Journal of Network and Computer Applications 64 (2016) 124–136

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

journal homepage: www.elsevier.com/locate/jnca

Optimal virtual network migration: A step closer for seamless resource mobility

Márcio Melo [a,b,*], Susana Sargento [b], Jorge Carapinha [a]

[a] Portugal Telecom Inovação, Aveiro, Portugal
[b] Instituto de Telecomunicações, University of Aveiro, Portugal

ARTICLE INFO

Article history:
Received 28 July 2014
Received in revised form
28 July 2015
Accepted 23 December 2015
Available online 6 February 2016

Keywords:
Migration
Network Virtualization
Mapping
NP-hard
Optimization
ILP model

ABSTRACT

Network Virtualization provides operators with the ability to move components of a virtual network (VN), or even the entire VN, between physical hosts, in real-time and seamlessly to the end-users.

This paper addresses virtual resource mobility from a fault management perspective: (i) it proposes VN Clone migration, which requires no assumptions regarding the protocols running inside the virtual networks or its own architecture; (ii) and it proposes virtual network re-embedding-node-link formulation (VNRE-NLF), an integer linear programming formulation to solve the online virtual network re-embedding problem as a simultaneous optimization of virtual nodes and virtual links, providing the optimal bound for the migration of virtual networks from physical resources prone to failure. This approach aims at minimizing the overall VN migration cost per re-embedding: (i) the number of virtual nodes migrated; and (ii) the physical bandwidth consumption.

The results are very promising: the VN Clone migration achieves no VN downtimes and it takes just a few seconds to be fully performed. This makes the VN Clone approach suitable both for non- and real-time traffic. The obtained results of VNRE-NLF show that VNs are highly resilient (i.e. resilience factor higher than 0.8) to fault events if no more than 2% of the physical resources are prone to failure. Moreover, it shows that it is not only important to have enough spare capacity to re-accommodate the virtual nodes and virtual links directly affected by the fault, but also to have additional capacity to accommodate virtual link re-assignments and virtual node migrations.

1. Introduction

One of the major features of Network Virtualization is the possibility to move virtual resources, i.e. Virtual Router (VR), on-demand and seamlessly from one physical host to another without losing network connectivity. The virtual resource migration was initially proposed by Wang et al. (2007, 2008) as a primitive for network management tasks; later on, Lo et al. (2005) proposed three algorithms to schedule which resources should be migrated firstly to minimize the overall migration cost.

The VN migration itself can be triggered by several reasons: (i) equipment/facilities maintenance; (ii) network performance; (iii) end-user requirement/service level agreement; (iv) energy saving; (v) security protection; (vi) fault management; (vii) end-user mobility. The VN migration process is composed by two distinct phases: (i) VN re-embedding – to obtain a new embedding solution of an existing VN, but taking into account the current constraints of the physical network (available bandwidth and available processing capacity); (ii) virtual resource migration – the operation of effectively moving the virtual resources, e.g. virtual links and virtual nodes.

This paper addresses the VN migration process from a fault management perspective. In the first phase, the VN re-embedding, an integer linear programming formulation is proposed, i.e. virtual network re-embedding - node-link formulation (VNRE-NLF), to obtain the optimal bound per VN re-embedding taking into account the VN migration cost: (i) number of virtual nodes migrated; (ii) overall bandwidth consumption. To address the second phase, the virtual resource migration, it is proposed VR Cloning as an alternative to VR live migration (Wang et al., 2007). VR Cloning is based on saving the current state of the VR and transferring the VR clone to the new physical host. Different triggers for VN migration are described and analyzed, which can be used as an input to define new heuristics and mathematical formulations to perform the VN re-embedding, taking into account the current constraints of the physical network (available bandwidth and available processing capacity).

* Corresponding author.
E-mail address: marciomelo@ua.pt (M. Melo).

http://dx.doi.org/10.1016/j.jnca.2015.12.015
1084-8045/

The performance of the VR Cloning is evaluated by means of real experiments. This proposed approach achieves no VN downtime and it takes 2.75 s to be fully performed. The performance of the VNRE-NLF is evaluated by means of simulation. The simulation results show that VNs resilience to fault events is dependent on the VN request rate and also on the percentage of physical resources prone to failure. The obtained results show that VNs are highly resilient (i.e. resilience factor higher than 0.8) to fault events if no more than 2% of the physical resources need to be shutdown.

The remainder of the paper is organized as follows. After summarizing the related work in Section 2, Section 3 presents the VN migration triggers and describes the virtual router migration process. Section 4 presents the Network Virtualization architecture proposed to support the VR cloning migration. Section 5 describes the virtual network re-embedding problem and the evaluation metrics. Section 6 describes the VNRE-NLF and presents one objective function that aims at minimizing the VN migration cost. Section 7 analyzes the performance of the VR clone method and evaluates the VNRE-NLF, and Section 8 concludes the paper and describes the future work.

2. Related work

2.1. VN migration

The Virtual Machine (VM) migration was initially proposed for data-centers as a way to move the CPU load from one physical server to another. This feature is not only important for load balancing purposes, but it can also be applicable for planned maintenance operations, i.e. moving away all the VMs from one physical server, that needs to be rebooted or shutdown, to different physical servers.

In order to reduce the downtime due to the VM migration process, Clark et al. (2005) proposed the live migration of VMs, within the same LAN, which allows a VM to be migrated while still running. This reduces significantly the downtime provoked by the VM migration, and it also makes the migration process seamless to the end-users or to the running applications. Ma et al. (2010) proposed enhancements to the live VM migration process, which both reduces the overall migration time and the total data transmitted in the order of 30%.

The VM migration approach is also applicable and important to the networking area, where it can be used for network maintenance operations or for network service deployment. For instance, it can be used to move critical (or non-critical) VRs from physical hosts that need some kind of maintenance operation without disrupting the routing protocols. This is much preferable and human error safer than manually configuring routing protocol metrics to move away the networking traffic from that physical router.

With that in mind, Wang et al. (2007) proposed Virtual ROuters On the Move (VROOM) as a primitive for network management tasks, which makes it possible to move virtual routers freely without changing the IP-layer topology. Wang et al. (2008) extended their prior work and proposed to decouple the data plane from the control plane of the VRs, which results in no performance impact on the data traffic when a hardware data plane is used, and very low impact when a software date plane is used.

The VR migration feature was also considered and evaluated on IPTV scenarios (Marquezan et al., 2009). Pisa et al. (2010) proposed a new migration model for XEN, using also data and control plane separation, which outperforms the XEN standard migration model. Lo et al. (2005) used the virtual router migration (Wang et al., 2008) as a primitive to perform the virtual network migration, i.e. the migration of an entire VN, and also proposed three algorithms to address the VN migration scheduling problem and to minimize the total migration cost.

Although the separation of the data plane from the control plane seems to be a very effective approach on the VR migration, since it reduces the downtime of the VN, it is also a limiting factor on the conception of new network architectures and on the deployment of new network protocols. We argue that not only the VN migration process should be independent of the networking protocols that are running on the virtual network, but also no assumption should be taken on the router architecture itself. Therefore, we consider each VR as a black-box and propose the VR Cloning as an alternative to the current VR live migration process (Wang et al., 2007).

2.2. Virtual network re-embedding (VNRE) problem

The VNRE problem can be formulated as an un-splittable flow problem (Zhu and Ammar, 2006) of resource re-allocation. In order to solve this problem, several centralized algorithms have been suggested (Lu and Turner, 2006; Zhu and Ammar, 2006; Yu et al., 2008; Farooq Butt et al., 2010; Nogueira et al., 2011b; Chowdhury et al., 2012) to address the online Virtual Network Embedding (VNE) problem, i.e. static provisioning of VNs; others consider, instead, a distributed algorithm (Houidi et al., 2008) to better handle scalability issues. However, these approaches neither take into account the VNRE problem, i.e. dynamic provisioning of VNs, nor consider the VN migration cost, i.e. virtual node migration and bandwidth re-allocation.

To handle physical network failures, Rahman et al. (2010) incorporated single substrate link failures on the VNE problem and proposed a heuristic to solve the problem. To handle resiliency protection against network failures during the process of online VN services provision, Chen et al. (2010) proposed an efficient resource allocation approach to balance the trade-off between service resource consumptions and service resiliency. However, both algorithms only support network link failures and do not incorporate the nodes.

Cai et al. (2010) proposed an algorithm to address network changes in response to network growth, node failures or node joining/leaving. Nevertheless, the algorithm focused on minimizing the upgrading cost of virtual networks, and not on the migration cost.

Moreover, centralized algorithms have been suggested (Yu et al., 2011; Rahman and Boutaba, 2013) to address the Embedding (SVNE) problem. The aim has been to reserve additional (or redundant) computing and bandwidth capacity for each VN when it is first provisioned, which can be used in situations of node and link failures. Nonetheless, these algorithms do not consider the scenario where the backup physical resources are not available, nor consider the VN migration cost.

To reduce the physical network fragmentation due to the VNs dynamic lifecycle, on a different perspective, Fajjari et al. (2011) proposed a virtual network reconfiguration algorithm, that relocates the VN star topology, instead of the whole VN topology (Zhu and Ammar, 2006), with the aim of minimizing the cost of reconfiguration.

Zhang and Qiu propose an algorithm that identifies virtual nodes and virtual links mapped in a non optimal way, and migrates them to better locations to save physical resources. However, this algorithm does not take into account the re-embedding of VNs due to link and node failures. Yeow et al. (2011) added node and link redundancy for reliability, and used a multi-commodity flow problem formulation to solve the VNE. To increase the resilience to link failures, they also considered path-splitting. However, this formulation does not contemplate the VN