Network Working Group
Internet-Draft
Intended status: Informational
Expires: 20 March 2026

O. G. D. Dios
Telefonica
S. Barguil
Nokia

16 September 2025

An Update of Service and Network YANG Data Models draft-bg-onions-update-network-service-models-00

Abstract

Service & Network data models have been implemented in recent years to facilitate the deployment of connectivity services such as Layer 2 and Layer 3 VPN services in provider networks. This document reports the findings from the implementations, including missing functionalities, configuration blocks aligment against recent network models published, operational issues/limitatations and enhancements.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://oscargdd.github.io/lxnm-bis/draft-bg-onions-update-network-service-models.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-bg-onions-update-network-service-models/.

Source for this draft and an issue tracker can be found at https://github.com/oscargdd/lxnm-bis.

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 https://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 20 March 2026.

Table of Contents

1. Introduction

Service and Network YANG data models [RFC8199][RFC8309] such as the Layer 3 Network Model (L3NM) [RFC9182] and the Layer 2 Network Model (L2NM) [RFC9291] have been implemented to automate the deployment of VPN services by providers. This document reports the findings from the implementations, deriving the functionalities required to update the Service and Network YANG data models.

[RFC8969] documents the automation framework. [RFC9315] documents Intent-based networking from IRTF perspective, with specific problems which are addressable today after the first deployments have been done.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Service and Network YANG Data Models in the IETF

Several IETF Working Groups have developed YANG modules in order to foster the provisioning and, more generally, the deployment of services. These modules focus on how the network operator intents to manage a network through protocols and devices to deliver a service. The intended configuration at the device level is derived from network YANG data models. The customer service YANG data models abstract the service for upper layers. The intended network service configuration is derived from the the service model.

A set of these models is listed here:

* [RFC9182] As a complement to the Layer 3 Virtual Private Network Service Model (L3SM), which is used for communication between customers and service providers, L3NM is a Network Model (L3NM) that can be used for the provisioning of BGP based Layer 3 Virtual Private Network (L3VPN) services within a service provider network.

* [RFC9291] documents a data model that describes the deployment of various types of L2VPN, including VPWS and BGP based L2VPN, such as EVPNs.

4. Observations and new requirements

4.1. Enhancements to LxSM and LxNM

Implementations of LxNM models in controllers required new functionalities which were not covered in [RFC9182] and [RFC9291] to deploy the missing functions in the Operator services. This section compiles the functions that were reported by those implementations.

4.1.1. L3NM Enhancements

* BFD parametrization of static routes (Github issue #1):

  - The L3NM Yang data model allows to manage static routes in a VPN. That is, for a particular VPN service, new Pv4 and IPv6 static routes can be added, modified or deleted. The data

model allows to specify whether BFD is desired in the static route. Whenever a controller derives the device configuration of the static route it will need to decide a particular BFD configuration, typically from a pre-defined template. Operators required, for different services, to customize the main BFD parameters to allow, for example, faster detection for critical services. The new requirement is the ability to specify BFD intended configuration in the IPv4 and IPv6 static routes, including a required-min-rx-interval and multiplier.

* Management of VLAN 0 in tagged interfaces (Github issue #2): LxNM Yang models have a range defined for cvlan between 1 and 4094. VLAN 0 should also be supported and is used in deployments.

* Missing BGP intended configuration blocks (in relation to Attachment Circuits) (Github issue #3).

  - There are a set of BGP configuration blocks required to manage BGP based services which are present now in the AC-Model but not in the L3NM:

    o BGP Peer group creation

    o BGP Redistribution rules

  - Missing pointer to ACL (also present in Github issue #3).

    o ACL pointer to attach forwarding filter

* SRv6 support for L3VPN (Github issue #15): SRv6-based BGP services including L3VPN, whose procedures are defined in [RFC9252]

* Improving Multicast Support:

  - For L3VPN with multicast, one implementation has reported that Cisco MVPN augmentation were added to include various profiles (ipmsi and spmsi) . There is no YANG module from IETF as of today that supports full MVPN/SPMSI/IPMSI under L3NM directly. Standardized profiles are required to be added.

* Extend guidance of how the network models can be used to/and operationalize Inter-AS VPN options (A, B, and C as defined in [RFC4364]) using the L3NM framework (Github issue #25).

4.1.2. L2NM Enhancements

* EVPN Remote and Local eth-tag (Github issue #6)

* Explicitly assign a RD at node level (Github issue #7)

  - In the model, a RD must be always assigned via profile at service level. It is useful to be able to set a explicit RD directly at node level overriding the value of the profile. This way, a common profile can be used for all the services for use cases where only RD changes per node.

* Add support for Flexible Cross-Connect (FXC) Service ([RFC9744]) (Github issue #8)

* Add explanatory text for EVPN multihoming using LAG (Github issue #9)

* support for vlan-lists/vlan-ranges (Github issue #10)

  - When defining a Layer 2 service, sometimes multiple VLANs are mapped into a given service. It would be good to support this in the L2NM encapsulation stanza. Examples are as follows:

    o Typically used in single-tagged scenarios: vlan-id-list [200 210-219 222 234 240-249];

    o Dual-tagged scenario, with s-vlan=430 and a list of c-vlans: vlan-tags outer 430 inner-list [200 210-219 222 234 240-249];

* SRv6 support for L2VPN (Github issue #15)

* Performance monitoring

  - ITU-T Y.1731 performance monitoring in Ethernet based networks . L2NM itself [RFC9291] doesn’t natively include OAM specifics.

* Add EVI identifier to differentiate it from VPN-ID. Each EVI maps to a specific EVPN service (e.g., a Layer 2 VPN bridging a particular VLAN across the EVPN fabric) (Github issue #24).

4.2. New Functionalities Required to Fully Support Connectivity Services

The realization of advanced connectivity services requires, in addition to the configurations expressed in LxNM models: * Definition of Access Control lists and prefix Sets: + Connectivity services often include mechanisms to filter forwarding packets. The LxNM models allow to include a ’forwarding-profile-identifier’. A forwarding profile refers to the policies that apply to the forwarding of packets conveyed within a VPN. Such policies may

consist, for example, of applying Access Control Lists (ACLs). However, currently there is not an ACL network service model (even though a device level ACL model exists) that allows to manipulate such ACLs and sets when creating the service. + ACLs and Prefix sets can be reused among services. Thus, they need to be handled at a network level, regardless of the actual service using them. * Definition of routing policies, including community sets, as path sets, etc: + Advanced connectivity services require the creation of complex policies. The LxNM models allows to indicate which policy (or policies) should be used. However, LxNM does not include the definition of policies. There are a set a device models which could be used as a base for the network model. * Pre and post checks before and after deploying a service in the network. * Preparation of the interface. Prior to the application of the entire configuration profile.

4.3. Status of the Intended Network Service

The Network Service Yang models represent an intent of the realization of service. A controller, after an instance of the network service yang intent has been created, will derive the necessary device level configurations and apply them to the necessary devices. However, the implementations reported a set of open issues which are related to the status. Those issues are not solved today and are left to implementation choices and solutions.

* Is the Service running on the Network? (Github issue #5)
  - How can the northbound system be assured with 100% certainty that a configuration has been successfully installed on the network device?
  - What mechanisms or feedback loops exist to confirm successful configuration deployment beyond simple acknowledgment?
  - How can the system can handle transient errors or partial configuration applications from the model perspective?
  - Are there specific operational attributes that can be used to reflect the real-time status of the configuration?
  - Does the network element provide any operational state parameters or notifications that indicate whether the configuration is active, pending, or has failed?
  - If a configuration is manually removed via CLI at the network device, is there a mechanism to reflect this change northbound?

* Operational Status Clarification (Github issue #4)

  - Understanding the interrelationships between the operational status of VPN services, VPN nodes, and VPN network access is another significant operational gap. The status of a VPN service may depend on the status of the underlying VPN nodes and the network access provided to the VPN.

  - To address this gap, IETF should establish clear dependencies and correlations between the various operational statuses. This could involve defining specific criteria for determining the overall status of a VPN service based on the status of its constituent VPN nodes and network access components. Moreover, real-time monitoring and correlation of status information can provide insights into the health and performance of VPN services.

4.4. Summary

* L3NM and L2NM need to be updated to cover new technologies (such as SRv6), incomplete support (such as multicast) and enhancements.

* New network YANG data models (such as ACLs and routing policies)

5. Security Considerations

TODO Security

6. IANA Considerations

This document has no IANA actions.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

[RFC9182] Barguil, S., Gonzalez de Dios, O., Ed., Boucadair, M., Ed., Munoz, L., and A. Aguado, "A YANG Network Data Model for Layer 3 VPNs", RFC 9182, DOI 10.17487/RFC9182, February 2022, <https://www.rfc-editor.org/rfc/rfc9182>.

[RFC9291] Boucadair, M., Ed., Gonzalez de Dios, O., Ed., Barguil, S., and L. Munoz, "A YANG Network Data Model for Layer 2 VPNs", RFC 9291, DOI 10.17487/RFC9291, September 2022, <https://www.rfc-editor.org/rfc/rfc9291>.

7.2. Informative References

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <https://www.rfc-editor.org/rfc/rfc4364>.

[RFC8199] Bogdanovic, D., Claise, B., and C. Moberg, "YANG Module Classification", RFC 8199, DOI 10.17487/RFC8199, July 2017, <https://www.rfc-editor.org/rfc/rfc8199>.

[RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", RFC 8309, DOI 10.17487/RFC8309, January 2018, <https://www.rfc-editor.org/rfc/rfc8309>.

[RFC8969] Wu, Q., Ed., Boucadair, M., Ed., Lopez, D., Xie, C., and L. Geng, "A Framework for Automating Service and Network Management with YANG", RFC 8969, DOI 10.17487/RFC8969, January 2021, <https://www.rfc-editor.org/rfc/rfc8969>.

[RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <https://www.rfc-editor.org/rfc/rfc9252>.

[RFC9315] Clemm, A., Ciavaglia, L., Granville, L. Z., and J. Tantsura, "Intent-Based Networking - Concepts and Definitions", RFC 9315, DOI 10.17487/RFC9315, October 2022, <https://www.rfc-editor.org/rfc/rfc9315>.

[RFC9744] Sajassi, A., Ed., Brissette, P., Uttaro, J., Drake, J., Boutros, S., and J. Rabadan, "EVPN Virtual Private Wire Service (VPWS) Flexible Cross-Connect (FXC) Service", RFC 9744, DOI 10.17487/RFC9744, March 2025, <https://www.rfc-editor.org/rfc/rfc9744>.

Acknowledgments

This documents is based on the issues compiled in Github lxnm repository. The authors would like to acknowledge the issues raised by Julian Lucek, Xiao Min, Daniele Ceccarelli and Sujay Murthy.

Authors' Addresses

Oscar Gonzalez de Dios
Telefonica
Email: oscar.gonzalezdedios@telefonica.com

Samier Barguil
Nokia
Email: samier.barguil_giraldo@nokia.com