Network Working Group
Internet Draft
Intended status: Informational
Expires: January 3, 2015

C. Huang
Carleton University
Jiafeng Zhu
Huawei
Peng He
Ciena
Shucheng (Will) Liu
Huawei
Baek-Yong Choi
University of Missouri-Kansas City

July 3, 2014

Use Cases on Application-centric Network Management and Service Provision
draft-huang-aponf-use-cases-01.txt

Status of this Memo

Copyright Notice

Abstract

New services such as virtual networks, service function chaining, and application-centric traffic steering bring new opportunities for network providers and service providers. With these new services, the interactions between applications and networks are becoming more critical in order to achieve satisfactory QoS, reliability, and security, which are tailored to each specific application. Application-based Policy for Network Function (APONF) is designed to facilitate these interactions. This internet draft describes some use cases that show how APONF may be deployed to support these services.

Table of Contents

1. Introduction

The wide deployment of optical networks and large-scale datacenters has enabled many new services such as network virtualization, service chaining, and application-centric traffic steering which bring new opportunities to network and service providers.

One of the primary new services that have been envisioned is the network virtualization service, which allows physical network provider to sell different virtual networks to different service network providers. Each service network provider can use its virtual network just like the way it uses its own private network while sharing underlying physical network resources with other

service network providers. The physical network provider, on the other hand, can enjoy new revenue growth through selling virtual networks with different granularities.

Traditionally a service chain consists of a set of dedicated network service boxes such as firewall, load balancers, and application delivery controllers that are concatenated together to support a specific application. With a new service request, new devices must be installed and interconnected in certain order. This can be a very complex, time-consuming, and error-prone process, requiring careful planning of topology changes and network outages and incurring high OPEX. This situation is exacerbated when a tenant requires different service sequences for different traffic flows or when multiple tenants share the same datacenter network.

Network Function Virtualization (NFV) is a concept built upon network virtualization. It involves the implementation of network functions such as load balancing, intrusion detection, firewall, monitoring, and accelerations in software that can run on a range of industry standard high volume servers, switches, and storage. Through NFV, service providers can dynamically create a virtual environment for a specific service chain and eliminate the dedicated hardware and complex labor work for supporting a new service chain request.

Both network virtualization and NFV require network management and service provisioning. Experiences from the services provided by modern datacenters (e.g. Amazon EC2) have shown that fast response time and easy-to-use are critical for the success of these new services. There is no existing management and service provisioning infrastructure to support network virtualization and NFV. However, these new services typically require complex interactions and orchestration between service subscribers and network infrastructure, which are missing in the existing Internet architecture.

There are many other use cases that can be better served with application-centric network management and service provisioning. Application service providers have tried various ways to differentiate their customers so that they can maximize their revenues and minimize their costs. For example, cookies have been used to track HTTP users. Unfortunately they are designed for specific applications. Because cookies only appear in HTTP headers, they will not be carried by all packets except the first one. Therefore they cannot be used by operators to provision switches. On the other hand, VLAN and DiffServ are hard to be maintained at the level of end-to-end, since they operate on Layers 2 and 3. In this

regard, application-centric network management and service provisioning, although widely desired, are still hard to achieve.

In this document, we describe some typical use cases that may be supported using APONF architecture.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119.

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC-2119 significance.

3. Use Cases

There are numerous use cases that APONF can be applied to. Some common use cases are described in this section.

3.1. Network Virtualization Service

The first use case is a network virtualization service. Here a physical network provider will serve as the service provider and virtual network providers will serve as clients. Virtual network providers request virtual networks by contacting Application Based Policy Decision (ABPD) module. ABPD module will check against its policy to see whether the request is acceptable. If not, it will reject the request. If acceptable, it will assign an ID to identify the request and map the virtual topology to physical networks with required performance. ABPD will then contact related physical network providers to provision the required resources. After all the required networks resources are reserved and setup, the infrastructure network providers will inform ABPD, which will then update its database about those physical networks and inform the virtual network provider about its requested virtual network and associated ID to identify the required virtual network.

When the users of the virtual network provider send traffic, they need to encapsulate the ID in their packets so that network providers can identify the traffic belonging to the same user group and treat them with required resources and policies. In the case that the users of the virtual network provider cannot encapsulate the ID into packets due to their software version, the ABDP, after learning this information from its interactions with the virtual network provider, will inform the border routers to classify traffic

and encapsulate and de-capsulate packets with the ID. Without ABDP, this kind of negotiation and configuration will not be possible.

3.2. Virtualized Enterprise Applications

Virtualized enterprise applications make the Virtualized Network Function (VNF) functionality available to enterprise users as a service, comparable to the cloud computing concept denoted as the Software as a Service (SaaS), see [NIST SP 800-146]. Virtualized enterprise application policies include dynamic orchestration of virtualized network functions, dynamic increase/decrease of network bandwidth, pay-as-you-go billing, etc.

GiLAN is another important application of a virtualized network function, as it is a boundary between mobile and data networks which operate on very different strategies. In mobile core networks, it is preferable that QoS provisioning and network function requirements are different for subscribers with different profiles. In such scenarios, specialized applications such as BSS/OSS can send service policies to a policy decision point, which further map these service policies to GiLAN specific VNF policies, and realize the required QoS and with appropriate network functions, for example, video transcoding parameters, traffic steering points, etc.

Consider a scenario where an enterprise requests such a service. The enterprise will contact ABDP for the required resources. After receiving the request, the ABDP will check its policy and database to see whether the required service can be supported or not. If yes, the ABDP will assign an ID for each of the requested service, map the request to network policies and related service functions in the infrastructure environment, and send a request to the infrastructure provider with related network policies and service functions. The infrastructure provider will install the requested policies and service functions and confirm the request. The ABDP will confirm to the enterprise with the associated IDs and update its database. When the users of the enterprise need to send traffic, they will encapsulate their packets with corresponding IDs. The infrastructure network can then apply appropriate service functions and policies to the corresponding packets. If the application users cannot add the IDs, the ABPD will inform edge service function nodes to classify traffic and encapsulate and de-capsulate packets.

3.3. Application-centric QoS and Reliability

Service providers are increasingly interested in providing different treatments to different types of customers, e.g. subscribers vs. casual users. User traffic flows need to be steered to different

environments with different networking and computing resources provisioned. Under this context, ABPD provides a simple and effective handle that connects applications to physical layer devices and enables application-centric network management and service provisioning.

There are many existing Quality of Service (QoS) schemes such as VLAN and DiffServ. However, they are Layer 2 or 3 mechanisms which are hard to scale to end-to-end applications without a network management and configuration agent like ABPD.

There are many application scenarios that can demonstrate the usage of ABPD. For example, a service provider may want some of its user traffic be protected from server or link failures while other traffic not. When a server or link failure happens, the traffic that needs protection is steered to a protection path. The ABDP provides an excellent option to achieve this function. Specifically, application users may send a request to ABPD for traffic that requires protection. The ABPD can check its policy and database to decide whether it can accept the request. If yes, the ABPD will decide the working and protection routes, assign an ID, and inform underlying physical networks to install the required mechanism and resources. After receiving confirmation from all the network elements, the ABPD will respond to the users with the ID. When the application users send traffic that requires protection, it will encapsulate the ID into their packets. When a network failure happens, network elements can route the traffic through backup path by identifying the traffic through the ID.

4. IANA Considerations

It is recommended that IANA assign a port in UDP and another port number in TCP to identify the existing of SFLs in Layer 5. The top level SFL of a SFL stack can use all existing port number assignments to identify various applications.

Authors' Addresses

Changcheng Huang
Department of Systems and Computer Engineering
Carleton University
1125 Colonel By Drive
Ottawa, ON K1S 5B6
Canada
Email: huang@sce.carleton.ca

Jiafeng Zhu
Huawei Technologies Inc
Santa Clara, CA
US
Email: Jiafeng.zhu@huawei.com

Peng He
Ciena Corp
Email: phe@ciena.com

Will(Shucheng) Liu
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China
Email: liushucheng@huawei.com

Baek-Young Choi
Department of Computer Science and Electrical Engineering
University of Missouri-Kansas City
550D FH, 5110 Rockhill Road
Kansas City, MO64110
Email: choiby@umkc.edu