Ecole Doctorale STIC

Sciences & Technologies de l'Information et de la Communication

ED STIC - Proposition de Sujets de Thèse

pour la campagne d'Allocation de thèses 2015

Axe Sophi@Stic :	SystèmesRéseaux|
Titre du sujet :	Towards NFV-based Green Networks
Mention de thèse :	Informatique
HDR Directeur de thèse inscrit à l'ED STIC :	Guillaume Urvoy-keller

Co-encadrant de thèse éventuel :

Nom :	Lopez-Pacheco
Prénom :	Dino
Email :	Dino.LOPEZ@unice.fr
Téléphone :	

Email de contact pour ce sujet :	urvoy@i3s.unice.fr
Laboratoire d'accueil :	I3S

Description du sujet :

Introduction

****

To optimize the energy consumption of networks, several strategies exist, like server consolidation, energy aware routing solution, traffic engineering, etc [1, 2, 3]. Migration of virtual machines (VMs) is the key to server consolidation. Those VMs might be the actual servers of users in an IaaS scenario, where a user fully controls his virtual server. They might also be instantiations of Virtual Network Functions, e.g., virtual switches, virtual firewall, virtual IDSor virtual load balancers as the cloud market evolves towards enriching the network functionalities available to the end-users. A consequence of the latter evolution is the increase of the number of VMs to be hosted in data centers. On the other hand, usage of data centers (and virtual networks) vary greatly over time and periods of inactivity are not rare, esp. when we further consider the

fact that physical networks and data centers are dimensioned based on the peak hour rule [4].

Objectives
****
By jointly considering the two above issues – increase of the number of VMs and frequent periods of inactivity – we would like to explore a more extreme approach to the issue of energy reduction in data centers by enabling a complete hibernations of VMs to further reduce the number of active physical servers. The latter should be transparent to the user of the service or the owner of a virtual network.
We aim at designing Virtualized Network Functions (VNFs) that will allow turning off unused virtual devices, while providing to the customer an “always available” network view, and fast availability. This means that during periods of inactivity, the networking stacks of the VMs could be migrated to a sink VNF that could, on one side, answers on behalf of the frozen servers in case of basic queries (e.g., keep-alive) and, on the other side, quickly restarts the frozen VMs at the end of an inactivity period.
Several challenges pave our route esp., (i) how to offload established transport layer connections from a VMs to the sink VNF and (ii) how to quickly restart a frozen server and export connections states back to it.
We will study and analyze the deployment speed offered by different virtualization solutions, especially Linux Containers (LXCs), and we will devise the needed VNFs to have a fully operational sink VNF. For that aim, the candidate will first identify the basic network services that must be migrated from a VM to a sink VNF, and propose strategies for the migration of the network protocol states. One might draw inspiration from know Migratory TCP [5] or a solution based on a TIFID (Transport Independent Flow Identifier) [6], even though both need some modifications of the TCP stack. Software Defined Networking [7], which provides a fined-control of network flows should also be a key enabler in the design of the sink VNF.
The performance of the proposed VNFs will be studied through different metrics, like reactivity and availability of services with respect to users requests. The robustness of the VNFs will be addressed by analyzing the consistency of the network protocol states during the migration between the VM and VNF.
Another key scenario of interest in which the above approach should be valuable is in large SaaS deployment, esp. in virtual Content Delivery Networks (vCDNs), which are heavily used for video and audio streaming traffic. Indeed, vCDNs entail the orchestration of a large number of VMs under the control of a central entity and in complex interaction with other services like the DNS. Again, the ability to turn-off VMs while leaving external entities/services unaware of this fact, is crucial to avoid service disruptions and could significantly reduce the carbon footprint of the service. An additional challenge of the sink NFV approach in this scenario is the impact on the QoS perceived by users, as services deployed with CDN (and vCDN) are in general QoS sensitive. The QoS will be studied here from a networking perspective: packet loss, jitter, throughput, play-out buffer level, etc.

Methodology
****
The evaluation of the performance of the proposed VNFs, vCDNs and QoS will de done

experimentally. The tools to be used are:
Linux containers -LXC-. Any VNF should be deployed on LXC.
Hypervisor-based virtualization solutions, like KVM.
C/C++ and Python programming language
Bash and Linux environments

Some references
****

[1] Xin Li, Zhuzhong Qian, Sanglu Lu, Jie Wu. Energy efficient virtual machine placement algorithm with balanced and improved resource utilization in a data center. Mathematical and Computer Modelling, Elservier, 58 (2013) 1222–1235.
[2] Yunfei Shang, Dan Li, and Mingwei Xu. 2010. Energy-aware routing in data center network. In Proceedings of the first ACM SIGCOMM workshop on Green networking (Green Networking '10). ACM, New York, NY, USA, 1-8.
[3] Lin Wang, Fa Zhang, Athanasios V. Vasilakos, Chenying Hou, and Zhiyong Liu. 2014. Joint virtual machine assignment and traffic engineering for green data center networks. SIGMETRICS Perform. Eval. Rev. 41, 3 (January 2014), 107-112.
[4] Minghong Lin; Wierman, A.; Andrew, L.L.H.; Thereska, E., Dynamic right-sizing for power-proportional data centers. In Proceedings of IEEE INFOCOM, 2011, vol., no., pp.1098,1106, 10-15 April 2011
[5] Sultan, F.; Srinivasan, K.; Iyer, D.; Iftode, L.; Migratory TCP: connection migration for service continuity in the Internet. In Proceedings of 22nd International Conference on Distributed Computing Systems, 2002. Vol., no., pp.469,470, 2002
[6] Umar Kalim, Mark K. Gardner, Eric J. Brown and Wu-chun Feng. Seamless Migration of Virtual Machines Across Networks. International Conference on Computer Communications and Networks (ICCCN), Nassau, Bahamas, Jul-Aug 2013.
[7] Kreutz, D., Ramos, F. M., Verissimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. proceedings of the IEEE, 103(1), 14-76.

English version:

Introduction
****

To optimize the energy consumption of networks, several strategies exist, like server consolidation, energy aware routing solution, traffic engineering, etc [1, 2, 3]. Migration of virtual machines (VMs) is the key to server consolidation. Those VMs might be the actual servers of users in an IaaS scenario, where a user fully controls his virtual server. They might also be instantiations of Virtual Network Functions, e.g., virtual switches, virtual firewall, virtual IDSor virtual load balancers as the cloud market evolves towards enriching the network functionalities available to the end-users. A consequence of the latter evolution is the increase of the number of VMs to be hosted in data centers. On the other hand, usage of data centers (and virtual networks) vary greatly over time and periods of inactivity are not rare, esp. when we further consider the fact that physical networks and data centers are dimensioned based on the peak hour rule [4].

Objectives
****

By jointly considering the two above issues – increase of the number of VMs and frequent periods of inactivity – we would like to explore a more extreme approach to the issue of energy reduction in data centers by enabling a complete hibernations of VMs to further reduce the number of active physical servers. The latter should be transparent to the user of the service or the owner of a virtual network.

We aim at designing Virtualized Network Functions (VNFs) that will allow turning off unused virtual devices, while providing to the customer an “always available” network view, and fast availability. This means that during periods of inactivity, the networking stacks of the VMs could be migrated to a sink VNF that could, on one side, answers on behalf of the frozen servers in case of basic queries (e.g., keep-alive) and, on the other side, quickly restarts the frozen VMs at the end of an inactivity period.

Several challenges pave our route esp., (i) how to offload established transport layer connections from a VMs to the sink VNF and (ii) how to quickly restart a frozen server and export connections states back to it.

We will study and analyze the deployment speed offered by different virtualization solutions, especially Linux Containers (LXCs), and we will devise the needed VNFs to have a fully operational sink VNF. For that aim, the candidate will first identify the basic network services that must be migrated from a VM to a sink VNF, and propose strategies for the migration of the network protocol states. One might draw inspiration from know Migratory TCP [5] or a solution based on a TIFID (Transport Independent Flow Identifier) [6], even though both need some modifications of the TCP stack. Software Defined Networking [7], which provides a fined-control of network flows should also be a key enabler in the design of the sink VNF.

The performance of the proposed VNFs will be studied through different metrics, like reactivity and availability of services with respect to users requests. The robustness of the VNFs will be addressed by analyzing the consistency of the network protocol states during the migration between the VM and VNF.

Another key scenario of interest in which the above approach should be valuable is in large SaaS deployment, esp. in virtual Content Delivery Networks (vCDNs), which are heavily used for video and audio streaming traffic. Indeed, vCDNs entail the orchestration of a large number of VMs under the control of a central entity and in complex interaction with other services like the DNS. Again, the ability to turn-off VMs while leaving external entities/services unaware of this fact, is crucial to avoid service disruptions and could significantly reduce the carbon footprint of the service. An additional challenge of the sink NFV approach in this scenario is the impact on the QoS perceived by users, as services deployed with CDN (and vCDN) are in general QoS sensitive. The QoS will be studied here from a networking perspective: packet loss, jitter, throughput, play-out buffer level, etc.

Methodology
****

The evaluation of the performance of the proposed VNFs, vCDNs and QoS will de done experimentally. The tools to be used are:

Linux containers -LXC-. Any VNF should be deployed on LXC.

Hypervisor-based virtualization solutions, like KVM.
C/C++ and Python programming language
Bash and Linux environments

Some references
****

[1] Xin Li, Zhuzhong Qian, Sanglu Lu, Jie Wu. Energy efficient virtual machine placement algorithm with balanced and improved resource utilization in a data center. Mathematical and Computer Modelling, Elservier, 58 (2013) 1222–1235.
[2] Yunfei Shang, Dan Li, and Mingwei Xu. 2010. Energy-aware routing in data center network. In Proceedings of the first ACM SIGCOMM workshop on Green networking (Green Networking '10). ACM, New York, NY, USA, 1-8.
[3] Lin Wang, Fa Zhang, Athanasios V. Vasilakos, Chenying Hou, and Zhiyong Liu. 2014. Joint virtual machine assignment and traffic engineering for green data center networks. SIGMETRICS Perform. Eval. Rev. 41, 3 (January 2014), 107-112.
[4] Minghong Lin; Wierman, A.; Andrew, L.L.H.; Thereska, E., Dynamic right-sizing for power-proportional data centers. In Proceedings of IEEE INFOCOM, 2011, vol., no., pp.1098,1106, 10-15 April 2011
[5] Sultan, F.; Srinivasan, K.; Iyer, D.; Iftode, L.; Migratory TCP: connection migration for service continuity in the Internet. In Proceedings of 22nd International Conference on Distributed Computing Systems, 2002. Vol., no., pp.469,470, 2002
[6] Umar Kalim, Mark K. Gardner, Eric J. Brown and Wu-chun Feng. Seamless Migration of Virtual Machines Across Networks. International Conference on Computer Communications and Networks (ICCCN), Nassau, Bahamas, Jul-Aug 2013.
[7] Kreutz, D., Ramos, F. M., Verissimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. proceedings of the IEEE, 103(1), 14-76.