SOFTWARE-DEFINED NETWORKS: DESIGN GLOBALLY, APPLY LOCALLY

СОФТУЕРНО ДЕФИНИРАНИ МРЕЖИ: ПРОЕКТИРАЙ ГЛОБАЛНО, ПРИЛАГАЙ ЛОКАЛНО

Ivan Iordanov, Stefan Dimitrov

Department of Computing Systems, Faculty of Mathematics and Informatics, Sofia University "St. Kliment Ohridski", 5, J. Bourchier, Blvd, 1164 Sofia, Bulgaria, phone: +359 87 8391862, e-mail: iiordanov@uni-sofia.bg, stefansd@fmi.uni-sofia.bg

Иван Йорданов, Стефан Димитров

Катедра "Изчислителни системи", Факултет по математика и информатика, СУ "Св. Кл. Охридски", бул. Джеймс Баучър 5 , 1164 София, България, тел.: +359 87 8391862, e-mail: iiordanov@uni-sofia.bg, stefansd@fmi.uni-sofia.bg

Keywords: Software-Defined Networking (SDN), Network function virtualization (NVF), Implementation strategy

Резюме – Този доклад представя преглед на препятствията, свързани с внедряването на софтуерно дефинирани мрежи (СДМ). Предлага се обща стратегия за въвеждане на виртуализиране на мрежовите функции (ВМФ) в средни до големи компании, както и мрежи на доставчици. Описваме ползите, които могат да бъдат натрупани чрез създаване на изцяло виртуализиран център за данни. Изредени са предимствата от замяната на всяко едно мрежово устройство в него. Анализирани са икономическата ефективност и повишената производителност, които биха последвали мигриране на терминиращите устройства за отдалечени корпоративни връзки към СДМ. Подчертана е ползата от въвеждане на унификация на дизайна, централизиране на контрола, стандартизиране на конфигурациите, софтуерните версии и хардуера в случаите, когато е възможно да разпространим изчислителните задачи сред географски разпръснати точки.

Abstract – In this paper, an overview of the challenges associated with Software-Defined Networking (SDN) implementation is provided and a general strategy for adoption of Network function virtualization (NVF) in medium to large size enterprise and service provider networks is proposed. We discuss the benefits that can be accrued by creating a fully virtualized data center. The advantages of replacing each type of network device therein are enumerated. The savings and efficiencies that can follow migrating the enterprise Wide Area Network (WAN) devices to SDN are expanded upon. The role of unified design, centralization of control and standardization in configuration, software and hardware combined with geographically distributed computing are highlighted.

1. INTRODUCTION

In the last decade we witnessed a redefinition of the concept of “server” and, indeed, data center (DC). The virtualization drive surpassed the halfway mark years ago – number of x86 virtual machines (VM) versus dedicated hardware servers in 2012 [1] and processed workload in 2014 [2]. Sadly, there has been no concomitant change on the network side of the equation. The current architecture is still inherited from the 20th century. This is true for the Internet, as well as the DC.

Most agree that the Future Internet will be, at least in part, based on Software-Defined Networking (SDN) [3]. It is the particulars and how to get there that pose a problem. Before a service provider (SP) can even consider migrating to an SDN core, they need to nurture in-house expertise in running a production software-defined network, form and test vendor relationships (even if the solution is entirely open source, having support contracts is only prudent), implement pilot projects, formulate a staged rollout plan, simulate it in a lab, etc. [4].

By far the less arduous of those tasks is the migration of DCs to SDN – the end devices in the data center are already largely virtualized, the network size is relatively modest, most vendors on the market offer solutions. Additionally, the expertise that will accumulate during this endeavor can later be used in re-imagining the SP core.

In section 2 of this paper, we will briefly go over the network components constituting a typical DC and some of benefits that can be accrued with migrating them to SDN.

In section 3, we offer an implantation strategy that can be applied by both service providers and medium to large enterprises.

2. BACKGROUND

Most current data centers enjoy some degree of centralized VM management, while the network side remains largely independent. This prevents the users and administrators from fully leveraging the benefits of VM orchestration (automated guest migration based on predefined conditions, one stop creation of new VM groups, etc.).

We will proceed with an overview of the typical hardware-centric network devices found in a DC and their SDN counterparts.

2.1 Switches

The use of purpose-built application-specific integrated circuits (ASICs) switches in DC environment imposes severe constraints.

Foremost of those is the need to create and modify virtual local area network (VLAN) entries and trunk port configuration directly on the devices. A partial solution can be applied by writing an API that translates the VLAN needs from the VM orchestration interface to device commands either via SNMP write credentials or purpose written configuration scripts (normally in Perl, Python or Bash). This workaround is more prone to errors (continuous in-house support for the code, keeping track of possible issues related to firmware updates) than a full blown SDN implementation.

Another major problem is the fact that a host server presents multiple VMs with the same MAC address. Hardware-centric OSI layer 2 network devices remain as yet unable to distinguish between traffic originating from two different VMs on the same host, because this information is not defined in the Ethernet header (Virtual Extensible LAN (VXLAN) relies on OSI layer 4 UDP encapsulation). This results in unnecessary traffic flows for all type of addressing – unicast, multicast, broadcast (Fig.1). Moving to a programmable overlay network will bring awareness to the intermediate devices.

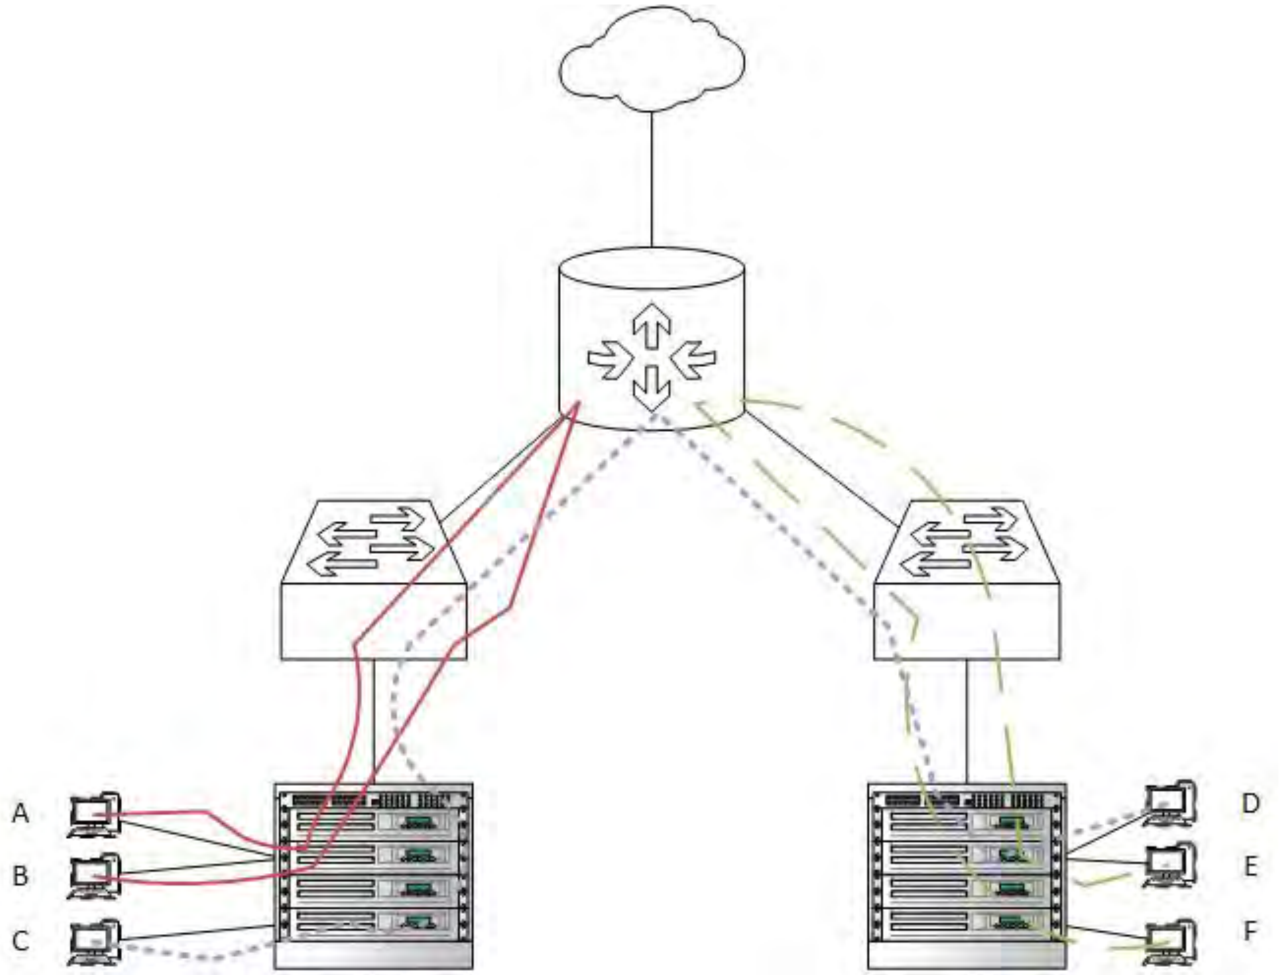

Fig. 1. In classical DC network design, all VM-to-VM communication must pass through a physical router. Even if the guests are on the same rack (D-E) or the same host server (A-B)

The final hurdle is the fact that guest-to-guest communication on the same host requires traversing the top-of-the-rack (ToR) switch, even if the VMs are part of the same subnet. Bringing switching functionality to the server hypervisor level ensures bus speed east-west connectivity (Fig. 2).

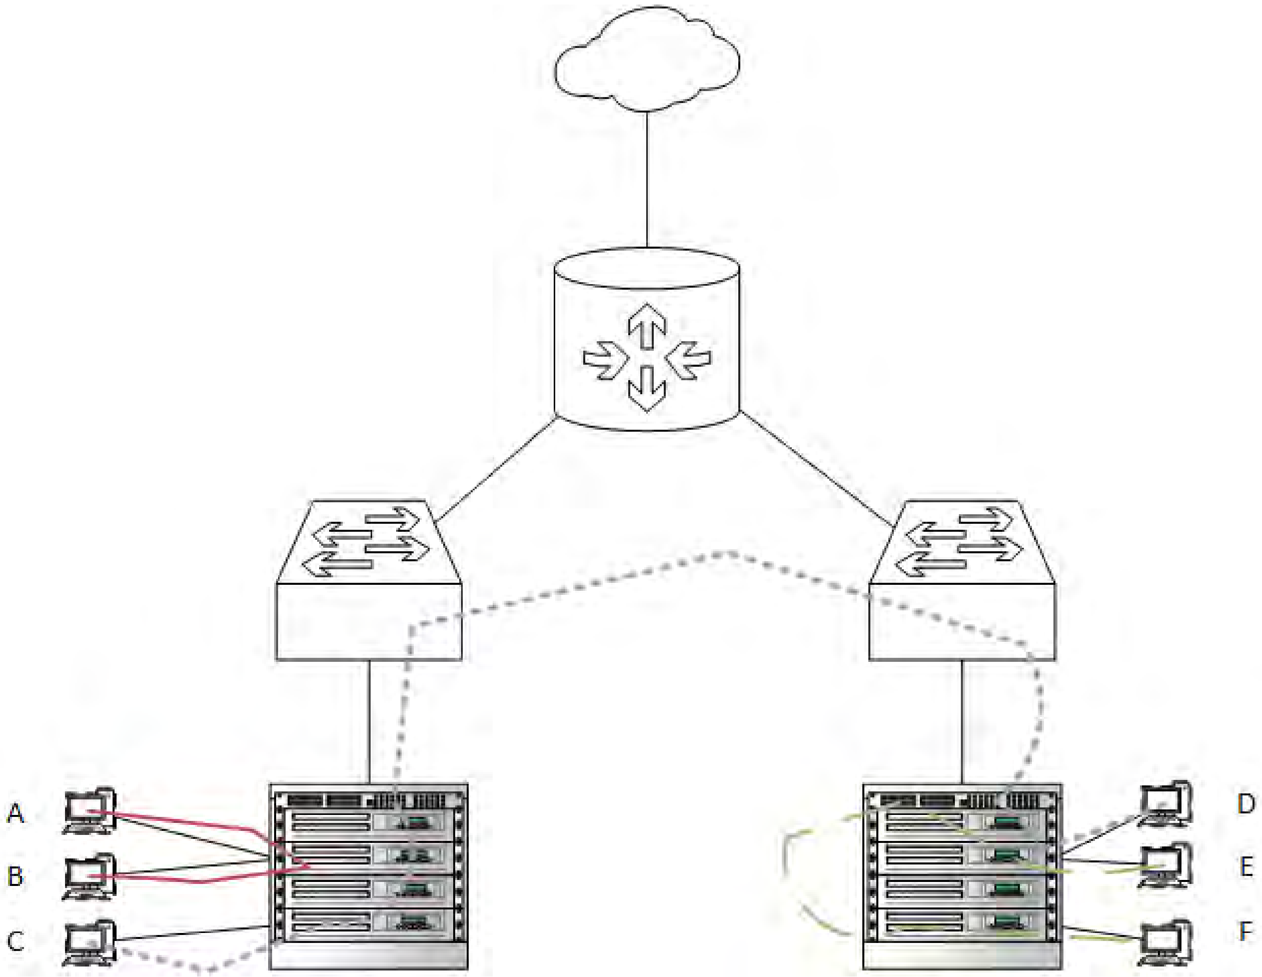

Fig. 2. In SDN DCs, the speed of lateral communication is greatly increased. Communication between two VMs on the same host (A-B) doesn't leave it. Traffic from guests on the same rack traverses only the ToR switch (E-F). Even communication between VMs on different racks can be accomplished by only passing through distribution/core switches (C-D). Only traffic towards external destinations, i.e. Internet, must exit via a physical router.

2.2 Routers

Utilizing hardware bound routing devices to service VMs relying on disparate computing and memory resources is equally problematic.

As with switches, a prevalent issue is the necessity of traffic flow being handled by network devices on at least five separate occasions in order to achieve east-west connectivity. By implementing distributed virtual routers this value can be reduced to anywhere from 0 (terminating VMs on the same host) to 3 (guests on different hosts located at different racks).

2.3 Firewalls

Another network device that greatly benefits from being implemented as overlay is the firewall.

In a traditional design, the DC firewall is a discrete hardware unit that filters traffic from the outside. The box itself is vulnerable to distributed denial of service (DDOS) attacks since it has a maximum number of concurrent connections that can be processed. Furthermore, if the increase in traffic is legitimate, improving the capacity necessitates installing additional hardware. Lastly, the only way to limit inter-DC VM-to-VM traffic is either with inserting yet another network device or configuring software firewall rules individually on each server.

An SDN firewall circumvents those issues. The outside traffic passes through a redundant set of appliances whose processing resources can be dynamically allocated (and increased, as needed), while the east-west flow is examined on host level based on globally configured rules. This allows for rapid change and expansion of capacity, yet ensures uniformity of configuration.

3. IMPLEMENTATION STRATEGY

As alluded in the previous section, novel technologies, such as SDN with network function virtualization (NFV), pose medium, to high, risk for the implementation project's success [5]. Bearing this in mind, when considering mid-to-large enterprise or service provider networks, it is prudent to separate the endeavor in self-contained stages.

3.1 SDN in the DC

The best place to start the rollout of NFV solutions in an organization is the DC. There are several key reasons for that:

- least capital investment
  - the overlay network can use the legacy hardware network devices
  - control and management nodes, as well as appliances, can be provisioned using the spare computing capacity that is maintained as a matter of course in a modern VM-based DC
- concentration of in-house expertise
  - the DC is supported by hosting engineers with experience in visualization who can directly assist with the initial appliance setup
  - on the pure IP networking and security side, the teams configuring and troubleshooting the legacy equipment are aware of the particular challenges a VM-based DC poses
- clear boundaries and project standardization
  - a data center is geographically contained entity which should prevent excessive scope creep
  - normally, if a company has several DCs, they are of the same mold which speeds up consecutive rollouts
  - even if the designs are disparate, the possible divergence is relatively limited (as compared to a continent spanning SP network, for example) and, as such, 3rd party experts can provide more thorough advise faster

If we turn to the benefits that can be accrued by such a migration, there are several that stand out:

- improved efficiency in east-west communication
- elimination of packets misdirection due to incompatibility between Ethernet and VXLAN
- increased resiliency of the overlay and single point of failure (SPOF) reduction due to moving away from hardware dependent networking

- greater flexibility in responding to unexpected events by software provisioning of new appliances and/or allocating additional computing resources to existing ones

3.1 SDN in the enterprise LAN

Another relatively straightforward place that can converted to NVF is the enterprise LAN, or the customer-premises equipment (CPE) for service providers. Here the transition would most benefit from global overall design. A massive advantage in SDN would be improved site autonomy – most enterprise networks desire conformity between several campuses (headquarters (HQ), regional HQs, branches). For practical reasons, however (difficulty in synchronizing configuration between hardware with different capabilities and keeping it up to date, etc.) the logical topologies for most services varies between star (all sites connect to HQ or a central DC) to exploded start (services are clustered around regional HQs or several DCs which rely on partial mesh or simple hub design to cross-connect over HQ).

Employing SDN will bring cost savings due to:

- reduced need for site to site connectivity – even though enterprises no longer exclusively use costly setups like leased lines and dark fibers to connect their branches, relying more and more on VPN-MPLS or MAN connectivity from SPs, that solution, while cheaper, still costs more per megabit than pure Internet uplink
- standardization on hardware – a solution with NVF will heavily rely on typical setups, i.e. a virtualization pod consisting of, for example, a host sever and uninterrupted power supply (UPS) unit. For the price of an advanced router, the enterprise edge and branch can blend while still allowing the IT team in the center to retain control

We shall also list a few of the improvement that some services will experience with the implementation of SDN:

- firewalls – most companies want consistency and heavy control over the Internet experience of their employees. This usually translates to routing all traffic through a central firewall, which is, as pointed above, incredibly costly both as financial budget and traffic latency. With the use of SDN, all rules from the central location can be implemented directly on branch level. Even if MPLS connectivity to HQ is lost, the branch edge will still continue to enforce them and provide all services that utilize pure Internet connectivity
- wan optimization – WANOPT solutions really come on their own when there is also an independent Internet uplink to offload web requests to (and having a web caching service, reduces the utilization of that link too)
- redundancy and load balancing – for optimal resiliency, a branch site will have two independent Internet connections – if one is negotiated through a global partner and the other a country wide provider, this could help reduce re-sale price over-inflation during negotiations; there should also be two links to the rest of the corporate network – an MPLS through the global partner and an IPSEC over the second Internet connection to HQ. Traffic can be allocated to each link dynamically to increase overall performance.

- centralized role base access – by consulting a unified authentication database, the security principle of "strap-and-suspenders" can be upheld. Only machines whose actively logged in users are authorized for certain systems will be allowed to reach their resources on IP connectivity level; each user will, of course, have to authenticate to the server for their level of access. This compartmentalizes the damage that a hijacked machine can wreak.
- improved administration – there is a litany of cases where being able to setup simple VMs on both ends of a misbehaving circuit to run a simple iperf test would have saved expensive and time consuming on-site interventions. Being able to dynamically configure test machines on the branch edge really opens up the possibilities for monitoring and troubleshooting thus improving reliability.

4. CONCLUSION

Year on year SDN adoption has been increasing [6]. It is still far from widespread, but as the technology matures its use moves to the mainstream. This is most notable in the DC market where more and more vendors are entering [7].

In this paper, we focused on the benefits offered by migrating to SDN in the data center and the overall enterprise network. We proposed that even service providers should start from those two steps before embarking on migration to pure SDN core, which should be the subject of a separate study.

In any case, unity of design and management control should be maintained in order to fully leverage the distributed features of SDN.

5. REFERENCES

[1] http://www.serverwatch.com/server-trends/survey-51-of-x86-servers-now-virtualized.html, last accessed on September 28, 2015.

[2] http://www.cioinsight.com/it-strategy/cloud-virtualization/slideshows/useful-virtualization-stats-trends-and-practices.html, last accessed on September 28, 2015.

[3] A. Hakiri, A. Gokhale, P. Berthou, D. Schmidt and T. Gayraud,"Software-Defined Networking: Challenges and research opportunities for Future Internet", *J. Computer Networks*, vol. 75, pp. 453–471, December 2014

[4] M. Kobayashi, S. Seetharaman, G. Parulkar, G. Appenzeller, J. Little, J. van Reijendam, P. Weissmann and N. McKeown, "Maturing of OpenFlow and Software-defined Networking through deployments", *J. Computer Networks*, vol. 61, pp. 151–175, March 2014

[5] Sebastián Martorell, Carlos Guedes Soares and Julie Barnett, *Safety, Reliability and Risk Analysis: Theory, Methods and Applications,* CRC Press, 2014, New York, ISBN 1482266482.

[6] J. Metzler, "The 2015 Guide to SDN and NFV Part 1: Software Defined Networking (SDN)", Webtorials, 2014, http://www.webtorials.com/content/2014/11/the-2015-guide-to-sdn-nfv.html, last accessed on September 28, 2015.

[7] M. Raza, S. Sivakumarb, A. Nafarieha and B. Robertsona, "A Comparison of Software Defined Network (SDN) Implementation Strategies", *Procedia Computer Science 32 – "The 2nd International Workshop on Survivable and Robust Optical Networks (IWSRON 2014)"*, Hasselt, Belgium, pp. 1050 – 1055, June 2-5, 2014