

Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: November 22, 2019

A. Aguado  
O. Gonzalez de Dios, Ed.  
V. Lopez  
Telefonica  
D. Voyer  
Bell Canada  
L. Munoz  
Vodafone  
May 21, 2019

**Layer 3 VPN Network Model**  
**draft-aguado-opsawg-l3sm-l3nm-00**

Abstract

[RFC 8299](#) [[RFC8299](#)] defines a L3VPN Service Model (L3SM) YANG data model that can be used for communication between customers and network operators. It assumes that there is a monolithic management system with full control of transport resources. This approach (that is valid for the customer to network operator conversation) limits the usage of the model to the role of a Customer Service Model, according to the terminology defined in [RFC 8309](#) [[RFC8309](#)].

There is a need for a YANG model for use between the entity that interacts directly with the customer (service orchestrator) and the entity in charge of network orchestration and control which, according to [RFC 8309](#) [[RFC8309](#)], can be referred as Service Delivery Model. In some cases, the control of the network is further expanded into per- domain control.

This document uses the L3SM model defined in [RFC 8299](#) [[RFC8299](#)], and extends it to facilitate communication between the service orchestrator and transport orchestrator (MSDC), and an MSDC and domain controllers. The resulting model is called the L3VPN Network Model (L3NM).

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 November 22, 2019.

Copyright Notice

Copyright (c) 2019 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 (<https://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 . . . . . [2](#)
- [1.1.](#) Terminology . . . . . [3](#)
- [1.2.](#) Requirements Language . . . . . [3](#)
- [2.](#) Reference architecture . . . . . [4](#)
- [3.](#) Yang model extensions . . . . . [7](#)
- [3.1.](#) Bearer ethernet Encapsulation . . . . . [8](#)
- [3.2.](#) Multi-Domain Resource Management . . . . . [8](#)
- [3.3.](#) Remote Far-End Configuration . . . . . [8](#)
- [3.4.](#) Provide Edge Identification Point . . . . . [9](#)
- [4.](#) Design of the data model . . . . . [9](#)
- [5.](#) Acknowledgements . . . . . [12](#)
- [6.](#) IANA Considerations . . . . . [12](#)
- [7.](#) Security Considerations . . . . . [12](#)
- [8.](#) References . . . . . [12](#)
- [8.1.](#) Normative References . . . . . [12](#)
- [8.2.](#) Informative References . . . . . [12](#)
- Authors' Addresses . . . . . [13](#)

**1. Introduction**

[RFC 8299](#) [[RFC8299](#)] defines a L3VPN Service Model (L3SM) YANG data model that can be used for communication between customers and network operators. Although the intention to provide an abstracted view of the customer's requested services is clear, the assumption is



that the model is applied at the top of a monolithic management system with full control of transport resources. That assumption substantially limits the usage of the L3SM to the role of a Customer Service Model, according to the terminology defined in [RFC 8309](#) [[RFC8309](#)].

This document defines a set of extensions of the YANG model described in [RFC 8299](#) [[RFC8299](#)] via augmentation. The augmentations facilitate the use the resulting model in communications with the transport orchestrator, also known as the MDSC (Multi-Domain Service Coordinator) in the terminology of the framework for Abstraction and Control of TE Networks (ACTN) defined in [RFC 8453](#) [[RFC8453](#)]. The MDSC is functional component responsible for orchestration of the network resources and instigate connections across the operator's networks.

The data model defined in this document is called the L3VPN Network Model (L3NM). It enables further capabilities, such as resource management or to serve as a multi-domain orchestration interface, where transport resources must be synchronized.

This document does not obsolete, but complements, the definitions in [RFC 8299](#) [[RFC8299](#)]. It aims to provide a wider scope for the L3SM via augmentation, but does not attempt to address all deployment cases especially those where the L3VPN connectivity is supported through the coordination of different VPNs in different underlying networks. More complex deployment scenarios involving the coordination of different VPN instances and different technologies to provide end-to-end VPN connectivity is out of scope of this document, but is discussed in [I-D.evenwu-opsawg-yang-composed-vpn].

### **1.1. Terminology**

This document assumes that the reader is familiar with the contents of [RFC 6241](#) [[RFC6241](#)], [RFC 7950](#) [[RFC7950](#)], [RFC 8299](#) [[RFC8299](#)], [RFC 8309](#) [[RFC8309](#)], and [[RFC8453](#)] and uses terminology from those documents. Tree diagrams used in this document follow the notation defined in [[RFC8340](#)].

### **1.2. Requirements Language**

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](#) [[RFC2119](#)].



## **2. Reference architecture**

Figure 1 shows where the L3NM is used in a management stack. The figure is an expansion of the architecture presented in [Section 5 of RFC 8299 \[RFC8299\]](#) and decomposes the box marked "orchestration" in that figure into three separate functional components called "Service Orchestration", "Network Orchestration", and "Domain Orchestration".

At the same time, terminology from [RFC 8309 \[RFC8309\]](#) is introduced to show the distinction between the "Customer Service Model", the "Service Delivery Model", the "Network Configuration Model", and the "Device Configuration Model". In that context, the "Domain Orchestration" and "Config Manager" roles may be performed by "Controllers".



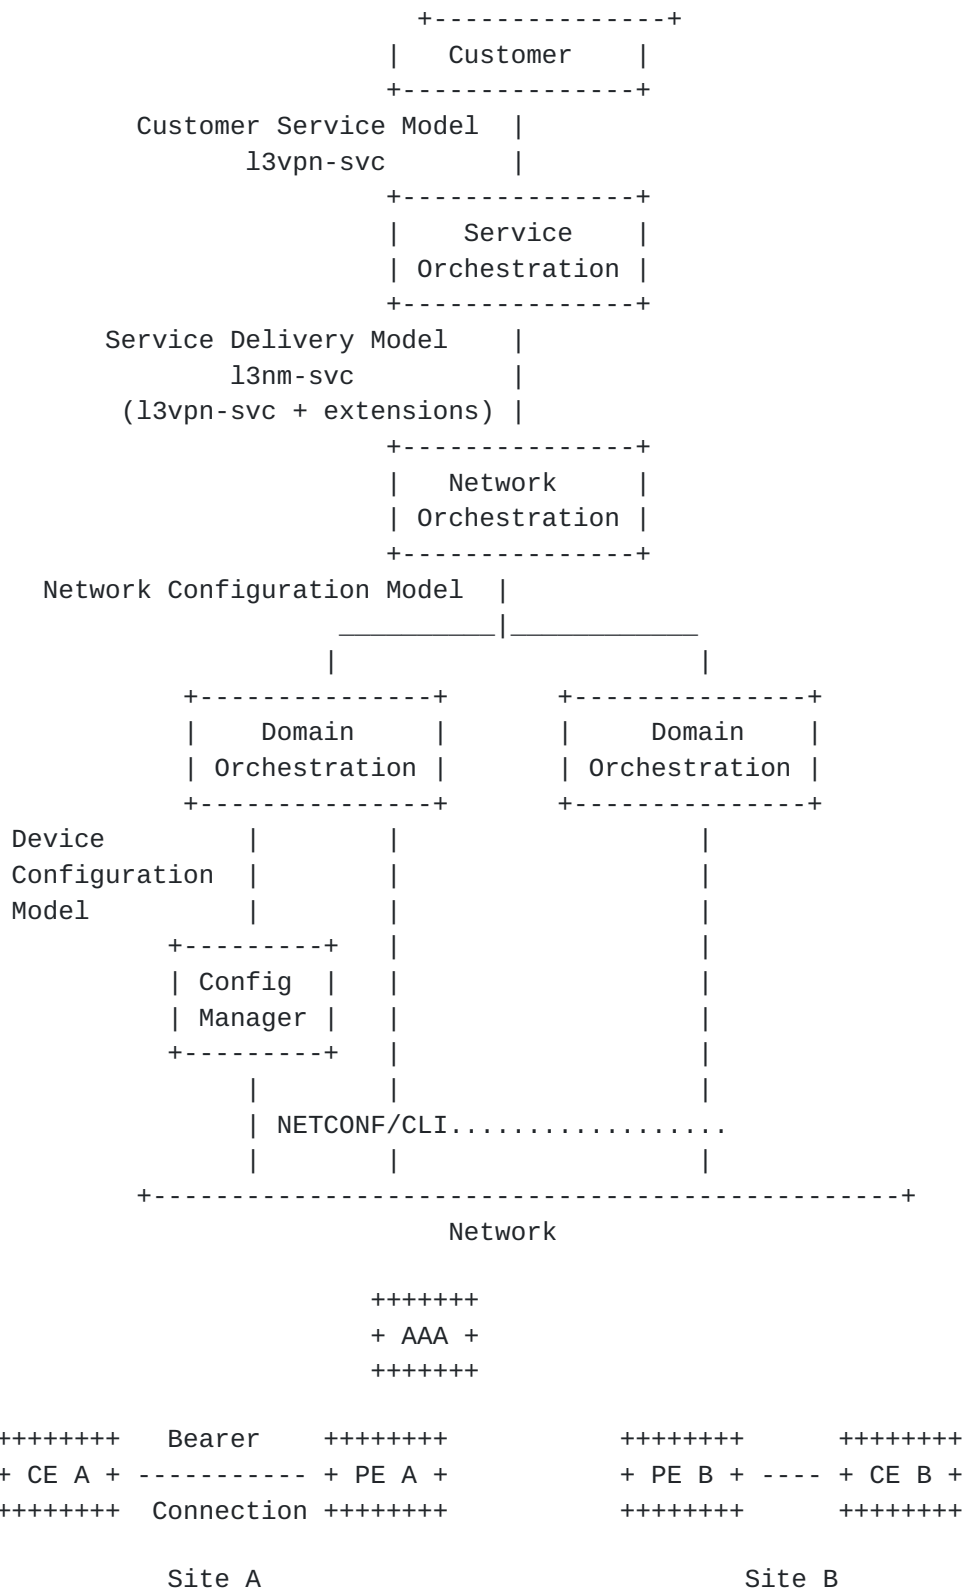


Figure 1: L3SM and L3NM





The L3SM and L3NM may also be set in the context of the ACTN architecture [[RFC8453](#)]. Figure 2 shows the Customer Network Controller (CNC), the Multi-Domain Service Coordinator (MDSC), and the Provisioning Network Controller (PNC). It also shows the interfaces between these functional units: the CNC-MDSC Interface (CMI), the MDSC-PNC Interface (MPI), and the Southbound Interface (SBI).





configuration of PEs not managed by the management system and the definition for PE identification.

### **3.1. Bearer ethernet Encapsulation**

The definition of a L3 VPN is commonly defined not only at the IP layer, but also requires to identify parameters at the Ethernet layer, such as encapsulation (e.g. VLAN, QinQ, QinAny, VxLAN, etc). This specification is not supported in [\[RFC8299\]](#), whilst it suggests that any extension on this direction shall be implemented via augmentation of the bearer container. The extension defined to cope with these parameters uses the connection container inside the site-network-access defined by the the [\[RFC8466\]](#). This container defines protocol parameters to enable connectivity at Layer 2. In the context of L3SM, the augmentation includes only mandatory parameters for the service configuration, which are mainly related to the interface encapsulation. Other definitions from L2SM connection container are left aside. For example, LAG information is not required and it shall be configured prior to the service configuration, being the aggregated interface identified in the model as the bearer-reference, as discussed later in [Section 4.4](#).

### **3.2. Multi-Domain Resource Management**

The implementation of L3 VPN services which spans across administratively separated domains (i.e. that under the administration of different management systems or controllers) requires some network resources to be synchronised between systems. Particularly, there are two resources that must be orchestrated and synchronised to avoid asymmetric (non-functional) configuration, or the usage of unavailable resources. For example, RTs shall be synchronised between PEs. When every PE is controlled by the same management system, RT allocation can be performed by the system. In cases where the service spans across multiple management systems, this task shall be synchronised and, therefore, the service model must allow this specification. In addition, RDs must be also synchronised to avoid collisions in RD allocation between separated systems. A incorrect allocation might lead into same RD and IP prefixes being exported by different PE routers.

### **3.3. Remote Far-End Configuration**

Depending on the control plane implementation, different network scenarios might require additional information for the L3 VPN service to be configured and active. For example, an L3 VPN Option C service, if no reflection of IPv4 VPN routes is configured via ASBR or route reflector, may require additional configuration (e.g. a new BGP neighbour) to be coordinated between both management systems.



This definition requires for every management system participant on the VPN to receive not just their own sites and site-network-accesses, but also to receive information about external ones, identified as an external site-network-access-type. In addition, this particular site-network-access is augmented to include the loopback address of the far-end (remote/external) PE router.

### **3.4. Provide Edge Identification Point**

[RFC8299](#) states that The "bearer-reference" parameter is used in cases where the customer has already ordered a network connection to the SP apart from the IP VPN site and wants to reuse this connection. The string used is an internal reference from the SP and describe the already-available connection. Oftenly, a client interface (either a customer one or an interface used by the SP) is already in place and connected, although it has not being used previously. In some other cases (e.g. for stitching purposes), the termination of a VPN service is done over logical terminations within a PE router.

The bearer-reference must serve as a strict unequivocal parameters to identify the connection between a PE and a client (CE). This means that, despite the type is maintained as a string and there is no restriction in the way this data is formed, the bearer-reference must serve as the unique way to identify the PE router and the client interface. This, together with the encapsulation augments proposed in 4.1, serves as the way to identify the client interface and configure L2 specific parameters.

## **4. Design of the data model**

The augments defined in this document are organised per scenario, as per defined in [Section 4](#). The case described 4.4 does not need any further extension of the data model and only requires a more restricted definition on how the data model is used for PE router and client port identification, so no augment is implemented for this scenario.

The augments implemented are distributed as follows. The first augment implements the extensions for RT and RD definition for the L3 VPN, following the YANG definitions from BESS-L3VPN. The second augment copes with the information from a remote PE not directly under the management system supervision. This augment does not follow any previously defined model and includes the loopback IP address of the external router. The last augment includes information below layer 3 that is required for the service. In particular, we include information related to clients interface encapsulation and aggregation.





The high-level model structure proposed by this document is as shown below:

|----- EXAMPLE -----|

```

module: ietf-l3vpn-svc-ext
  augment /l3vpn-svc:l3vpn-svc/l3vpn-svc:sites/l3vpn-svc:site/l3vpn-svc:site-
network-accesses/l3vpn-svc:site-network-access:
  +-rw transport
    +-rw vpn-targets
      +-rw vpn-target* [route-target]
        | +-rw route-target          rt-types:route-target
        | +-rw route-target-type     rt-types:route-target-type
      +-rw route-policy? -> /rt-pol:routing-policy/policy-definitions/
policy-definition/name
  augment /l3vpn-svc:l3vpn-svc/l3vpn-svc:sites/l3vpn-svc:site/l3vpn-svc:site-
network-accesses/l3vpn-svc:site-network-access:
  +-rw far-end
    +-rw address?  inet:ip-address
  augment /l3vpn-svc:l3vpn-svc/l3vpn-svc:sites/l3vpn-svc:site/l3vpn-svc:site-
network-accesses/l3vpn-svc:site-network-access/l3vpn-svc:bearer:
  +-rw ethernet
    +-rw connection
      +-rw encapsulation-type?  identityref
      +-rw eth-inf-type?        identityref
      +-rw tagged-interface
        | +-rw type?            identityref
        | +-rw dot1q-vlan-tagged {dot1q}?
        | | +-rw tg-type?      identityref
        | | +-rw cvlan-id      uint16
        | +-rw priority-tagged
        | | +-rw tag-type?     identityref
        | +-rw qinq {qinq}?
        | | +-rw tag-type?     identityref
        | | +-rw svlan-id      uint16
        | | +-rw cvlan-id      uint16
        | +-rw qinany {qinany}?
        | | +-rw tag-type?     identityref
        | | +-rw svlan-id      uint16
        | +-rw vxlan {vxlan}?
        | +-rw vni-id          uint32
        | +-rw peer-mode?     identityref
        | +-rw peer-list* [peer-ip]
        | +-rw peer-ip        inet:ip-address
      +-rw untagged-interface
        | +-rw speed?          uint32
        | +-rw mode?           neg-mode
        | +-rw phy-mtu?        uint32
        | +-rw lldp?           boolean
        | +-rw oam-802.3ah-link {oam-3ah}?
        | | +-rw enabled?     boolean
        | +-rw uni-loop-prevention?  boolean

```

Figure 3

Aguado, et al.

Expires November 22, 2019

[Page 11]

## **5. Acknowledgements**

Thanks to Adrian Farrel and Miguel Cros for the suggestions on the document

## **6. IANA Considerations**

This memo includes no request to IANA.

## **7. Security Considerations**

All the security considerations of [RFC 8299](#) [[RFC8299](#)] apply to this document. Subsequent versions will provide additional security considerations.

## **8. References**

### **8.1. 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/info/rfc2119>>.

### **8.2. Informative References**

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", [RFC 6241](#), DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.

[RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", [RFC 7950](#), DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.

[RFC8299] Wu, Q., Ed., Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", [RFC 8299](#), DOI 10.17487/RFC8299, January 2018, <<https://www.rfc-editor.org/info/rfc8299>>.

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



Authors' Addresses

Alejandro Aguado  
Telefonica  
Madrid  
ES

Email: [alejandro.aguadomartin.ext@telefonica.com](mailto:alejandro.aguadomartin.ext@telefonica.com)

Oscar Gonzalez de Dios (editor)  
Telefonica  
Madrid  
ES

Email: [oscar.gonzalezdedios@telefonica.com](mailto:oscar.gonzalezdedios@telefonica.com)

Victor Lopez  
Telefonica  
Madrid  
ES

Email: [victor.lopezalvarez@telefonica.com](mailto:victor.lopezalvarez@telefonica.com)

Daniel Voyer  
Bell Canada  
CA

Email: [daniel.voyer@bell.ca](mailto:daniel.voyer@bell.ca)

Luis Angel Munoz  
Vodafone  
ES

Email: [luis-angel.munoz@vodafone.com](mailto:luis-angel.munoz@vodafone.com)

