

Network Working Group
INTERNET-DRAFT
Intended Status: Informational

J. Bi
Tsinghua University
H. Rafiee
V. Choudhary
J. Strassner
Huawei
D. Romascanu
Avaya
May 19, 2015

Expires: November 19, 2015

Simplified Use of Policy Abstractions (SUPA) Gap Analysis
<draft-bi-sup-a-gap-analysis-03.txt>

Abstract

As operators struggle to optimize their network for different applications while maximizing network resources usage, there's growing business pressure to minimize operational tasks and the deployment time of new services. New automation paradigms are meant to help reach these goals, including the optimization of network functions through application control. This control could be signaled directly by an application, through a proxy or orchestrated in a centralized manner. The purpose of SUPA is to develop a methodology by which network services can be managed using standardized policy rules. SUPA will focus in the first phase on inter-datacenter traffic management as part of the distributed data center use case, including the automated provisioning of site-to-site virtual private networks of various types. This memo analyses the current state of the art of the industries in IETF and outside IETF.

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 November 19, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved. This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Scope and target for SUPA	3
3. Related work within the IETF	4
3.1. I2RS Working Group	4
3.2. L3SM Working Group	4
3.3. ALTO Working Group	5
3.4. TEAS Working Group	5
3.5. BESS Working Group	5
3.6. SFC Working Group	6
3.7. NVO3 Working Group	6
3.8. ACTN Proposed Working Group	6
4. Related work outside the IETF	6
4.1. TM Forum	6
4.2. MEF	7
4.3. Open Daylight	8
4.3.1. Network Intent Composition (NIC)	8
4.3.2. Group Policy	8
4.4. Open Networking Foundation	8
4.5. OpenStack	8
4.5.1. Group-Based Policies	9
4.5.2. Congress	9
4.6. The NEMO Project	9
4.7. The Floodlight Project	9
4.8. The ONOS Project	10
5. Discussion	10
6. Security Considerations	10
7. IANA Considerations	11
8. Acknowledgements	11
9. References	11
9.1. Normative	11
Authors' Addresses	12

1. Introduction

Network operators, including Internet Service Providers, Datacenters operators and others, are under constant pressure to optimize the usage of their installed network resources while maintaining high availability, complexity and deploying new services at an ever-increasing pace. The introduction of new paradigms aims at reducing these efforts, optimized network resource usage and minimize operational overhead. Such a new paradigm is the deployment of automated network configuration and optimization through the use of two complementary mechanisms that are software abstractions to simplify monitoring and control operations and the increase in programmatic control over the configuration and operation of such networks. Policy-based management can be used to combine these two mechanisms into an extensible framework.

Management applications would benefit from a view of the network that is adapted to their needs and from a policy framework that is efficient and simple to use. Several organizations have started working on protocols and models to be used between controllers and network devices, either physical ones or virtualized. This work started some years ago in a number of different organizations and has spawned a large amount of interest in the networking community. However the definition of interfaces between controllers and applications, the so-called "northbound" side, has seen a lot less progress during the same time. There's a need for management applications to interface with controllers in a simple and elegant way. For this purpose, applications require a way to express their requirements in the form of simple policy statements applied to network elements. These network elements should be as simplified (abstracted) as possible for their manipulation by the application. The responsibility of providing an abstract and simple view adapted to each application need is the burden of the controller. The goal of the Simplified Use of Policy Abstractions (SUPA) group is to develop a methodology by which network services can be managed and automated by using a set of information policy model and how these model can map to YANG-based service and policy data models. It also focus on how to communicate these policy models. SUPA will focus in the first phase on inter- datacenter traffic management as part of the distributed data center use case, including the set of information models required to construct an extensible, policy-based framework. These information models will lead to a set of core YANG data models for a policy-based management framework to monitor and control network services.

2. Scope and target for SUPA

SUPA introduces the concepts of multi-level (multiple levels of abstraction) and multi-technology (e.g., IP, VPNs, MPLS) network abstractions to address the current separation between development and deployment operations. Multiple levels of abstraction enable

common concepts present in different technologies and implementations to be represented in a common manner. This facilitates using diverse components and technologies to implement a network service. The following standard generic YANG-based service and policy data models are within the scope of SUPA working group:

- o model of the physical and virtual network topology including the resources (e.g., data rate or latency of links) and operational parameters needed to support service deployment over the network topology.
- o model of the network service (e.g., VPNs) and the network resources required by the network service to be correctly deployed and executed on the physical and/or virtual topology.
- o model of policy rules for managing the network service and mapping services dynamically to the network topology and network resources.

Using the above models, service specific policy data models will be derived from a generic policy model, ensuring that policies have a common structure and can be easily reused as managed objects.

3. Related work within the IETF

3.1. I2RS Working Group

They are not working on interconnection of data centers and considering multi-tenant environment where having a possibility that each tenant control (config, modify, etc.) its whole network that might be physically located on different data centers simply without even the need to involve in its complex communication processes. In other word, SUPA wants to serve a user a service and the interaction to a user is also important. This is not true for I2RS as it focuses on the processes and uses programmable synchronize interfaces to transfer states to and out of the internet routing systems. This is true that I2RS WG also uses the Yang model, however, the model introduced in [yang-i2rs] is so general and not only specific to use cases defined in I2RS. In other word, for I2RS, yang model can help a network controller to understand the topology of the whole network and compare it with what it has and update the topology as needed. Therefore, the general model introduced in [yang-i2rs] can also be used as a base for SUPA.

3.2. L3SM Working Group

This working group focuses on communication of operators and customers by allowing customers to configure the network elements via layer 3 VPN technology. The proposal is very specific about using layer 3 VPN technology via MPBGP. This group also wants to use Yang model to be able to configure network devices.

The differences of this group with SUPA is as followings:

- SUPA proposes a generic proposal for various VPN technologies like L2VPN, L3VPN and composite VPNs. Moreover, the proposed framework is

flexible enough to meet the requirement of any of the existing or upcoming VPN technologies.

- L3SM is more inclined towards MPLS/BGP VPN usecase but SUPA does not focus on a specific use case.
- L3SM focus only on configuration and has no provision for monitoring but SUPA provides service monitoring flexibility.
- L3SM charter did not explain anything about having a network controller and only focuses on device configuration via a L3VPN. In other word, customers might need to have different L3SM to configure different devices. While in SUPA, a management system would allow a customer to configure all or any selected devices concurrently via a network control.

The result of L3SM might be able to feed SUPA with their model to support policy information exchange in Layer 3 and SUPA might want to extend their model to use for SUPA-specific purposes.

3.3. ALTO Working Group

The ALTO working group defined an architecture for exposing topology information, more specifically the cost of paths through an infrastructure, as defined in [RFC7285]. ALTO services are able to provide network maps defined as groups of endpoints. Endpoints are providers-defined entities and can therefore represent any granularity of network, from the physical to groups of networks following similar paths or restrains. Although this model can represent different levels of abstraction at multiple granularities, it's not clear if it could be adapted easily for other purposes than providing cost maps in the context of ALTO. The ALTO model is meant to be used outside of the trust domain of an ISP toward external clients.

3.4. TEAS Working Group

The Traffic Engineering Architecture and Signaling working group is responsible of MPLS-based Traffic Engineering, in other words the control of traffic flows in an MPLS network. It covers YANG models for a traffic engineering database, in coordination with other working groups (I2RS) providing YANG models for network topologies.

3.5. BESS Working Group

The BGP Enabled Services working groups aims at providing a protocol for the provisioning of L3VPN and L2VPN solutions based on BGP. This includes BGP-enabled solutions for datacenter networking and extensions to BGP-enabled solution to support Service Function Chaining. The working group is also chartered to work on on BGP

extensions to YANG models and data models for BGP-enabled services.

3.6. SFC Working Group

The Service Function Chaining (SFC) working group defines a mechanism where traffic is classified before going through an ordered set of services. The set of services is a definite and ordered group of instances defining a service function path. More than one instance may exist for each service in order to allow for load-balancing. A YANG definition for SFC is already proposed in [sfc-yang] and has been subject to an early implementation in Open Daylight. This interface and its model, as currently defined, is an abstraction limited to the scope of service chains.

3.7. NVO3 Working Group

The NVO3 group proposes a way to virtualize the network edge for datacenters in order to be able to move virtual instances without impacting their network configuration. This is realized through a centrally controlled overlay layer-3 network, as described in draft-lasserre-nvo3-framework. At first sight, there doesn't seem to be an overlap between this work and what is being proposed in SUPA. This type of architecture could support a virtual tenant model similar to what is proposed in Open Daylight, but does not offer policing or new models for applications to use.

3.8. ACTN Proposed Working Group

The ACTN proposed work, as described in [actn] framework, has two main goals, the abstraction of multiple optical transport domains into a single controller offering a common abstract topology and the splitting of that topology into abstract client views, which are usually a fraction of the complete network. The ACTN work is therefore about unification of several physical controllers in a virtual one and also about the segmentation, isolation and sharing of network resources. The ACTN work is not explicitly about policies, but some level of policing is applied in the creation of a client view and the way it interacts with the virtual controller beneath. One point where overlap may exist with some of the work proposed in SUPA is in the definition of multiple levels of abstract topologies.

4. Related work outside the IETF

4.1. TM Forum

The TM Forum (a.k.a., the TeleManagement Forum) develops standards and best practices, research, and collaborative programs focused on digital business transformation. It consists of three major programs:

- 1) Agile Business and IT
- 2) Customer Centricity (experience)
- 3) Open Digital Ecosystem

Of these, the ZOOM (Zero-touch Orchestration, Operations, and Management) project, located in the Agile Business and IT project, is the main sub-project in this area that is of interest to SUPA.

Within ZOOM, the Foundational Studies project contains work on an information model and management architecture that are directly relevant to SUPA. The Information Model, Policy, and Security working groups are involved in this work.

The ZOOM information model updates the existing Shared Information and Data (SID) information model to add support for the management of physical and virtual infrastructure, event- and data-driven systems, policy management (architecture and model), metadata for describing and prescribing behavior that can support changes at runtime, and access control.

The policy information model defines event-condition-action (ECA), declarative (intent-based), utility function, and promise policies. The work in draft-strassner-supra-generic-policy-info-model-01 is based on this work. It currently extends the ECA model and provides additional detail not currently present in ZOOM; the next version of this draft will do the same for declarative policies.

There is currently no plan to use the utility function and promise policies.

Finally, it should be noted that the data model work planned for SUPA is not currently planned for the ZOOM project.

4.2. MEF

The MEF (originally named the Metro Ethernet Forum) develops architecture, service and management specifications related to Carrier Ethernet (CE). The CE architecture includes the definition of several interfaces specific to CE like the User Network Interface (UNI) and External Network Network Interface (ENNI). Specifications developed in this space include the definitions of CE services, CE service attributes, Ethernet Access Services, Class of Service, OAM and Management interfaces, Service Activation and Test.

The more recent vision of the MEF related to the future of networking is described as The Third Network and includes plans to develop Lifecycle Service Orchestration with APIs for existing network, NFV and SDN implementations enabling Agile, Assured and Orchestrated

Services. This stage of the MEF activity is now in early phases with focus on architectural work.

The MEF has developed a number of Information and Data Models, and has recently started a project that used YANG to model and manage the services covered by the MEF. Although the MEF has created quite rigorous definitions of these services, these are transport technology specific, and they do not include and rely on policies.

4.3. Open Daylight

Open Daylight network controller implements a number of models through its service abstraction Layer (MD-SAL) based on draft IETF Yang models. Few of the below mentioned Open Daylight projects provides policy abstraction and better flexibility to the user.

4.3.1. Network Intent Composition (NIC)

Network Intent Composition project aims at providing better flexibility and high-level interface for the specification of policies. The intents-based interface would provide a high level of abstraction easy to formulate by an application developer and completely detached from the underlying implementation details. By making intents portable and composable, the project aims at making intents a more scalable approach than existing interfaces.

4.3.2. Group Policy

The group-based policy project defines an application-centric policy model for Open Daylight that separates information about application connectivity requirements from information about the underlying details of the network infrastructure.

4.4. Open Networking Foundation

The ONF created a group responsible of defining northbound interfaces, but this hasn't lead to the publication of standards in this area so far. A blog entry on the ONF web site showed an interest in using the principle of intents at ONF, but no details were provided on the status of this project. The membership of this group being closed in nature, the status of current draft proposals is not known.

4.5. OpenStack

OpenStack software controls large pools of compute, storage, and networking resources throughout a datacenter, managed through a dashboard or via the OpenStack API. OpenStack works with popular

enterprise and open source technologies making it ideal for heterogeneous infrastructure. Few of the below mentioned OpenStack projects provides policy abstraction and better flexibility to the user.

4.5.1. Group-Based Policies

The Group-Based Policies project for OpenStack Neutron is built around entities assembled in Endpoints Groups (EPG) that provide or consume Contracts. Such Contracts are hierarchical entities containing policy rules. A first version was released in January 2015, based on the Juno release. This type of approach is more relational than declarative, but could be used to describe a large amount of possible scenarios. It has the advantage of providing a relatively simple policy model that covers a large applicability. From an OpenStack point of view, the scope of GBP is limited to networking within the Neutron module.

4.5.2. Congress

The Congress project within OpenStack provides a way to formulate complex policies using the Datalog language, a derivate of Prolog. Datalog is entirely declarative and first-order logic, which gives it interesting properties, such as providing the same result no matter the order in which the statements are made. The language allows for the definition of types and for active enforcement or verification of the policies. Using Datalog also allows Congress to take advantage of the significant body of knowledge and experience relating to declarative languages and their implementation. The Congress policies aim at manipulating objects exposed by multiple OpenStack modules and is therefore larger in scope than network policying only. The only drawback of this approach lies in its potentially large computational complexity, which could limit its ability to react in real time fast events such as those relating to the network.

4.6. The NEMO Project

The NEMO project is a research activity aiming at defining a simple declarative framework for networking. The NEMO syntax is not based on an existing language and covers the basic elements for network manipulation such as nodes, links and flows. The NEMO project has been successfully demonstrated at IETF-91, along with a companion graphical user interface, and this work now being proposed as the base for a new group called Intent-Based NEMO (IBNEMO) within the IETF.

4.7. The Floodlight Project

The Floodlight is an openflow enabled SDN controller. it uses another

open source project called Indigo to support openflow and manage southbound devices. Indigo agent also supports abstraction layer to make it easy to integrate with physical and virtual switches. It supports configuration of abstraction layer so that it can configure openflow in hybrid mode.

4.8. The ONOS Project

The ONOS is a SDN controller. It supports abstraction for both southbound and northbound interfaces. This is because NFV used in ONOS can reduce CapEx because each service can be virtualized, but it increases OpEx as service providers now have to contend with increased management complexity due to the management and orchestration of a large numbers of VMs on commodity servers, the management of network function software on the VMs and how VMs must be interconnected based on subscriber's service contract. This is why, ONOS uses Network Function as a Service (NFaaS) that not only virtualized the components and Virtual Network Functions (VNFs) but also introduces an abstractive unit for a collection of VMs and their interconnecting network(s). Being able to create and manipulate these units, rather than handling individual components, significantly simplifies operation. NFaaS manipulates these units in the enhanced form of a service.

5. Discussion

The ongoing projects outside of the IETF (see Section 4) demonstrate that there is a need to develop service level abstractions and policies that govern their implementation and mapping to the underlying network infrastructure. While different approaches are currently being prototyped, it is desirable from an operator's perspective and of likely also of strategic importance from an IETF's perspective to host work in this area within the IETF with a goal to drive progress towards a common standardized solution in this space. Generic policy driven service management is not directly worked on by existing IETF working groups. Several working groups provide technology specific mechanisms (TEAS, BESS, ACTN) that ideally can be leveraged by a generic policy driven service management solution. Other working groups provide key building blocks (e.g., the generic topology work recently chartered in the I2RS working group) or they look at specific aspects such as the chaining of data plane traffic manipulation functions (SFC) or the movement of virtual machines (NVO) or the export of typically aggregated topology information to distributed file sharing or streaming applications (ALTO).

6. Security Considerations

TBD

7. IANA Considerations

There is no IANA consideration

8. Acknowledgements

Jean-Francois Tremblay from Viagenie contributed to some significant portions of this text. James Huang, Oliver Huang, Will Liu, Yiyong Zha and Dacheng Zhang helped in providing valuable comments and text.

9. References

9.1. Normative References

- [RFC6020] Bjorklund, M., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, October 2010.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.
- [RFC7285] Alimi, R., Penno, R., Yang, Y., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, September 2014
- [yang-i2rs] Medved, J., Varga, R., Tkacik, T., Bahadur, N., Ananthakrishnan, H., "A Data Model for Network Topologies", <https://tools.ietf.org/html/draft-ietf-i2rs-yang-network-topo-00>, April 2015
- [yang-l3sm] Litkowski, S., Shakir, R., Tomotaki, L., D'Souza, K., "YANG Data Model for L3VPN service delivery", <https://tools.ietf.org/html/draft-l3vpn-service-yang-00>, February 2015

Authors' Addresses

Jun Bi
Tsinghua University
Network Research Center, Tsinghua University
Beijing 100084
China
E-mail: junbi@cernet.edu.cn

Hosnieh Rafiee
Huawei Technologies Duesseldorf GmbH
Munich, Germany
Phone: +49 (0)162 204 74 58
E-mail: ietf@rozanak.com

Vikram Choudhary
Huawei Technologies
E-mail: vikram.choudhary@huawei.com

Dan Romascanu
Avaya
Atidim Technology Park, Bldg. #3
Tel Aviv 61581, Israel
Phone: +972-3-6458414
E-mail: dromasca@avaya.com

John Strassner
Futurewei Technologies
US R&D Center
2330 Central Expressway
Building A, office A2-2143
Santa Clara, California 95050
E-mail: john.sc.strassner@huawei.com

