate with the destination Domain 0 after the migration. The communication between the virtual machine and the Domain 0 uses XML messages exchanged over TCP.

Our migration approach is a zero packet-loss process that improves the Xen live migration tool. Packet forwarding happens in the Domain0, which is not suspended during the virtual machine migration. While the virtual machine is migrated, Domain0 continues forwarding packets, resulting in zero downtime for data packets. Although the control plane still has a downtime, it is not prejudicial because routing protocols implement their own retransmission mechanisms or run over TCP.

III. Evaluation and Concluding Remarks

We compare the Xen original live migration tool with our approach. In the evaluation scenario, we set a virtual machine to route the network traffic between two physical machines. We then migrate the virtual machine to a third physical machine in the meanwhile the traffic is forwarded. Preliminary results point out that the Xen original live migration tool loses packets for about 0.5 second and our implementation does not lose any packets during the migration time.

This paper shows that the Xen built-in live migration tool has to suspend the process during the migration and, consequently, for virtual routers applications there are packet losses. We implement a migration process with separated control and data planes to avoid packet losses during migration.

Acknowledgments

This work was supported by FINEP, FUNTTEL, CNPq, CAPES, and FAPERJ.

References

[1] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield. Live migration of virtual machines. In *NSDI'05*. USENIX Association, 2005.

[2] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford. Virtual routers on the move: live router migration as a network-management primitive. In *SIGCOMM '08*. ACM, 2008.