REARM: Renewable Energy based Resilient Deployment of Virtual Network Functions

Sameer G Kulkarni*, Mayutan Arumaithurai*, K.K. Ramakrishnan† and Xiaoming Fu*
*University of Göttingen, Germany, †University of California, Riverside.

***Abstract*—Network Function Virtualisation (NFV) is becoming more prevalent in Data Center, Telecommunication and Enterprise networks, and the Virtual Network Functions (VNFs) are fast replacing the traditional dedicated hardware based middleboxes. Also, the inclination towards employing renewable (green) resources to power up the Data Centers is also increasing. Mitigating the carbon footprint and curbing the energy costs have been the driving factors for push towards employing the green energy resources. However, the Green energy supply is rather intermittent and unstable. In this work, we study the impact of deploying VNFs in Green Data Centers (GDCs) and make a case for addressing the VNF reliability and high availability despite the stability concerns of GDC. To this extent, we propose the concept of Transient VNFs that rely on very short advance warning time to seamlessly migrate the VNFs from GDC to the more reliable and stable Data Centers.**

I. Introduction

Network functions such as Firewalls, Deep packet inspection (DPI), cache optimization, load balancing, *etc.* have become an integral part of large scale enterprise and Data center (DC) networks. NFV enables the deployment of VNFs (software based middleboxes) on top of the commercial off-the-shelf (COTS) hardware rather than using the dedicated hardware appliances. This allows for flexible realization of network services with greater cost optimization both in-terms of capital expenses (CapEx) and operational expenses (OpEx). More importantly, NFV caters towards better energy efficiency due to consolidation of compute and network resources [1].

It has been studied that the Data center industry accounts to over 30 Gigawatts of energy per year [2] and the demand for power keeps increasing every day. The carbon footprint of a medium 10 Megawatt data center can range from 3,000,000 to over 130,000,000 kilograms of CO2 [3]. Depending on the electric grid region, Power Usage Effectiveness (PUE) improvements can eliminate millions of pounds of CO2 emissions [4]. These factors have led to tremendous increase in the widespread adoption of renewable resources for powering the Data centers. The recent study [5] indicates a phenomenal increase in the investments ($285.9 billion) for harnessing renewable energy, which is more than double ($130.6 billion) the investments on non-renewable energy resources in 2016. It is also noteworthy that the amount of renewable energy generation capacity has increased by nearly 56 percent over last two years. Despite growth, the nature of renewable resource based power is i) not sufficient to fully power the large data centers, ii) highly intermittent and unstable, hence pose a greater challenge in adopting them for the large Data Centers which require stable and sustained power resources in-order to avoid any service disruptions. However the Green energy could be used to adequately power a small DC with reasonable degree of reliability.

To this end, we present our work REARM[1], that aims to enable running the VNFs in renewable energy backed Data centers while providing sufficient degree of reliability and high availability. The key contributions of our work include:

- Distinguish and illustrate the VNF deployment model and associated challenges in providing resiliency and high availability.
- Our proposal REARM, that builds on top of YANK [6], to provide a framework for transient VNFs that can be efficiently and reliably run in GDCs.

II. Related work

In [1], authors analyze the prospects of energy efficiency by employing the VNFs for the Evolved Packet Core (EPC), Customer Premise Equipment (CPE) and Radio Access Network (RAN) in telecommunication networks. This study is seminal in terms of establishing the energy efficiency prospects of VNFs. In [7], authors analyze and present the energy efficiency implications of NFV for different packet processing mechanisms. In contrast, we consider a more broader perspective and target towards achieving energy efficient network infrastructure that can be powered by renewable resources and still be able to meet the high availability and resiliency requirements. The prospects of employing Green energy for the VNFs is a less studied topic. However, plethora of work exist in the context of Virtual Machines (VMs), which can be pertinent in the context of VNFs.

III. Research Proposal

First, we characterize the VNFs and list the fundamental differences w.r.t the VMs, and the associated challenges towards deploying the VNFs in GDCs. Then, we briefly present the architecture of REARM, that aims to overcome these challenges and enable for resilient deployment of VNFs.

A. NFV Deployment Model and Usage Scenario

Generally the VM's are application processing engines characterized by the application states, while VNFs are essentially the high speed packet processing engines that maintain flow/packet specific states and tend to serve millions of packets

[1]REARM: **R**enewable **E**nergy b**A**sed **R**esilient deploy**M**ent of virtual network functions

per second at 10G/40G/100Gbps line rates and depending on the type of processing, VNFs can be either stateful or stateless. This means the frequency at which the VNF state changes is too high and even a sub-second of downtime leads to severe service disruptions. Hence, to achieve high-availability, consistent updates need to be done more frequently than compared to the traditional VMs. In addition to VNF state, the network routing state also needs to be updated for reliable processing of subsequent packets. Hence, only the stateful class of NFVs need to snapshot their internal states, while stateless NFVs need only the routing state update. Also, the amount of state that need to be transferred to back-up the stateful VNF is minimal (few mega bytes) compared to VMs that range in several giga bytes.

Second, typical VNF deployment models include i) as dedicated VMs, ii) as Container or Docker based applications, iii) Packaged network appliances iv) as binaries that can be run as dedicated processes. This diversity not only hinders portability - since the Docker based and process based VNFs need to be backed-up on nodes matching the hardware and Operating system requirements, but also pose a challenge towards achieving generalized framework for replication as the needs and means to snapshot and back-up the VNFs significantly differ. However, the promising part of the VNF diversity is that the amount of data that need to be backed-up is significantly lower compared to the traditional VMs.

Finally, the flows served by the VNFs are typically subject to more than one network functions, processed in a specific order, referred as Service Function Chain (SFC). *e.g.,* NAT, Firewall, IDS, and Load-balancer. This implies the VNFs cannot be treated in isolation, but the chain (ordered list) of network functions are to be treated as a group of services. Hence the back-up and snapshot mechanism should consider the periodicity for group of VNFs.

B. REARM: An extension of YANK

Figure 1 shows the high level architecture of REARM. Key components of our architecture include i) NFV Orchestrator: responsible to manage, co-ordinate and communicate with the VNF managers to handle snapshot and restoration of stateful VNFs in a data center. ii) SDN Controller: responsible to setup paths for the flows, and to migrate the flows to the appropriate VNFs in the data center. iii) VNF Managers: responsible to instantiate and manage the VNFs and to perform snapshot and restoration of VNFs on a physical node.

1) Advance Warning Time: We leverage YANK [6] concept of advance warning time to backup the transient VNFs on the stable servers. On receiving the advance warning for power disruption events, REARM performs i) VNF state update for all the stateful VNFs and ii) routing updates using the SDN controller to setup up the path to the steer the flows to the new stable VNFs, as the amount of state update is too minimal (few Mbs), we believe a few milliseconds (50 to 100ms) of warning time is sufficient to backup the VNFs.

2) Snapshot and Restore Mechanism: Backups on stable nodes are only maintained for Stateful NFs. We make use of

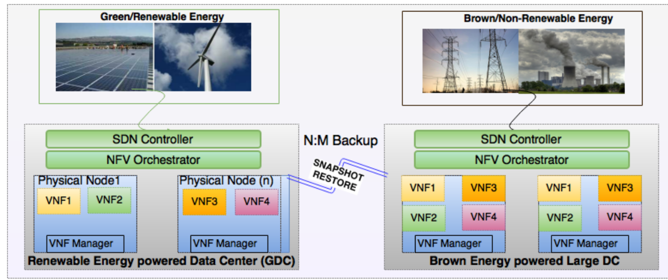

Fig. 1: REARM Architecture

the N:1 backup model, so that more than one transient VNFs in GDC can be backed up on a single node in the brown energy powered data center. This helps to save on brown energy, but at the cost of increased latency and lower throughput for short period of time. Once the power levels are back to normal operating conditions, NFV Orchestrator triggers for the restoration of transient VNFs from the stable data center. YANK additionally presents several optimization techniques for efficient disk state and memory replication, which we plan to account for VNF state transfer so as to minimize the overall transfer of states. As an added optimization, we consider batching the updates for an entire chain of VNFs, so that the state of VNFs in a service chain is coherently synchronized and backed-up on stable services.

IV. Conclusion & Future Work

We have characterized and analyzed the benefits and challenges in employing the Green Data Centers (GDCs) for VNFs. In order to meet the high availability and resiliency requirements, we make a proposal of ***REARM***, that leverages and enhances YANK, especially to caters towards the special needs of VNFs. We seek to discuss further and incorporate improvements in our proposal based on the feedback from the community. Next, we plan to prototype our solution and study the associated trade-offs and quantify the benefits through thorough evaluation on our SDN/NFV test-bed.

References

[1] R. Mijumbi, "On the energy efficiency prospects of network function virtualization," *CoRR*, vol. abs/1512.00215, 2015. [Online]. Available: http://arxiv.org/abs/1512.00215

[2] W. V. Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, "Trends in worldwide {ICT} electricity consumption from 2007 to 2012," *Computer Communications*, vol. 50, pp. 64 – 76, 2014, green Networking. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0140366414000619

[3] J. Koomey, "Growth in data center electricity use 2005 to 2010," 2011.

[4] "Green house data, a report on energy consumption," https://www.greenhousedata.com/green-data-centers, accessed: 2016-10-02.

[5] "Global trends in renewable energy investment 2016," United Nations Environment Programme, Bloomberg New Energy Finance, accessed: 2016-11-10.

[6] R. Singh, D. Irwin, P. Shenoy, and K. Ramakrishnan, "Yank: Enabling green data centers to pull the plug," in *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*. Lombard, IL: USENIX, 2013, pp. 143–155. [Online]. Available: https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/singh

[7] Z. Xu, F. Liu, T. Wang, and H. Xu, "Demystifying the energy efficiency of network function virtualization," in *2016 IEEE/ACM 24th International Symposium on Quality of Service (IWQoS)*, June 2016, pp. 1–10.