Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 23, 2015

G. Karagiannis
Huawei Technologies
Q. Sun
China Telecom
Luis M. Contreras
Telefonica
P. Yegani
Juniper Networks
JF Tremblay
Viagenie
J.Bi
Tsinghua University
January 23, 2015

Problem Statement for Shared Unified Policy Automation (SUPA)
draft-karagiannis-supa-problem-statement-04

Abstract

The rapid increase in the amount and type of traffic makes it significantly more challenging for operational and management applications to maintain the network and deploy new services. This is the root cause of one of the major challenges that network operators (service providers, SME, etc) are facing today. The operators are obliged to create a simplified view of their network infrastructure that can help network engineers to use such a simplified model rather than manipulating individual devices. In this context, providing network operators with a set of standard generic YANG-based data models that enable management and automation of services on their network is essential.
This document describes what has to be addressed in order to equip service providers with the means to quickly and dynamically create/query/scale/update/delete the services they want to offer. This may include a variety of different service enabling scenarios and in particular VPN management within a data center or among a set of physical or virtualized data centers that belong to an organization or different organizations.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any

time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 23, 2015.

Table of Contents

1. Introduction

Network operators are faced with networks of increasing size and complexity while trying to improve their quality and availability, as more and more services depend on them. Programmatic ways to configure networks, often called software-defined, are considered by many network operators in order to shift the balance in their favor.

Currently, the separation of development and operation of network technologies leads to slow deployment of network functions/devices and poor user experiences.

Providing means of exposing a view of the network to applications may provide significant improvements in configuration agility, error detection and uptime for operators.

However, the real value behind central configuration schemes lies

within the possible simplification through abstract models provided by such systems to applications and network services running above them (on the so-called northbound side). Well-designed simplified models are able to provide a wide range of granularity for various applications and network services needs, from the lower-level physical network to high-level application services.

1.1 Motivation

The rapid increase in the amount and type of traffic makes it significantly more challenging for operational and management applications to maintain the network and deploy new services. Programmatic ways to configure and operate networks, often called software-defined, are one means used by many network operators to provide significant benefits in design and deployment agility.

The purpose of the SUPA (Shared Unified Policy Automation) working group is to introduce the concepts of multi-level and multi-technology network abstractions to address the current separation between development and deployment operations. Business agility, along with OpEx reduction, cannot be obtained unless it becomes possible to deploy changes as products.
Policy-based management is one way to help do this, but better abstractions of network resources and services are needed to achieve these goals.

Several working groups in IETF such as I2RS (L3/ routing topologies), ALTO (cost maps), SFC (service chain), have already defined various schemes for the configuration of network devices and specific network controllers. However, none of these efforts offer (1) a vendor-neutral standardized scheme for applications to transmit their needs to controllers and (2) a set of generic YANG-based data models that enable management and automation of services.

Figure 1 is copied from [ID.draft-zhou-supa-framework] to show the SUPA framework where applications can communicate with management agents of all types, which can be for example single or multiple management agents. These management agents can use any type of mechanisms for excahning information to and from NEs. In this framework NEs can interact with local or remote management agents (e.g., exchange configuration information, status, etc).

Management agents, exchange configuration information with NEs and derive the actual and detailed network topology model. When an application needs to use this network topology it applies NETCONF [RFC6241] or RESTCONF [ID.draft-ietf-netconf-restconf] and it sends a request to receive a service specific abstraction from the network controller(s). Subsequently, the management agent(s) provides, a service specific abstraction of the network topology to the application, which should be able to meet the requirements imposed by this application. Different types of applications may get different service specific abstractions of the same network topology from the management agent (s). For example, for the same actual network topology, a VPN network service will receive a different service specific abstraction of the network topology, than an inter Data Center (DC) network service. By using policies, e.g., for traffic

steering, the application can instruct the management agent(s) to map the service specific abstractions to the actual (detailed) network topology and NE specific configuration.

The main goal of the SUPA working group is to develop a methodology by which the management and monitoring of network services can be done using standardized policy rules. Three types of YANG data models [RFC6020], [RFC6991] are envisioned, each at a different level of abstraction:

1) model of network at the protocol level (logical topology)
2) model of the service that relates the needs of the service to the physical and/or virtual topology used by the service
3) model of the policy rules for managing the service

In particular, the network is first defined as a topology. A service is then defined as a graph that uses that topology. A set of policy rules is then defined to manage the service. In this approach, the service data models, as well as the policy model, will be derived from a single information model, ensuring that each can be shared and reused as managed objects.
Policy rules will be used to define the operational aspects of both the "southbound" (e.g., controller to network device) and "northbound" (e.g., controller to network application) portions of the service environment. The first example that the working group will focus on will be VPN management.

Following the above described methodology, services can be quickly and dynamically created/deleted/updated, using proper mechanisms for exchanging information between the appropriate NEs. Examples of YANG-based data models for network topologies are provided in [ID.draft-contreras-supa-yang-network-topo].

A YANG Data model for SUPA configuration is provided in [ID.draft-zaalouk-supa-configuration-model].

The document [ID.draft-pentikousis-supa-mapping] describes guidelines for mapping high-level configuration and policy information into device-level configuration.

This document is organized as follows. Section 2 presents the terminology. Section 3 provides a brief overview of the use cases associated with SUPA. The requirements/objectives are provided in Section 4. Section 5 provides the security considerations. The IANA considerations are given in Section 6. Section 7 gives the acknowledgements and Section 8 provides the list of references.

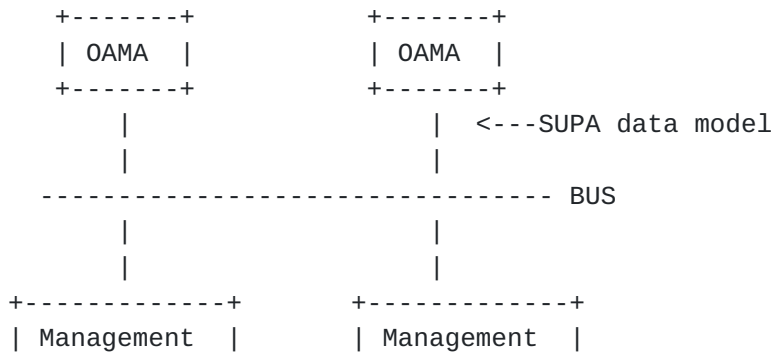

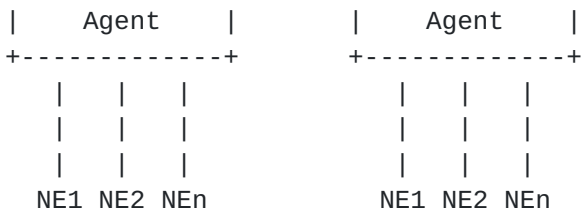

Figure 1: SUPA Framework overview

2. Terminology

Network Service: is the composition of network functions and defined by its functional and behavioural specification. The network service contributes to the behaviour of the higher layer service, which is characterized by at least performance, dependability, and security specifications.

Network Element: a physical or virtual entity that implements one or more network function(s). NEs can interact with local or remote network controllers in order to exchange information, such as configuration information and status.

Service specific abstraction: an abstract view of the actual topology of a network, which is associated with a specific network service type, e.g., VPN or Inter-DC.

3. Use Cases

This section briefly describes the use cases that are associated with different types of network services. The detailed description of these use cases is provided in other Internet draft(s).

A large-scale IDC (Inter Data Center) operator provides server hosting, bandwidth, and value-added services to enterprises and ISPs, and has more than 10 data centers and more than 1Tbs bandwidth in a capital city. In current IDC networks, traffic is routed by configuring policy routes and adjusting routes prioritization to choose an outgoing link. Furthermore, the link bandwidth resources in the data centers are not efficiently utilized. Services usually do not have consistent bandwidth requirements at all times of a day, e.g. video ISP usually require more bandwidth at non-working hours but require less bandwidth at working hours. Some customers have relative high QoS requirement for their services, e.g. IM (Instant Messaging). Such scenarios may be worth modeling since Static bandwidth and QoS provisioning for all the customers and services is not reasonable and not a cost-effective solution.

Operators such as China Telecom, are testing and implementing the DTS(DC Traffic Schedule) schema now. Due to the rapid development of Internet services, each single physical DC can not meet the requirements of an given service system, a general model is that service instances hosted in multiple DCs collaborate to provide services to end-users correspondently,inter-DC traffic increase dramatically during the last several years. More specifically, the first driven factor to implement DTS is the scale of services, generally, single physical DC can not provide all the resources for large-scale service at all time which require the physical DCs to form a virtual DCso the system can apply for the resource for a

'single' virtual DC more flexible and scalable. Another factor is for reliability and security of services, for instance service instances of an given service located in different DCs will exchange large volume of production data for backup and filing, which may occur at a fixed or non-fixed time of each day. In such a case, an management system monitors traffic volume on the link conveying the exporting traffic of a DC. When the volume exceeds the threshold set by the system, the system designs traffic adjustment system to move the overflowling traffic from that link to another exporting link in

order to make sure that the traffic volume on the first link below the threshold. Such scenarios are well worth modeling as operators need to design flexible adjustment policies for optimizing the throughput of DC exporting router.

There are requirements from campus network operators to flexibly manage traffic for multiple functions in a building, such as traffic for network operation, traffic for building monitoring network, traffic for professor working on test-bed/data for different research projects. Traditionally, the operation staffs manually set up VLANs for different users. However, the increasing number of projects/users makes it becoming very hard to manually set up those different network/test-beds in the shared building LAN, because sometimes one office having multiple rights to access different networks/projects. Therefore, SUPA could potentially proving flexible VPN set up on the shared infrastructure (based on IP/MAC address, VLAN ID, etc.). In this case, a controller and standardized northbound APIs could serve for an application for operators to flexibly set up the access to different resources. In general, SUPA will help operators or service providers to design flexible adjustment policies for optimizing the throughput of layer two devices.

SUPA can be used to request the optimization of the traffic paths dynamically and has the ability to request load balancing between data centers and links, and direct customer traffic via network management policies. Path optimization can be accomplished using data models or software programs routines to differentiate customer based on their service class and/or QoS requirements.

Moreover, when VPN tunnels are interconnecting DCs, SUPA can be used to dynamically reconfigure these VPN tunnels, e.g., L2VPN or L3VPN in order to avoid possible congested communication paths and improve end to end latency. Detailed descriptions of these use cases are provided in [ID.draft-cheng-supa-ddc-use-cases].

Currently, there are VPCs (Virtual Private Clouds) that can support a various number of applications. The VPCs need to securely access services running in a data center, i.e., services that are not being exposed to the general Internet. VPNs have been the mechanism of choice to secure these connections. The number of VPCs accessing services running on data centers is significantly increasing. This increases the complexity of managing the VPNs supporting the secure connections between the large number of VPCs and the data centre.

4. Requirements/Objectives

The SUPA architectural framework must support the following capabilities:

1) Define a Yang model that represents the logical topology and capabilities of a network within a single administrative domain. This will later be expanded to accommodate multiple administrative domains.

2) Define a Yang model that maps network services to the capabilities of a network.

3) Define a Yang model that specifies how policy rules may control the operational, administrative, and management aspects of a network service.

5. Security Considerations

Security is a key aspect of any protocol that allows state installation and extracting detailed configuration states of network elements. This places additional security measures on SUPA (e.g., authorization, and authentication of network services) that needs further investigation.

6. IANA Considerations

This document has no actions for IANA.

7. Acknowledgements

The authors of this draft would like to thank the following persons for the provided valuable feedback and contributions: Diego Lopez, Spencer Dawkins, Jun Bi, Xing Li, Chongfeng Xie, Benoit Claise, Ian Farrer, Marc Blancet, Zhen Cao, Hosnieh Rafiee, Mehmet Ersue, Simon Perreault, Fernando Gont, Jose Saldana, Tom Taylor, Kostas Pentikousis, Juergen Schoenwaelder, Eric Voit, Scott O. Bradner.

Tina Tsou and Will Liu contributed to an early version of this draft.

8. References

8.1. Normative References

8.2. Informative References

[ID.draft-cheng-supa-ddc-use-cases] Y. Cheng, C. Zhou, G. Karagiannis, JF. Tremblay, "Use Cases for Distributed Data Center Applicatinos in APONF", IETF Internet draft (Work in progress), draft-cheng-supa-ddc-use-cases-00, September 17, 2014

[ID.draft-contreras-supa-yang-network-topo] L.Contreras, Andrew Qu, "A YANG Data Model for Network Topologies", IETF draft (work in progress), draft-contreras-supa-yang-network-topo, September 18, 2004.

[ID. draft-pentikousis-supa-mapping] K. Pentikousis, Junru Lin, Yiyong Zha, "SUPA Configuration and Policy Mapping", IETF Internet draft, draft-pentikousis-supa-mapping-00, September 23, 2014

[ID.draft-zaalouk-supa-configuration-model] A. Zaalouk, K. Pentikousis, W. Liu, "YANG Data Model for Configuration of Shared Unified Policy Automation (SUPA)", IETF draft, draft-zaalouk-supa-configuration-model-00, 22 September, 2014

[ID.draft-zhou-supa-framework] C. Zhou, L. M. Contreras, Q. Sun, P. Yegani, "The Framework of Shared Unified Policy Automation (SUPA)", IETF draft, draft-zhou-supa-framework-00, 15 January, 2015

[ID.draft-ietf-netconf-restconf] A. Bierman, M. Bjorklund, K. Watsen, R. Fernando, "RESTCONF Protocol", IETF Internet draft (work in progress), draft-ietf-netconf-restconf-01, July 2014

[RFC6020] M. Bjorklund, "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, October 2010.

[RFC6241] R. Enns, M. Bjorklund, J. Schoenwaelder, A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.

[RFC6991] J. Schoenwaelder, "Common YANG Data Types", RFC 6991, July 2013.

Authors' Addresses

Georgios Karagiannis
Huawei Technologies
Hansaallee 205,
40549 Dusseldorf,
Germany
Email: Georgios.Karagiannis@huawei.com

Qiong Sun
China Telecom
No.118 Xizhimennei street, Xicheng District
Beijing 100035
P.R. China
Email: sunqiong@ctbri.com.cn

Luis M. Contreras
Telefonica I+D
Ronda de la Comunicacion, Sur-3 building, 3rd floor
Madrid 28050
Spain
Email: luismiguel.contrerasmurillo@telefonica.com
URI: http://people.tid.es/LuisM.Contreras/

Parviz Yegani
JUNIPER NETWORKS
1133 Innovation Way
Sunnyvale, CA 94089
Email: pyegani@juniper.net

Jean-Francois Tremblay
Viagenie inc.
Email: jean-francois.tremblay@viagenie.ca

Jun Bi
Tsinghua University
Network Research Center, Tsinghua University
Beijing 100084
China

EMail: junbi@tsinghua.edu.cn