

L3SM Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 4, 2016

S. Litkowski
Orange Business Service
R. Shakir
BT
L. Tomotaki
Verizon
K. D'Souza
ATT
August 03, 2015

YANG Data Model for L3VPN service delivery
draft-ietf-l3sm-l3vpn-service-model-01

Abstract

This document defines a YANG data model that can be used to deliver a Layer 3 Provider Provisioned VPN service. The document is limited to the BGP PE-based VPNs as described in [RFC4110](#) and [RFC4364](#). This model is intended to be instantiated at management system to deliver the overall service. This model is not a configuration model to be used directly on network elements. This model provides an abstracted view of the Layer 3 IPVPN service configuration components. It will be up to a management system to take this as an input and use specific configurations models to configure the different network elements to deliver the service. How configuration of network elements is done is out of scope of the document.

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 [[RFC2119](#)].

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 February 4, 2016.

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
1.1.	Terminology	3
1.2.	Tree diagram	4
2.	Definitions	4
3.	Layer 3 IP VPN service model	4
4.	Service data model usage	5
5.	Design of the Data Model	6
5.1.	VPN service overview	12
5.1.1.	VPN service topology	13
5.1.1.1.	Route Target allocation	13
5.1.1.2.	Any to any	14
5.1.1.3.	Hub and Spoke	14
5.1.1.4.	Hub and Spoke disjoint	15
5.1.2.	Cloud access	16
5.1.3.	Multicast service	18
5.2.	Site overview	19
5.2.1.	Deciding where to connect the site	20
5.2.1.1.	Site location	20
5.2.1.2.	Site availability	21
5.2.1.3.	Site diversity	23
5.2.1.4.	Route Distinguisher and VRF allocation	25
5.2.2.	VPN policy	25
5.2.3.	Security	28
5.2.3.1.	Encryption	28
5.2.4.	Management	28
5.2.5.	Attachment	28
5.2.5.1.	Bearer	29
5.2.5.2.	Connection	29
5.2.6.	Service	35

5.2.6.1.	QoS	35
5.2.6.2.	Multicast	41
5.2.6.3.	Traffic protection	41
5.2.7.	Customer specific information	42
5.2.7.1.	Customer cascaded LANs with provider managed CE .	43
5.2.7.2.	Customer LANs with provider managed CE	43
5.2.7.3.	Customer LANs with customer managed CE	43
5.3.	Carrier Supporting Carrier	44
5.4.	Using configuration templates	45
6.	Service model usage example	49
7.	Interaction with Other YANG Modules	53
8.	YANG Module	57
9.	To do list / Open questions / Open issues	92
10.	Security Considerations	92
11.	Contributors	92
12.	Acknowledgements	92
13.	IANA Considerations	93
14.	Normative References	93
Appendix A.	Example: NETCONF <get> Reply	94
	Authors' Addresses	94

[1.](#) Introduction

This document defines a YANG data model for Layer 3 IPVPN service configuration.

[1.1.](#) Terminology

The following terms are defined in [[RFC6241](#)] and are not redefined here:

- o client
- o configuration data
- o server
- o state data

The following terms are defined in [[RFC6020](#)] and are not redefined here:

- o augment
- o data model
- o data node

The terminology for describing YANG data models is found in [\[RFC6020\]](#).

1.2. Tree diagram

A simplified graphical representation of the data model is presented in [Section 5](#).

The meaning of the symbols in these diagrams is as follows:

- o Brackets "[" and "]" enclose list keys.
- o Curly braces "{" and "}" contain names of optional features that make the corresponding node conditional.
- o Abbreviations before data node names: "rw" means configuration (read-write), and "ro" state data (read-only).
- o Symbols after data node names: "?" means an optional node and "*" denotes a "list" or "leaf-list".
- o Parentheses enclose choice and case nodes, and case nodes are also marked with a colon (":").
- o Ellipsis ("...") stands for contents of subtrees that are not shown.

2. Definitions

Customer Edge (CE) Device: The equipment on the customer side of the SP-customer boundary (the customer interface).

Provider Edge (PE) Device: The equipment on the SP side of the SP-customer boundary (the customer interface).

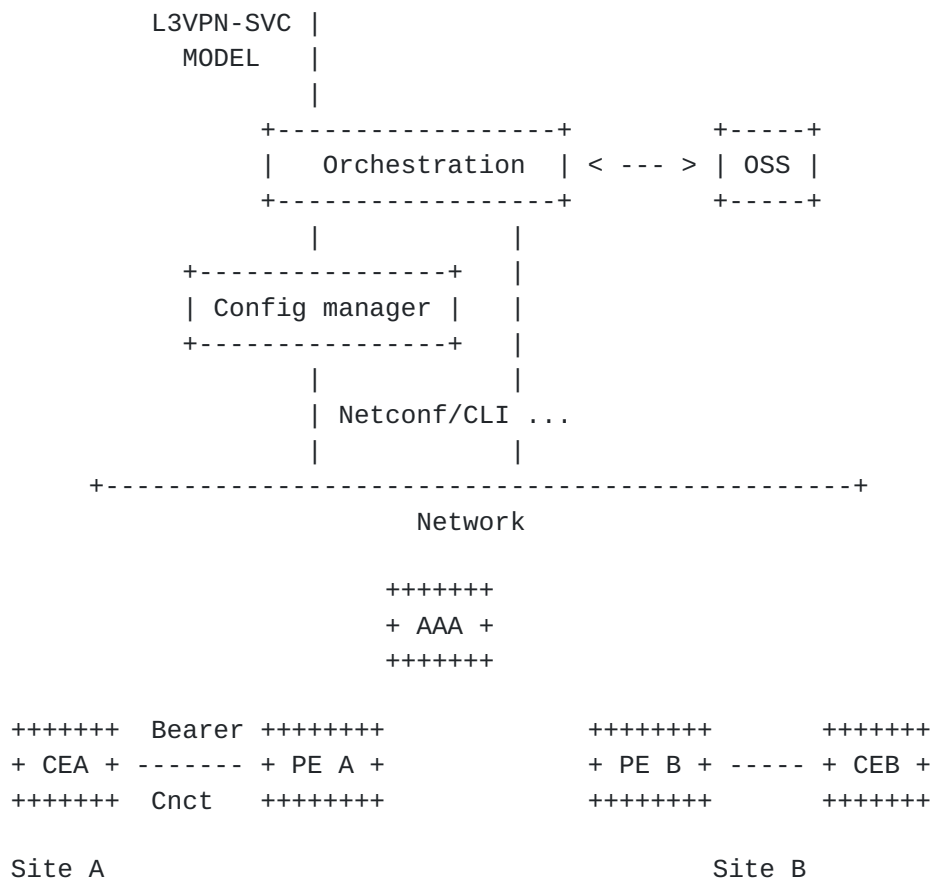
PE-Based VPNs: The PE devices know that certain traffic is VPN traffic. They forward the traffic (through tunnels) based on the destination IP address of the packet, and optionally on based on other information in the IP header of the packet. The PE devices are themselves the tunnel endpoints. The tunnels may make use of various encapsulations to send traffic over the SP network (such as, but not restricted to, GRE, IP-in-IP, IPsec, or MPLS tunnels).

3. Layer 3 IP VPN service model

A Layer 3 IPVPN service is a collection of sites that are authorized to exchange traffic between each other over a shared IP infrastructure. This layer 3 VPN service model aims at providing a

common understanding on how the corresponding IP VPN service is to be deployed over the shared infrastructure. This service model is limited to BGP PE-Based VPNs as described in [RFC4110] and [RFC4364].

4. Service data model usage



The idea of the L3 IPVPN service model is to propose an abstracted interface to manage configuration of components of a L3VPN service. A typical usage is to use this model as an input for an orchestration layer who will be responsible to translate it to orchestrated configuration of network elements who will be part of the service. The network elements can be routers, but also servers (like AAA), and not limited to these examples. The configuration of network elements MAY be done by CLI, or by NetConf/RestConf coupled with specific configuration YANG data models (BGP, VRF, BFD ...) or any other way.

The usage of this service model is not limited to this example, it can be used by any component of the management system but not directly by network elements.

5. Design of the Data Model

The YANG module is divided in two main containers : vpn-svc, sites.

The vpn-svc defines global parameters for the VPN service for a specific customer.

A site is composed of a customer edge router (CE) attached to a provider edge router (PE). The attachment is done through a bearer with a connection (transport protocol) on top. The bearer refers to physical properties of the attachment while the connection refers to more protocol oriented properties.

Authorization of traffic exchange is done through what we call a VPN policy or VPN topology defining routing exchange rules between sites.

The figure below describe the overall structure of the YANG module:

```

module: ietf-l3vpn-svc
  +--rw l3vpn-svc
    +--rw vpn-svc* [name]
      | +--rw name          string
      | +--rw id?          uint32
      | +--rw customer-name? string
      | +--rw topology?    identityref
      | +--rw cloud-access* [cloud-identifier]
      | | +--rw cloud-identifier    string
      | | +--rw authorized-sites* [site-id]
      | | | +--rw site-id    leafref
      | | | +--rw denied-sites* [site-id]
      | | | +--rw site-id    leafref
      | | +--rw nat-enabled?      boolean
      | | +--rw customer-nat-address? inet:ipv4-address
      | +--rw mpls?              boolean
      | +--rw multicast
      | | +--rw tree-flavor*      identityref
      | | +--rw rp
      | | | +--rw static-rps* [rp-address group]
      | | | | +--rw rp-address    union
      | | | | +--rw group         union
      | | | +--rw rp-discovery?   identityref
      | | +--rw anycast-rp-location* string
    +--rw sites* [site-id]
      | +--rw site-id          string
      | +--rw apply-template?  leafref
      | +--rw requested-site-start? yang:date-and-time
      | +--rw requested-site-stop?  yang:date-and-time
      | +--rw actual-site-start?    yang:date-and-time

```



```
| +--rw actual-site-stop?                yang:date-and-time
| +--rw location
| | +--rw address?          string
| | +--rw zip-code?         string
| | +--rw city?             string
| | +--rw country-code?     string
| +--rw site-diversity
| | +--rw type?              enumeration
| | +--rw site-group*       uint32
| +--rw availability
| | +--rw availability-type
| | | +--rw service-type?    identityref
| | | +--rw access-type?     identityref
| | +--rw loadsharing-type?  identityref
| +--rw management
| | +--rw type?              identityref
| | +--rw management-transport? identityref
| | +--rw address?          union
| +--rw vpn-policy
| | +--rw entries* [id]
| | | +--rw id              uint32
| | | +--rw filter
| | | | +--rw lan-prefixes
| | | | | +--rw ipv4-lan-prefixes* [lan]
| | | | | | +--rw lan      inet:ipv4-prefix
| | | | | | +--rw ipv6-lan-prefixes* [lan]
| | | | | | | +--rw lan      inet:ipv6-prefix
| | | | | +--rw lan-tag*      string
| | | +--rw vpn
| | | | +--rw vpn?            leafref
| | | | +--rw site-role?     identityref
| +--rw maximum-routes
| | +--rw address-family* [af]
| | | +--rw af                identityref
| | | +--rw maximum-routes?   uint32
| +--rw customer-specific-information
| | +--rw name?                string
| | +--rw autonomous-system?   uint32
| | +--rw interface?           string
| | +--rw customer-lan-connection* [address]
| | | +--rw address            union
| | | +--rw lan-protocol?      identityref
| | +--rw cascaded-lan-prefixes
| | | +--rw ipv4-lan-prefixes* [lan]
| | | | +--rw lan              inet:ipv4-prefix
| | | | +--rw lan-tag?         string
| | | | +--rw next-hop?        inet:ipv4-address
| | | +--rw ipv6-lan-prefixes* [lan]
```



```

| |      +--rw lan          inet:ipv6-prefix
| |      +--rw lan-tag?     string
| |      +--rw next-hop?    inet:ipv6-address
| +--rw security
| |   +--rw apply-template?  leafref
| |   +--rw authentication
| |   +--rw encryption
| | |   +--rw enabled?        boolean
| | |   +--rw layer?          enumeration
| | |   +--rw encryption-profile
| | |   |   +--rw (profile)?
| | |   |   |   +--:(provider-profile)
| | |   |   |   |   +--rw profile-name?  string
| | |   |   |   |   +--:(customer-profile)
| | |   |   |   |   +--rw algorithm?      string
| | |   |   |   |   +--rw (key-type)?
| | |   |   |   |   ...
| | |   +--rw access-control-list
| +--rw attachment
| |   +--rw apply-template?  leafref
| |   +--rw bearer
| | |   +--rw type?           string
| | |   +--rw bearer-reference? string
| | +--rw ip-connection
| | |   +--rw ipv4
| | | |   +--rw address-allocation-type?  identityref
| | | |   +--rw subnet-prefix?            inet:ipv4-prefix
| | |   +--rw ipv6
| | | |   +--rw address-allocation-type?  string
| | | |   +--rw subnet-prefix?            inet:ipv6-prefix
| | |   +--rw oam
| | | |   +--rw bfd
| | | | |   +--rw bfd-enabled?            boolean
| | | | |   +--rw (holdtime)?
| | | | |   +--:(profile)
| | | | |   |   ...
| | | | |   +--:(fixed)
| | | |   ...
| | +--rw routing-protocols* [type]
| | |   +--rw type            identityref
| | |   +--rw ospf
| | | |   +--rw address-family*  identityref
| | | |   +--rw area-address?    yang:dotted-quad
| | | |   +--rw metric?          uint16
| | | |   +--rw sham-link* [target-site]
| | | | |   +--rw target-site    leafref
| | | | |   +--rw metric?        uint16
| | +--rw bgp

```



```

| | | +--rw address-family* identityref
| | +--rw static
| | | +--rw address-family* identityref
| | +--rw rip
| | | +--rw address-family* identityref
| | +--rw vrrp
| | +--rw address-family* identityref
| +--rw service
| | +--rw apply-template? leafref
| | +--rw svc-input-bandwidth? uint32
| | +--rw svc-output-bandwidth? uint32
| | +--rw svc-mtu? uint16
| | +--rw qos
| | | +--rw qos-classification-policy
| | | | +--rw rules* [id]
| | | | | +--rw id uint16
| | | | | +--rw match-flow
| | | | | | +--rw ipv4-src-prefix? inet:ipv4-prefix
| | | | | | +--rw ipv6-src-prefix? inet:ipv6-prefix
| | | | | | +--rw ipv4-dst-prefix? inet:ipv4-prefix
| | | | | | +--rw ipv6-dst-prefix? inet:ipv6-prefix
| | | | | | +--rw l4-src-port? uint16
| | | | | | +--rw l4-dst-port? uint16
| | | | | | +--rw protocol-field? union
| | | | | +--rw target-class-id? string
| | | +--rw std-qos-profile? string
| | | +--rw custom-qos-profile
| | | | +--rw class* [class-id]
| | | | | +--rw class-id string
| | | | | +--rw rate-limit? uint8
| | | | | +--rw priority-level? uint8
| | | | | +--rw guaranteed-bw-percent? uint8
| | +--rw traffic-protection
| | | +--rw link-local-protection? boolean
| | | +--rw node-local-protection? boolean
| | | +--rw node-global-protection? boolean
| | +--rw mpls
| | | +--rw signalling-type? enumeration
| | +--rw multicast
| | | +--rw multicast-site-type? enumeration
| | | +--rw multicast-transport-protocol
| | | | +--rw ipv4? boolean
| | | | +--rw ipv6? boolean
| | | +--rw protocol-type? enumeration
+--rw site-templates* [site-template-id]
| | +--rw site-template-id string
| | +--rw requested-site-start? yang:date-and-time
| | +--rw requested-site-stop? yang:date-and-time

```



```

+--rw actual-site-start?          yang:date-and-time
+--rw actual-site-stop?          yang:date-and-time
+--rw location
|  +--rw address?                string
|  +--rw zip-code?               string
|  +--rw city?                   string
|  +--rw country-code?           string
+--rw site-diversity
|  +--rw type?                    enumeration
|  +--rw site-group*             uint32
+--rw availability
|  +--rw availability-type
|  |  +--rw service-type?         identityref
|  |  +--rw access-type?         identityref
|  +--rw loadsharing-type?       identityref
+--rw management
|  +--rw type?                    identityref
|  +--rw management-transport?   identityref
|  +--rw address?                union
+--rw vpn-policy
|  +--rw entries* [id]
|  |  +--rw id                    uint32
|  |  +--rw filter
|  |  |  +--rw lan-prefixes
|  |  |  |  +--rw ipv4-lan-prefixes* [lan]
|  |  |  |  |  +--rw lan          inet:ipv4-prefix
|  |  |  |  |  +--rw ipv6-lan-prefixes* [lan]
|  |  |  |  |  |  +--rw lan        inet:ipv6-prefix
|  |  |  +--rw lan-tag*           string
|  |  +--rw vpn
|  |  |  +--rw vpn?                leafref
|  |  |  +--rw site-role?         identityref
+--rw maximum-routes
|  +--rw address-family* [af]
|  |  +--rw af                    identityref
|  |  +--rw maximum-routes?       uint32
+--rw customer-specific-information
|  +--rw name?                    string
|  +--rw autonomous-system?       uint32
|  +--rw interface?               string
|  +--rw customer-lan-connection* [address]
|  |  +--rw address                union
|  |  +--rw lan-protocol?         identityref
|  +--rw cascaded-lan-prefixes
|  |  +--rw ipv4-lan-prefixes* [lan]
|  |  |  +--rw lan                inet:ipv4-prefix
|  |  |  +--rw lan-tag?           string
|  |  |  +--rw next-hop?          inet:ipv4-address
```



```

|     +--rw ipv6-lan-prefixes* [lan]
|         +--rw lan          inet:ipv6-prefix
|         +--rw lan-tag?     string
|         +--rw next-hop?    inet:ipv6-address
+--rw security
|   +--rw authentication
|   +--rw encryption
|   |   +--rw enabled?        boolean
|   |   +--rw layer?          enumeration
|   |   +--rw encryption-profile
|   |       +--rw (profile)?
|   |           +--:(provider-profile)
|   |               |   +--rw profile-name?    string
|   |               +--:(customer-profile)
|   |                   +--rw algorithm?        string
|   |                   +--rw (key-type)?
|   |                       ...
|   +--rw access-control-list
+--rw attachment
|   +--rw bearer
|   |   +--rw type?            string
|   |   +--rw bearer-reference? string
|   +--rw ip-connection
|       +--rw ipv4
|       |   +--rw address-allocation-type?    identityref
|       |   +--rw subnet-prefix?              inet:ipv4-prefix
|       +--rw ipv6
|       |   +--rw address-allocation-type?    string
|       |   +--rw subnet-prefix?              inet:ipv6-prefix
|       +--rw oam
|       |   +--rw bfd
|       |       +--rw bfd-enabled?            boolean
|       |       +--rw (holdtime)?
|       |           +--:(profile)
|       |               |   ...
|       |               +--:(fixed)
|       |                   ...
|       +--rw routing-protocols* [type]
|           +--rw type            identityref
|           +--rw ospf
|           |   +--rw address-family*    identityref
|           |   +--rw area-address?      yang:dotted-quad
|           |   +--rw metric?            uint16
|           |   +--rw sham-link* [target-site]
|           |       +--rw target-site    leafref
|           |       +--rw metric?        uint16
|           +--rw bgp
|           |   +--rw address-family*    identityref

```



```

|         +--rw static
|         | +--rw address-family*  identityref
|         +--rw rip
|         | +--rw address-family*  identityref
|         +--rw vrrp
|         +--rw address-family*  identityref
+--rw service
  +--rw svc-input-bandwidth?  uint32
  +--rw svc-output-bandwidth? uint32
  +--rw svc-mtu?              uint16
  +--rw qos
  | +--rw qos-classification-policy
  | | +--rw rules* [id]
  | |   +--rw id                uint16
  | |   +--rw match-flow
  | |     +--rw ipv4-src-prefix? inet:ipv4-prefix
  | |     +--rw ipv6-src-prefix? inet:ipv6-prefix
  | |     +--rw ipv4-dst-prefix? inet:ipv4-prefix
  | |     +--rw ipv6-dst-prefix? inet:ipv6-prefix
  | |     +--rw l4-src-port?     uint16
  | |     +--rw l4-dst-port?     uint16
  | |     +--rw protocol-field?  union
  | |   +--rw target-class-id?  string
  | +--rw std-qos-profile?      string
  | +--rw custom-qos-profile
  |   +--rw class* [class-id]
  |     +--rw class-id          string
  |     +--rw rate-limit?       uint8
  |     +--rw priority-level?   uint8
  |     +--rw guaranteed-bw-percent? uint8
+--rw traffic-protection
  | +--rw link-local-protection?  boolean
  | +--rw node-local-protection?  boolean
  | +--rw node-global-protection? boolean
+--rw mpls
  | +--rw signalling-type?  enumeration
+--rw multicast
  +--rw multicast-site-type?      enumeration
  +--rw multicast-transport-protocol
  | +--rw ipv4?  boolean
  | +--rw ipv6?  boolean
  +--rw protocol-type?            enumeration

```

5.1. VPN service overview

The vpn-svc top container contains generic information about the VPN service. The name of the vpn-svc refers to an internal reference for this VPN service, while customer name refers to a more explicit

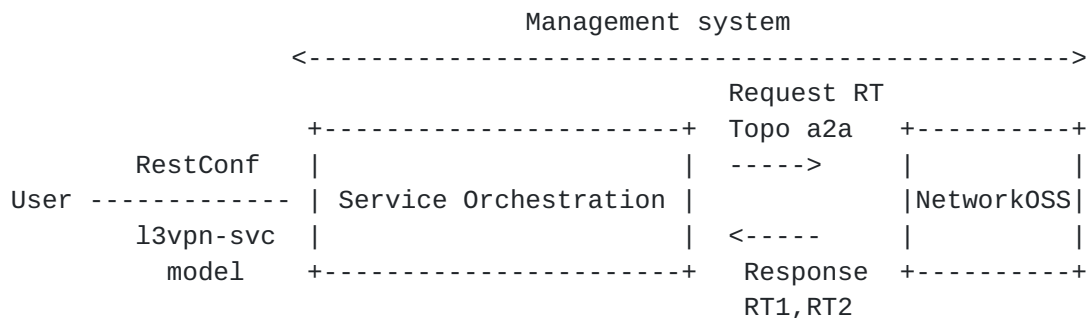
reference to the customer. An identifier (id) is also present for information systems and configuration that requires this information. This identifier is purely internal to the organization responsible for the VPN service. The vpn-svc name and id MUST be unique.

5.1.1.1. VPN service topology

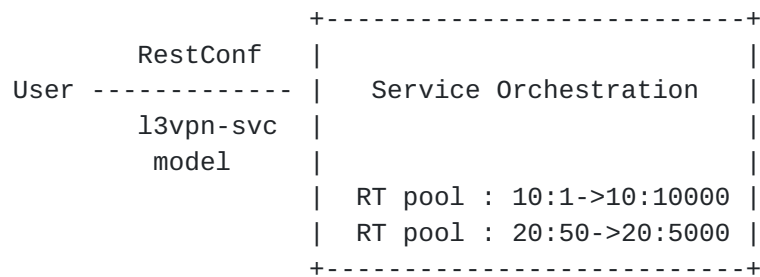
The type of topology of the VPN is required for configuration. Current proposal supports : any-to-any, hub and spoke (where hubs can exchange traffic), and hub and spoke disjoint (where hubs cannot exchange traffic). New topologies could be added by augmentation.

5.1.1.1.1. Route Target allocation

Layer 3 PE-based VPN is built using route-targets as described in [RFC4364]. It is expected management system to allocate automatically set of route-targets upon a VPN service creation request. How management system allocates route-targets is out of scope of the document but multiple ways could be envisaged as described below.



In the example above, a service orchestration, owning the instantiation of this service model, request route-targets to the network OSS. Based on the requested VPN topology, the network OSS replies with one or multiple route-targets. The interface between this service orchestration and network OSS is out of scope of this document.



In the example above, a service orchestration, owning the instantiation of this service model, owns one or more pools of route-target (filled by service provider) that can be allocated. Based on the requested VPN topology, it will allocate one or multiple route-targets from the pool.

The mechanism displayed above are just examples and SHOULD NOT be considered as exhaustive list of solutions.

[5.1.1.2.](#) Any to any



Figure - Any to any VPN topology

In the any to any topology, all VPN sites can discuss between each other without any restriction. It is expected that the management system that owns a any to any IPVPN service request through this model, needs to assign and then configure the VRF and route-targets on the appropriate PEs. In case of any to any, in general a single route-target is required and every VRF imports and exports this route-target.

[5.1.1.3.](#) Hub and Spoke

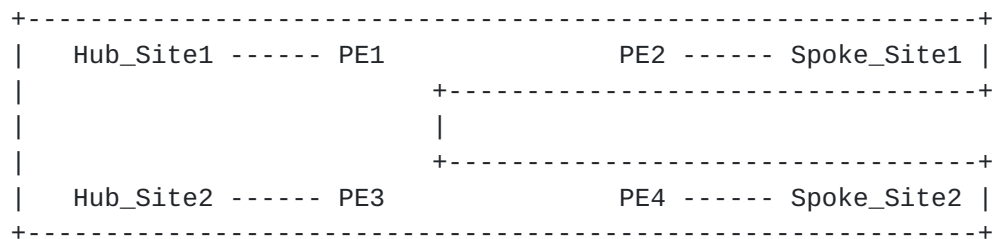
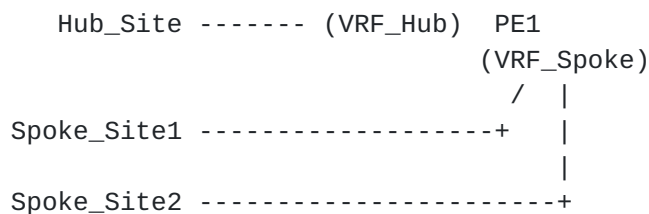


Figure - Hub and Spoke VPN topology

In the hub and spoke topology, all spoke sites can discuss only with Hub sites but not between each other. Hubs can discuss also between each other. It is expected that the management system that owns a any to any IPVPN service request through this model, needs to assign and then configure the VRF and route-targets on the appropriate PEs. In case of hub and spoke, in general a two route-targets are required (one route-target for Hub routes, one route-target for spoke routes). A Hub VRF, connecting Hub sites, will export Hub routes with Hub route-target, and will import Spoke routes through Spoke route-target. It will also import the Hub route-target to allow Hub to Hub communication. A Spoke VRF, connecting Spoke sites, will export Spoke routes with Spoke route-target, and will import Hub routes through Hub route-target.

The management system MUST take into account Hub and Spoke connections constraints. For example, if management system decides to mesh a spoke site and a hub site on the same PE, it needs to mesh connections in different VRFs as displayed in the figure below.



[5.1.1.4.](#) Hub and Spoke disjoint

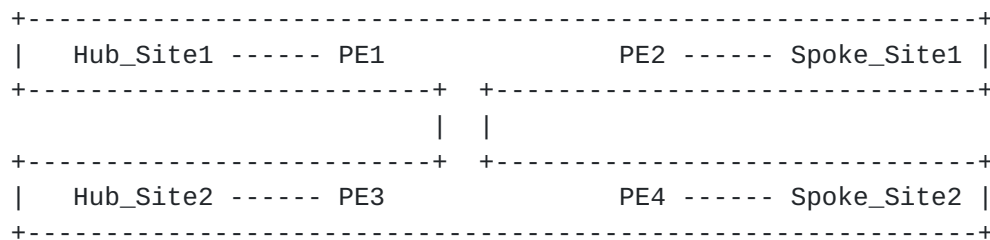


Figure - Hub and Spoke disjoint VPN topology

In the hub and spoke disjoint topology, all spoke sites can discuss only with Hub sites but not between each other. Hubs cannot discuss between each other. It is expected that the management system that owns a any to any IPVPN service request through this model, needs to assign and then configure the VRF and route-targets on the appropriate PEs. In case of hub and spoke, in general a two route-targets are required (one route-target for Hub routes, one route-target for spoke routes). A Hub VRF, connecting Hub sites, will export Hub routes with Hub route-target, and will import Spoke routes

through Spoke route-target. A Spoke VRF, connecting Spoke sites, will export Spoke routes with Spoke route-target, and will import Hub routes through Hub route-target.

The management system MUST take into account Hub and Spoke connections constraints as in the previous case.

5.1.2. Cloud access

The proposed model provides cloud access configuration through the cloud-access container. Internet access can typically be considered as cloud access service. The cloud-access container provides parameters for network address translation and authorization rules.

A cloud identifier is used to reference the target service. This identifier is local to each administration.

If NAT is required to access to the cloud, the nat-enabled leaf MUST be set to true. A NAT address can be provided in customer-nat-address, in case the customer is providing the public IP address for the cloud access. If service provider is providing the NAT address, customer-nat-address MAY not be necessary as it can be picked from a service provider pool.

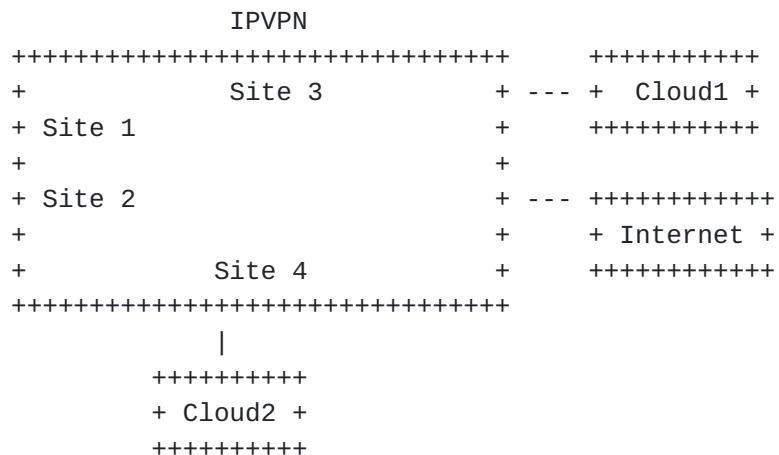
By default, all sites in the IPVPN MUST be authorized to access to the cloud. In case restrictions are required, a user MAY configure the authorized-sites and denied-sites list. The authorization-sites defines the list of sites authorized for cloud access. The denied-sites defines the list of sites denied for cloud access. The model supports both "deny all expect" and "accept all expect" authorization.

The "deny all expect" behavior is obtained by filling only the authorized-sites. All the sites listed will be authorized, all others will be denied.

The "accept all expect" behavior is obtained by filling only the denied-sites. All the sites listed will be denied, all others will be authorized.

Defining both denied-sites and authorized-sites MUST be processed as "deny all expect", so the denied-sites will have not effect.

How the restrictions will be configured on network elements is out of scope of this document and will be specific to each deployment.



In the example above, we may configure the global VPN to access Internet by creating a cloud-access pointing to the cloud identifier for Internet service. No authorized-sites will be configured as all sites are required to access to Internet. NAT-enabled will be set to true and a nat-address will be configured.

```

<vpn-svc>
  <name>ZKITYHJ054687</name>
  <id>12456487</id>
  <customer-name>CUSTOMER_1</customer-name>
  <topology>any-to-any</topology>
  <cloud-access>
    <cloud-identifier>51</cloud-identifier>
    <nat-enabled>true</nat-enabled>
    <nat-address>1.1.1.1</nat-address>
  </cloud-access>
</vpn-svc>

```

If Site1 and Site2 requires access to Cloud1, a new cloud-access will be created pointing to the cloud identifier of Cloud1. Authorized sites will be filled with reference to Site1 and Site2.


```
<vpn-svc>
  <name>ZKITYHJ054687</name>
  <id>12456487</id>
  <customer-name>CUSTOMER_1</customer-name>
  <topology>any-to-any</topology>
  <cloud-access>
    <cloud-identifier>1111111</cloud-identifier>
    <authorized-sites>
      <site-id>site1</site-id>
      <site-id>site2</site-id>
    </authorized-sites>
  </cloud-access>
</vpn-svc>
```

If all sites except Site1 requires access to Cloud2, a new cloud-access will be created pointing to the cloud identifier of Cloud2. denied-sites will be filled with reference to Site1.

```
<vpn-svc>
  <name>ZKITYHJ054687</name>
  <id>12456487</id>
  <customer-name>CUSTOMER_1</customer-name>
  <topology>any-to-any</topology>
  <cloud-access>
    <cloud-identifier>2222222</cloud-identifier>
    <denied-sites>
      <site-id>site1</site-id>
    </denied-sites>
  </cloud-access>
</vpn-svc>
```

5.1.3. Multicast service

Multicast in IP VPN is described in [[RFC6513](#)].

If IPVPN supports multicast service, it is expected to provide inputs on global multicast parameters.

The user of this model will need to fill the flavor of trees that will be used by customer within the IPVPN (Customer tree). The proposed model supports ASM, SSM and BiDirectional trees (and can be augmented). Multiple flavors of tree can be supported simultaneously.


```

                                (SSM tree)
Recv (IGMPv3) -- Site2 ----- PE2
                                PE1 --- Site1 --- Source1
                                \
                                -- Source2

                                (ASM tree)
Recv (IGMPv2) -- Site3 ----- PE3

                                (SSM tree)
Recv (IGMPv3) -- Site4 ----- PE4
                                /
Recv (IGMPv2) -- Site5 -----
                                (ASM tree)

```

In case of ASM flavor, this model requires to fill the rp and rp-discovery parameters. The rp-discovery supports the auto-rp, static-rp, anycast-rp and bsr-rp modes.

Work on multicast service modeling is not finalized yet.

5.2. Site overview

The L3VPN service is really attached to the notion of sites. A site is composed of some characteristics :

- o Unique identifier (site-id) : to uniquely identify the site within the overall network infrastructure. The identifier is a string allowing to any encoding for the local administration of the VPN service.
- o Location (location) : site location informations to allow easy retrieval on nearest available ressources.
- o Site constraints (site-diversity) : site-diversity container allow to define some constraints for the setup of the site, for example : PE disjointness or PoP disjointness. A site-group identifier allow to manage the disjointness. Two sites with the same group and requiring PE disjointness cannot be connected on the same PE.
- o Site availability (site-availability) : allow to define the availability service type (single, loadsharing, primary/backup scenarios) and also the access-type within the service (single, primary, backup ...).

- o Management (management) : defines the model of management of the site, for example : comanaged, customer managed or provider managed.
- o Attachment (attachment) : defines parameters of the attachment of the site, especially bearer, connection and service parameters.

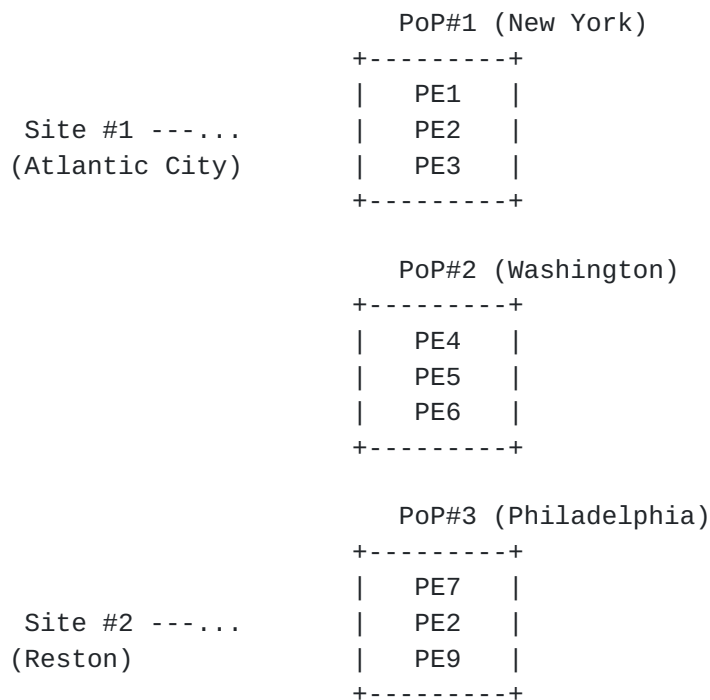
The site configuration is viewed as a global entity, we assume that it is mostly the role of the management to split the parameters between the different elements within the network. For example, in the case of the attachment configuration, the management system needs to split the overall parameters between PE configuration and CE configuration.

5.2.1. Deciding where to connect the site

The management system will have to decide where to connect the site in the provider network (PE, aggregation switch ...). This decision MAY be based on any constraint that are up to the service provider : least load, distance ... The current model proposes some parameters that will help the management system to decide where to attach the customer site. It would be up to the service provider to define which on those parameters are relevant for placing the site, moreover the service provider can decide to rely also on other internal parameters.

5.2.1.1. Site location

The location information provided in this model MAY be used by a management system to decide the target PE to mesh the site.



In the example below, the management system may decide to mesh Site #1 on a PE from Philadelphia PoP for distance reason. It may also take in account resources available on PEs to decide the exact target PE (least load). In case of shortest distance PE used, it may also decide to mesh Site #2 on Washington PoP.

[5.2.1.2.](#) Site availability

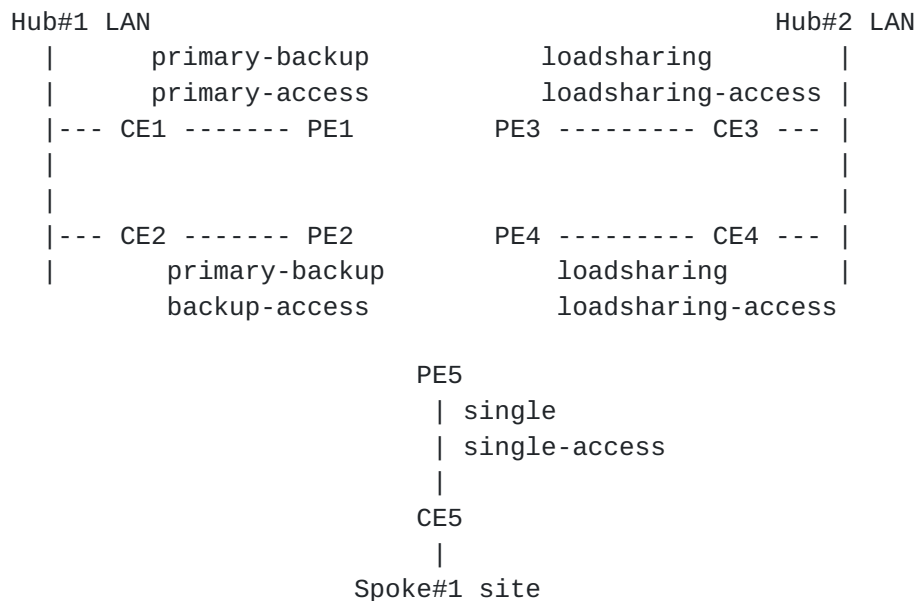
The site availability defines parameters for the site redundancy. The YANG model proposes three models of redundancy (service-type) that MAY be extended by augmentation :

- o single : single homing scenario, no attachment redundancy required.
- o primary-backup : dual homing scenario, one attachment is primary, when the attachment goes down, traffic goes to the backup attachment.
- o loadsharing : multihoming scenario, both attachment are used at the same time. How loadsharing is done is not part of service-type.

The site availability also defines access-type which defines the role of the site in the availability system. The YANG model proposes four models of access-type :

- o single-access : used to indicate the single access function. Used for single-homed sites and coupled with "single" service-type.
- o primary-access : used to indicate the primary access in a primary-backup service-type.
- o backup-access : used to indicate the backup access in a primary-backup service-type.
- o loadsharing-access : used to indicate any access in a loadsharing service-type.

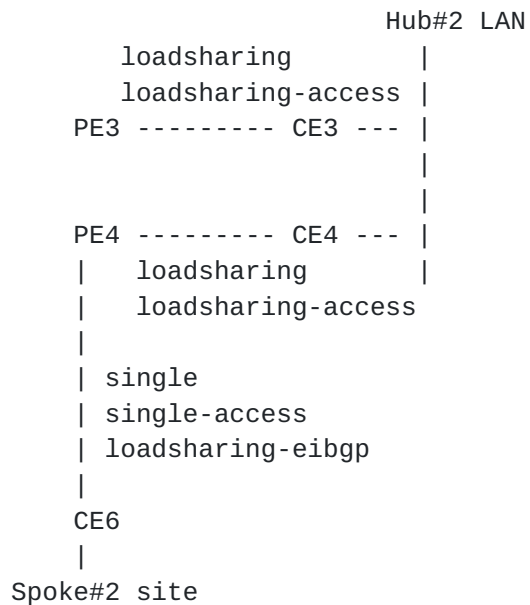
The figure below describes how site availability attributes can be used.



In case of loadsharing, the YANG model proposes to define the loadsharing policy using loadsharing-type. The model proposes the following types that could be augmented. These option MAY be used on any remote sites (single, primary-backup ...) facing a loadsharing site.

- o loadsharing-ibgp : loadbalancing will be done between multiple received iBGP paths.
- o loadsharing-eibgp : loadbalancing will be done between multiple received iBGP paths or eBGP paths : typical application is when a remote site is connected on the same PE than the hub requiring loadsharing.

Considering the diagram above, to enable loadsharing for Hub#2 sites, Hub#1 sites and Spoke#1 site can be configured with loadsharing-ibgp as loadsharing-type. If a Spoke#2 site was connected also on PE4, Spoke#2 site may be configured with loadsharing-eibgp as loadsharing-type.



The site availability will have an impact on site meshing location, as dual homing may require meshing at least on different network elements (see next section).

[5.2.1.3.](#) Site diversity

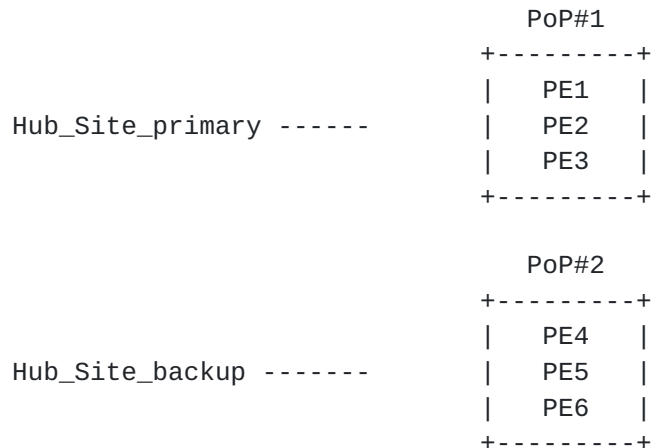
The site diversity defines what is the acceptable fate sharing level in case multiple sites for a single VPN must be provisioned in a common location. The site diversity introduces the notion of site-group. Sites belonging to the same site-group cannot share the same fate. We propose to introduce two constraints :

PoP diverse : site belonging to the same site-group MUST be provisioned on different PoPs.

PE diverse : site belonging to the same site-group MUST be provisioned on different PE routers.

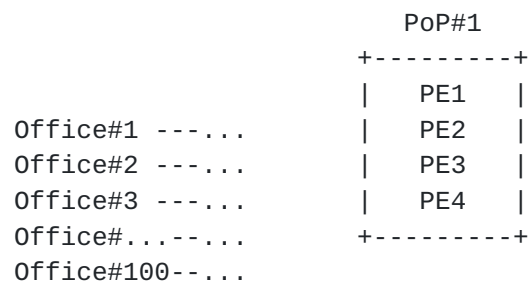
How these diversity constraints are applied is out of scope of the document. As an example, the management system receiving the request for diversity, MAY exchange information with some OSS components to define the best target PEs based on location and ressource availability.

Consider a dual homed hub site, it is desirable for redundancy to provision the two VPN access connections on two different PEs or two different PoPs.



In a PoP diverse scenario, the management system may decide to mesh Hub_Site_primary on any PE of PoP#1 and Hub_Site_backup on any PE of PoP#2. In a PE diverse scenario, if the management system decides to mesh Hub_Site_primary on PE1, it is require to mesh Hub_Site_backup on any PE different from PE1.

Site diversity is not limited to multihoming scenarios. If a company has multiple small branch offices (single homed) that requires to be connected in the same location, it is desirable to dispatch the attachment on multiple PEs. So in case of PE crash, only some offices will be impacted.



In the figure above, it may be good to mesh 25 offices on each PE of PoP#1 to prevent concentration of two many customer offices on common network elements.

5.2.1.4. Route Distinguisher and VRF allocation

Route distinguisher is also a critical parameter of PE-based L3VPN as described in [[RFC4364](#)] that will allow to distinguish common addressing plans in different VPNs. As for Route-targets, it is expected management system to allocate a VRF on the target PE and a route-distinguisher for this VRF.

If a VRF exists on the target PE, and the VRF fulfils the connectivity constraints for the site, there is no need to recreate another VRF and the site MAY be meshed within this existing VRF. How the management system checks that an existing VRF fulfils the connectivity constraints for a site is out of scope of this document.

If no VRF exists on the target PE, filling the site constraints, the management system will have to initiate a new VRF creation on the target PE and will have to allocate a new route distinguisher for this new VRF.

The management system MAY apply a per-VPN or per-VRF allocation policy for the route-distinguisher depending of the service provider policy. In a per-VPN allocation policy, all VRFs (dispatched on multiple PEs) within a VPN will share the same route distinguisher value. In a per-VRF model, all VRFs will always have a unique route-distinguisher value. Some other allocation policies are also possible, and this document does not restrict the allocation policies to be used.

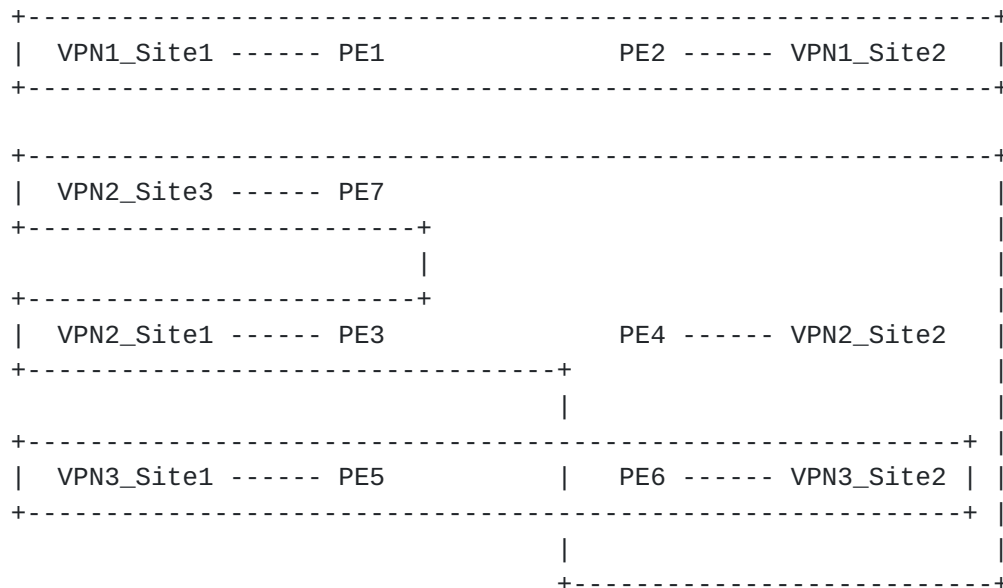
Allocation of route-distinguisher MAY be done in the same way as the route-targets. The example provided in [Section 5.1.1.1](#) could be reused.

Note that a service provider MAY decide to configure target PE for automated allocation of route distinguisher. In this case, there will be no need for any backend system to allocate a route-distinguisher value.

5.2.2. VPN policy

The VPN policy defines the route exchange between multiple VPNs. A vpn-policy configuration MUST be configured to attach a site to one or multiple VPNs. The vpn-policy container defines relations of the site (or specific LAN of the sites) with VPNs. When vpn-policy is defined, the management system will built the route-target policy configuration from a combination of both the vpn-policy and the vpn-topology of the VPNs listed in the policy.

As a site can belong to multiple VPNs, the vpn-policy may be composed of multiple entries. A filter can be applied to specify that only some LANs of the site should be part of a particular VPN. Each time a site (or LAN) is attached to a VPN, we must precise its role within the targeted VPN topology.



Let us consider service VPN1 with a any-to-any topology. For site#1, we define a VPN policy that attach the site to the VPN1 and setting the role of the site as "any-to-any-role".

```

<vpn-policy>
  <entries>
    <id>1</id>
    <vpn>
      <vpn>VPN1</vpn>
      <site-role>any-to-any-role</site-role>
    </vpn>
  </entries>
</vpn-policy>

```

Now let us consider service VPN2 with a hub and spoke disjoint topology (Site1, Site3 are Hubs). For site#1, we define a VPN policy that attach the site to the VPN2 and setting the role of the site as "hub-role".


```
<vpn-policy>
  <entries>
    <id>1</id>
    <vpn>
      <vpn>VPN2</vpn>
      <site-role>hub-role</site-role>
    </vpn>
  </entries>
</vpn-policy>
```

When the vpn services are provisioned a route-target value will be affected by the OSS of the service provider for these VPNs. Let's call RT1 the route-target of VPN1, and RT21(spoke)/RT22(hub) the route-target of VPN2. Now we consider a new VPN service VPN3 (any to any) that must be provisioned, RT3 will be allocated by the OSS for proper configuration on network elements.

Consider a site#1 in VPN3 that must communicate only in VPN3, in this case, the vpn-policy will be similar to Site#1 of VPN1. The VRF on PE5 for VPN3 will be so provisioned by the management system using RT3 value as import and export value.

Consider a site#2 in VPN3 that must communicate in VPN3, and a specific LAN LAN1 must communicate with VPN2 as a spoke (as VPN2 is hub and Spoke disjoint). Below is a sample configuration of the VPN policy for site#2 :

```
<vpn-policy>
  <entries>
    <id>1</id>
    <filter>
      <lan-tag>LAN1</lan-tag>
    </filter>
    <vpn>
      <vpn>VPN2</vpn>
      <site-role>spoke-role</site-role>
    </vpn>
  </entries>
  <entries>
    <id>2</id>
    <vpn>
      <vpn>VPN3</vpn>
      <site-role>any-to-any-role</site-role>
    </vpn>
  </entries>
</vpn-policy>
```


The VRF on PE6 for VPN3 will be so provisioned by the management system using RT3 value as import and RT3 as export value plus RT21 for LAN1 prefix.

5.2.3. Security

Security container defines customer specific security parameters for the site. This section will more be detailed in future revision.

5.2.3.1. Encryption

Encryption can be requested on the connection. It may be performed at layer 2 or layer 3 by selecting the appropriate enumeration in "layer" leaf. The encryption profile can be a service provider defined template or customer specific.

5.2.4. Management

The model proposes three types of common management options :

- o comanaged : the CE router is managed by the provider and also by the customer.
- o provider-managed : the CE router is managed only by the provider.
- o customer-managed : the CE router is managed only by the customer.

Based on the management model, different security options MAY be derived.

In case of "provider-managed" or "comanaged", the model proposes some option to define the management transport protocol (IPv4 or IPv6) and the associated management address.

5.2.5. Attachment

The attachment defines how the service is connected on the network and is splitted in three class of parameters :

- o bearer : defines physical parameters of the attachment.
- o connection : defines protocol parameters of the attachment (transport layer and routing protocols).

5.2.5.1. Bearer

TBD.

5.2.5.2. Connection

The connection defines the protocol parameters of the attachment (IPv4 and IPv6) as well as routing.

5.2.5.2.1. IP addressing

IP subnet can be configured for either transport protocols. For a dual stack connection, two subnets will be provided, one for each transport layer.

The address-allocation-type will help in defining how the address allocation MUST be done. The current model proposes three ways of IP address allocation :

- o pe-dhcp : the PE router will provide DHCP service for CE router, this is applicable to both IPv4 and IPv6 addressing.
- o static-address : Addresses will be assigned manually on both sides, this is applicable to both IPv4 and IPv6 addressing.
- o slaac : enables stateless address autoconfiguration ([\[RFC4862\]](#)). This is applicable only for IPv6.

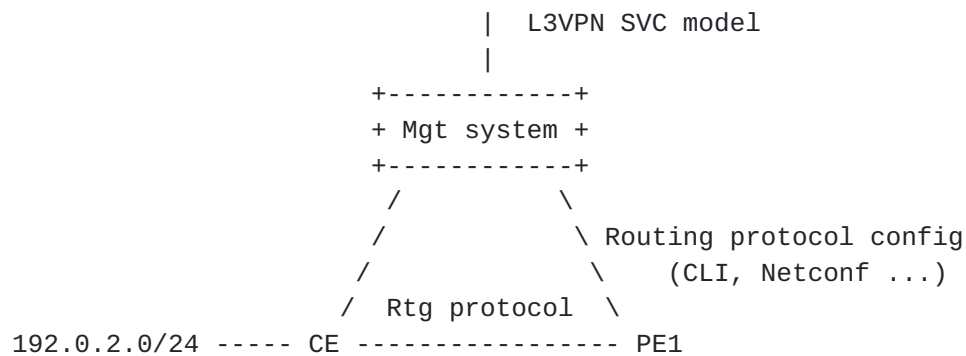
5.2.5.2.2. Routing protocols

Routing-protocol defines which routing protocol must be activated between the provider edge and the customer. The current model support : bgp, rip, rip-ng, ospf, static, direct, vrrp.

```

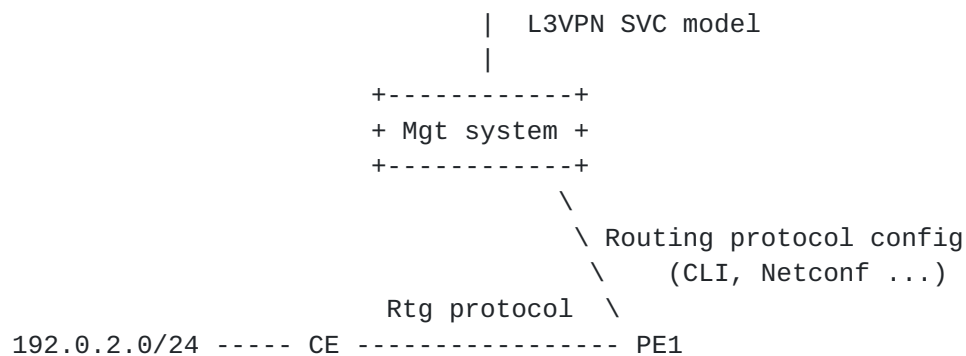
                                Rtg protocol
192.0.2.0/24 ----- CE ----- PE1
```

If the CE is comanaged or provider managed, the management system will be required to configure the routing protocol on both side on the connection.



Routing protocol config - provider managed or comanaged CE

If the CE is customer managed, the management system will be required to configure the routing protocol only the PE router.



Routing protocol config - customer managed CE

[5.2.5.2.2.1](#). Dual stack handling

All routing protocol types support dual stack by using address-family leaf-list.

Example of Dual stack using the same routing protocol :

```

<routing-protocol>
  <type>static</type>
  <static>
    <address-family>ipv4-unicast</address-family>
    <address-family>ipv6-unicast</address-family>
  </static>
</routing-protocol>

```


Example of Dual stack using two different routing protocols :

```
<routing-protocol>
  <type>rip</type>
  <rip>
    <address-family>ipv4-unicast</address-family>
  </rip>
</routing-protocol>
<routing-protocol>
  <type>ospf</type>
  <ospf>
    <address-family>ipv6-unicast</address-family>
  </ospf>
</routing-protocol>
```

5.2.5.2.2.2. Direct LAN connection onto SP network

Routing-protocol "direct" SHOULD be used when a customer LAN is directly connected to the provider network and must be advertised in the IPVPN.

LAN attached directly to provider network :

192.0.2.0/24 ----- PE1

5.2.5.2.2.3. Direct LAN connection onto SP network with redundancy

Routing-protocol "vrrp" SHOULD be used when a customer LAN is directly connected to the provider network and must be advertised in the IPVPN and LAN redundancy is expected.

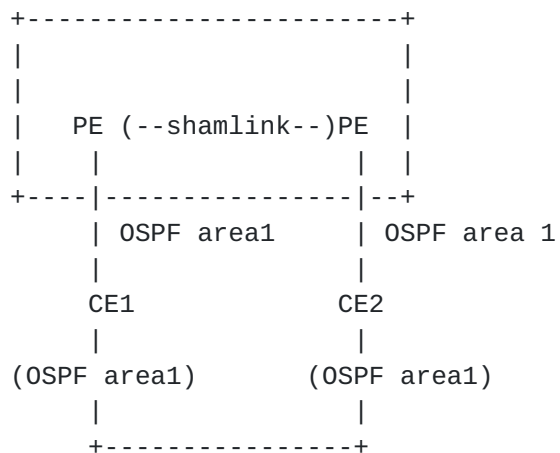
LAN attached directly to provider network with LAN redundancy:

192.0.2.0/24 ----- PE1
 |
 +-- PE2

5.2.5.2.2.4. Static routing

Routing-protocol "static" MAY be used when a customer LAN is connected to the provider network through a CE router and must be advertised in the IPVPN.

The model also proposes an option to create an OSPF sham-link between two sites sharing the same area and having a backdoor link. The sham-link is created by referencing the target site sharing the same OSPF area. The management system will be responsible to check if there is already a shamlink configured for this VPN and area between the same pair of PEs. If there is no existing shamlink, the management system will provision it, this shamlink MAY be reused by other sites.



Regarding Dual stack support, user MAY decide to fill both IPv4 and IPv6 address families, if both protocols SHOULD be routed through OSPF. As OSPF is using two different protocol for IPv4 and IPv6, the management system will need to configure both ospf version 2 and version 3 on the PE-CE link.

Example of OSPF routing parameters in service model.

```

<routing-protocol>
  <type>ospf</type>
  <ospf>
    <area-address>1.1.1.1</area-address>
    <address-family>ipv4-unicast</address-family>
    <address-family>ipv6-unicast</address-family>
  </ospf>
</routing-protocol>

```

Example of PE configuration done by management system :

```

router ospf 10
  area 0.0.0.0
  interface Ethernet0/0
!
router ospfv3 10
  area 0.0.0.0
  interface Ethernet0/0
!

```

[5.2.5.2.2.7.](#) BGP routing

Routing-protocol "bgp" MAY be used when a customer LAN is connected to the provider network through a CE router and must be advertised in the IPVPN.


```
                BGP rtg
10.0.0.0/24 ----- CE ----- PE
```

The AS numbers, peerings addressing will be derived from connection parameters or customer-specific-information as well as internal knowledge of SP.

In case of dual stack access, user MAY request BGP routing for both IPv4 and IPv6 by filling both address-families. It will be up to SP and management system to decide how to decline the configuration (two BGP sessions, single, multisession ...).

The service configuration below activates BGP on PE-CE link for both IPv4 and IPv6.

```
<routing-protocol>
  <type>bgp</type>
  <bgp>
    <address-family>ipv4-unicast</address-family>
    <address-family>ipv6-unicast</address-family>
  </bgp>
</routing-protocol>
```

This service configuration can be derived by management system into multiple flavors depending on SP flavor.

Example #1 of PE configuration done by management system (single session IPv4 transport):

```
router bgp 100
  neighbor 203.0.113.2 remote-as 65000
  address-family ipv4 vrf Cust1
    neighbor 203.0.113.2 activate
  address-family ipv6 vrf Cust1
    neighbor 203.0.113.2 activate
    neighbor 203.0.113.2 route-map SET-NH-IPV6 out
```

Example #2 of PE configuration done by management system (two sessions):

```
router bgp 100
  neighbor 203.0.113.2 remote-as 65000
  neighbor 2001::2 remote-as 65000
  address-family ipv4 vrf Cust1
    neighbor 203.0.113.2 activate
  address-family ipv6 vrf Cust1
    neighbor 2001::2 activate
```


Example #3 of PE configuration done
by management system (multisession):

```
router bgp 100
  neighbor 203.0.113.2 remote-as 65000
  neighbor 203.0.113.2 multisession per-af
  address-family ipv4 vrf Cust1
    neighbor 203.0.113.2 activate
  address-family ipv6 vrf Cust1
    neighbor 203.0.113.2 activate
    neighbor 203.0.113.2 route-map SET-NH-IPV6 out
```

5.2.6. Service

The service defines service parameters associated with the site.

5.2.6.1. QoS

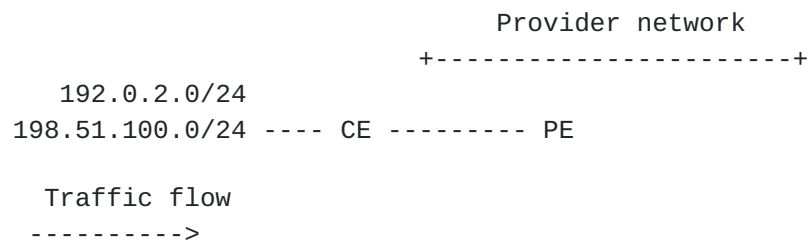
The model proposes to define QoS parameters in an abstracted way :

- o qos-classification-policy : define a set of ordered rules to classify customer traffic.
- o std-qos-profile : provider standard outgoing QoS profile to be applied. This is a reference to a well known profile in Service provider administration.
- o custom-qos-profile : defines customer specific profiles.

5.2.6.1.1. QoS classification

QoS classification rules are handled by qos-classification-policy. The qos-classification-policy is an ordered list of rules that match a flow and set the appropriate target class of service (target-class-id). The match criterions provide a basic infrastructure for defining flows : layer 3 source and destination address, layer 4 ports, layer 4 protocol.

Where the classification is done depends on the SP implementation of the service, but classification concerns the flow coming from the customer site and entering the network.



In the figure above, the management system can decide :

- o if the CE is customer managed, to implement the classification rule in the ingress direction on the PE interface.
- o if the CE is provider managed, to implement the classification rule in the ingress direction on the CE interface connected to customer LAN.

The figure below describes a sample service description of qos-classification for a site :


```
<service>
  <qos>
    <qos-classification-policy>
      <rules>
        <id>1</id>
        <match-flow>
          <ipv4-src-prefix>192.0.2.0/24</ipv4-src-prefix>
          <ipv4-dst-prefix>1.1.1.1/32</ipv4-dst-prefix>
          <l4-dst-port>80</l4-dst-port>
          <l4-protocol>tcp</l4-protocol>
        </match-flow>
        <target-class-id>DATA2</target-class-id>
      </rules>
      <rules>
        <id>2</id>
        <match-flow>
          <ipv4-src-prefix>192.0.2.0/24</ipv4-src-prefix>
          <ipv4-dst-prefix>1.1.1.1/32</ipv4-dst-prefix>
          <l4-dst-port>21</l4-dst-port>
          <l4-protocol>tcp</l4-protocol>
        </match-flow>
        <target-class-id>DATA2</target-class-id>
      </rules>
      <rules>
        <id>3</id>
        <target-class-id>DATA1</target-class-id>
      </rules>
    </qos-classification-policy>
  </qos>
</service>
```

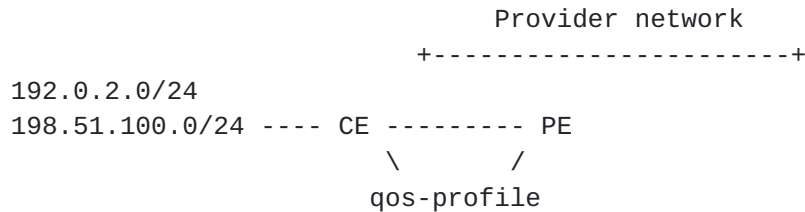
In the example above :

- o HTTP traffic from 192.0.2.0/24 LAN destined to 1.1.1.1/32 will be classified in DATA2.
- o FTP traffic from 192.0.2.0/24 LAN destined to 1.1.1.1/32 will be classified in DATA2.
- o All other traffic will be classified in DATA1.

The order of rules is really important. The management system responsible for translating those rules in network element configuration MUST keep the same processing order in element configuration. The order of rule is defined by the "id" leaf. The lowest "id" MUST be processed first.

5.2.6.1.2. QoS profile

std-qos-profile and custom-qos-profile define output QoS profiles to be used between PE and CE.



std-qos-profile and custom-qos-profile cannot be used at the same time. If both are present, the management system SHOULD keep only the custom-qos-profile.

A custom-qos-profile is defined as a list of class of services and associated properties. The properties are :

- o rate-limit : used to rate-limit the class of service. The value is expressed as a percentage of the global service bandwidth.
- o priority-level : used to define priorities between class of services. The value of the priority to be used is dependant of each administration. The higher the priority-level is, the higher the priority of the class will be. Priority-level is used to define strict priority queueing. A priority-level 250 class will be served before a priority-level 100 class until there is no more packet to process or until rate-limit does not allow anymore packets from the higher priority class.
- o guaranteed-bw-percent : used to define a guaranteed amount of bandwidth for the class of service. It is expressed as a percentage. The guaranteed-bw-percent uses available bandwidth at the priority-level of the class.

Example of service configuration using a standard qos profile :


```
<sites>
  <site-id>1245HRTFGJGJ154654</site-id>
  <service>
    <svc-input-bandwidth>100000000</svc-input-bandwidth>
    <svc-output-bandwidth>100000000</svc-output-bandwidth>
    <qos>
      <std-qos-profile>PLATINUM</std-qos-profile>
    </qos>
  </service>
</sites>
<sites>
  <site-id>555555AAAA2344</site-id>
  <service>
    <svc-input-bandwidth>2000000</svc-input-bandwidth>
    <svc-output-bandwidth>2000000</svc-output-bandwidth>
    <qos>
      <std-qos-profile>GOLD</std-qos-profile>
    </qos>
  </service>
</sites>
```

Example of service configuration using a custom qos profile :


```
<sites>
  <site-id>Site1</site-id>
  <service>
    <svc-input-bandwidth>100000000</svc-input-bandwidth>
    <svc-output-bandwidth>100000000</svc-output-bandwidth>
    <qos>
      <custom-qos-profile>
        <class>
          <class-id>REAL_TIME</class-id>
          <rate-limit>10</rate-limit>
          <priority-level>10</priority-level>
        </class>
        <class>
          <class-id>DATA</class-id>
          <priority-level>5</priority-level>
        </class>
      </custom-qos-profile>
    </qos>
  </service>
</sites>
<sites>
  <site-id>Site2</site-id>
  <service>
    <svc-input-bandwidth>20000000</svc-input-bandwidth>
    <svc-output-bandwidth>20000000</svc-output-bandwidth>
    <qos>
      <custom-qos-profile>
        <class>
          <class-id>REAL_TIME</class-id>
          <rate-limit>30</rate-limit>
          <priority-level>10</priority-level>
        </class>
        <class>
          <class-id>DATA1</class-id>
          <priority-level>5</priority-level>
          <guaranteed-bw-percent>80</guaranteed-bw-percent>
        </class>
        <class>
          <class-id>DATA2</class-id>
          <priority-level>5</priority-level>
          <guaranteed-bw-percent>20</guaranteed-bw-percent>
        </class>
      </custom-qos-profile>
    </qos>
  </service>
</sites>
```


The custom-qos-profile for site1 defines that traffic from REAL_TIME class will have a higher priority than traffic from DATA class. The REAL_TIME traffic will be rate-limit to 10% of the service bandwidth (10% of 100Mbps = 10Mbps) to let some place for DATA traffic.

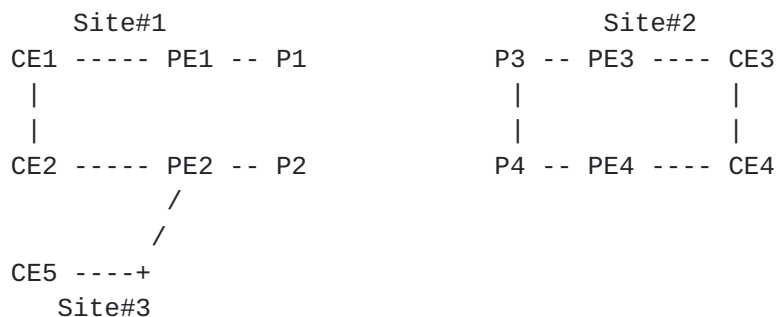
The custom-qos-profile for site2 defines that traffic from REAL_TIME class will have a higher priority than traffic from data traffic. Data traffic will be splitted in two class of service DATA1 and DATA2 that will share bandwidth between them according to the percentage of guaranteed-bw-percent. The maximum of percentage to be used is not limited by this model but MUST be limited by the management system according to the policies authorized by the service provider. The REAL_TIME traffic will be rate-limit to 30% of the service bandwidth (30% of 100Mbps = 30Mbps) to let some place for data traffic. In case of congestion of the access, the REAL_TIME traffic can go up to 30Mbps (Let's assume that 20Mbps only are consumed). The DATA1 and DATA2 will share remaining bandwidth (80Mbps) according to their percentage. So DATA1 will be served with at least 64Mbps of bandwidth.

[5.2.6.2. Multicast](#)

The multicast section defines the type of site in the customer multicast topology : source, receiver, or both. These parameters will help management system to optimize the multicast service.

[5.2.6.3. Traffic protection](#)

The service model supports the ability to protect traffic from or to a site.



In the figure above, we consider an IPVPN service with three sites including two dual homed sites (site#1 and #2). For dual homed sites, we consider PE1-CE1 and PE3-CE3 as primary, and

PE2-CE2,PE4-CE4 as backup for the example (even if protection also applies to loadsharing scenarios.)

In order to protect Site#2 for a PE-CE link failure, a service provider MAY configure link-local-protection on the primary access of site#2 (PE3-CE3). In such case, if we consider traffic coming from a remote site (site#1 or site#3), primary path is to use PE3 as egress PE. PE3 has preprogrammed a backup forwarding entry pointing to backup path (through PE4-CE4) for all prefixes going through PE3-CE3 link. How backup path is computed is out of scope of the document. When PE3-CE3 link fails, traffic is still received by PE3 but PE3 switch automatically traffic to the backup entry, path will so be PE1-P1-(...)-P3-PE3-PE4-CE4 until remote PEs reconverge and use PE4 as egress PE.

In order to protect Site#2 for a egress PE node failure, a service provider may configure node-local-protection on the primary access of site#2 (PE3-CE3). This protection flavor is so called egress PE node protection. In such case, if we consider traffic coming from a remote site (site#1 or site#3), primary path is to use PE3 as egress PE. P3 has preprogrammed a backup forwarding entry pointing to backup path. How backup path is computed is out of scope of the document. When PE3 node fails, traffic is still received by P3 but P3 switch automatically traffic to the backup entry, path will so be PE1-P1-(...)-P3-P4-PE4-CE4 until remote PEs reconverge and use PE4 as egress PE.

In order to protect Site#2 for a egress PE node failure, another flavor of protection can be used, a service provider MAY configure node-global-protection on all remote sites. This protection flavor is generally called PIC Edge (Prefix Independent Convergence for Edge). In such case, if we consider traffic coming from a remote site (site#1 or site#3), primary path is to use PE3 as egress PE. All remote PEs (PE1,PE2) have preprogrammed a backup forwarding entry pointing to backup path (through PE4). How backup path is computed is out of scope of the document. When PE3 node fails, remote PEs detect through IGP convergence that PE3 is no more reachable and decide to switch automatically traffic to the backup entry, path will so be PE1-P1-(...)-P4-PE4-CE4. In this case, as BGP convergence is involved, there is no other need of reconvergence.

5.2.7. Customer specific information

Customer specific informations are used to store informations owned by customers. The customer may have some LANs directly connected to CE routers managed by Service providers. In this case, customer needs to allocate IP addresses from its LAN addressing plan for the CE router. This is the purpose of the customer-lan-connection list

which provides information about addresses allocated from each customer LAN connected to the CE router.

In case, provider needs to route some LAN that are cascaded behind those connected LANs, cascaded-lan-prefixes can be used. In this version of the document, static routing is assumed between CE and customer LANs.

5.2.7.1. Customer cascaded LANs with provider managed CE

```
192.0.2.0/24 ---- CustFW1 ---- 203.0.113.0/24 --- CE1 ----- PE1
                               /
198.51.100.0/24 ----
```

In the example above, CE1 owned by service provider is attached to 203.0.113.0/24 LAN of the customer. Customer allocated 192.0.2.254 for CE1, this information will have to be reported as a customer-lan-connection in the service model. Customer also have two LANs cascaded behind a firewall (192.0.2.0/24 and 198.51.100.0/24), those LANs MUST be routed within the IPVPN. cascaded-lan-prefixes MUST be used to fill those informations. The IP address of the firewall MUST be filled as nexthop of those LANs.

5.2.7.2. Customer LANs with provider managed CE

```
192.0.2.0/24 ---- CE1 ----- PE1
                               /
198.51.100.0/24 ----
```

In the example above, CE1 owned by service provider is attached to 203.0.113.0/24 LAN of the customer. Customer allocated 192.0.2.254 and 198.51.100.254 for CE1, this information will have to be reported as a customer-lan-connection in the service model.

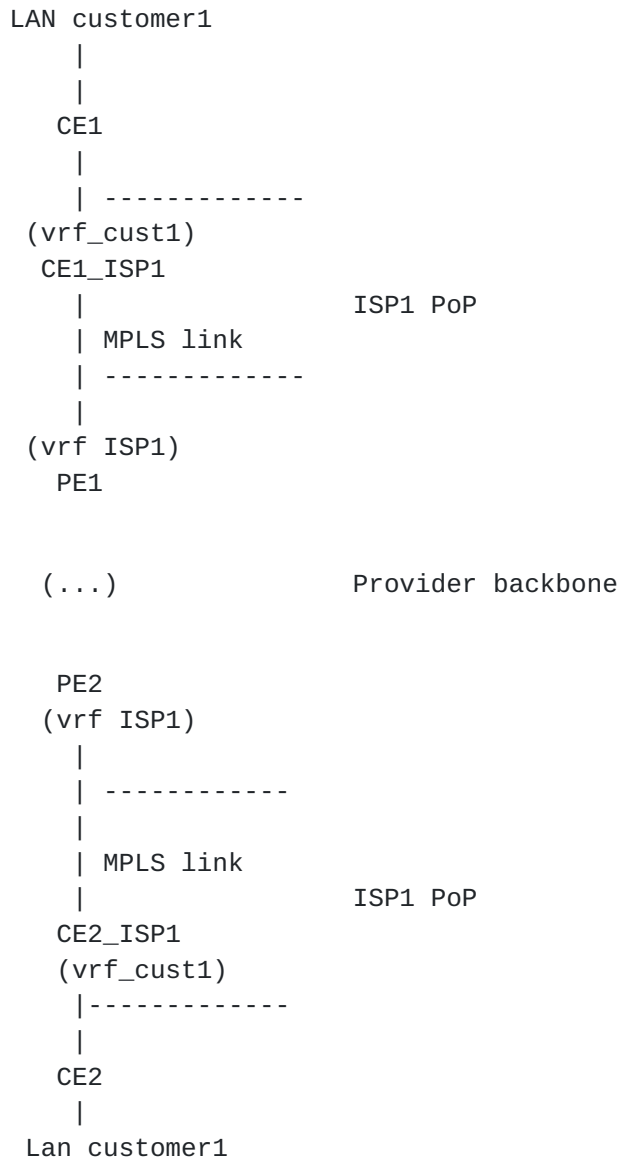
5.2.7.3. Customer LANs with customer managed CE

```
192.0.2.0/24 ---- CE1 ----- PE1
                               /
198.51.100.0/24 ----
```

In the example above, CE1 owned by customer is attached to 203.0.113.0/24 LAN. The service provider only have to deal with routing with customer CE, the routing-protocol information on the connection will define this.

5.3. Carrier Supporting Carrier

In case of Carrier Supporting Carrier (CsC), a customer MAY want to build MPLS service using an IPVPN as transport layer.



In the figure above, ISP1 resells IPVPN service but has no transport infrastructure between its PoPs. ISP1 uses an IPVPN as transport infrastructure (belonging to another provider) between its PoPs.

In order to support CsC, the VPN service must be declared MPLS support using the "mpls" leaf set to true in vpn-svc. The link

between CE1_ISP1/PE1 and CE2_ISP1/PE2 must also run a MPLS signalling protocol. This configuration is done at the site level.

In the proposed model, LDP or BGP can be used as MPLS signalling protocol. In case of LDP, an IGP routing protocol MUST also be activated. In case of BGP signalling, BGP MUST also be configured as routing-protocol.

5.4. Using configuration templates

The proposed model supports the creation and application of configuration templates for sites.

A template can be configured by creating adding a site in the site-template list. The "site-templates" list contains only templates. Real sites are part of the "sites" list.

Multiple templates can be configured. Templates can be applied at multiple levels referenced by apply-template leaf. The apply-template references the site-id of the template to be called. The location of the apply-template within the sites hierarchy defines which parameters must be inherited. For example, if apply-template is done on service container of a site, only service container parameters (and childs) from the template will be applied.


```
<site-templates>
  <site-id>Template-VoiceCoS-Cust1</site-id>
  <service>
    <qos>
      <custom-qos-profile>
        <class>
          <class-id>REAL_TIME</class-id>
          <rate-limit>30</rate-limit>
          <priority-level>10</priority-level>
        </class>
        <class>
          <class-id>DATA1</class-id>
          <priority-level>5</priority-level>
          <guaranteed-bw-percent>80</guaranteed-bw-percent>
        </class>
        <class>
          <class-id>DATA2</class-id>
          <priority-level>5</priority-level>
          <guaranteed-bw-percent>20</guaranteed-bw-percent>
        </class>
      </custom-qos-profile>
    </qos>
  </service>
</sites>
```

```
<site-templates>
  <site-id>Template-VPNsite-Customer1</site-id>
  <service>
    <qos>
      <std-qos-profile>PLATINUM</std-qos-profile>
    </qos>
  </service>
  <attachment>
    <connection>
      <routing-protocol>
        <name>static</name>
        <type>static</type>
      </routing-protocol>
    </connection>
  </attachment>
</sites>
```

Site-templates allow to define configuration blocks that will be inherited by one or multiple sites in order to speed up configuration. For example, if all the sites of an IPVPN service have the almost same configuration (routing-protocol, qos, management ...), a template can be created and each site of the VPN will

reference the template. If a site has some particular parameters, specific parameters within the site MUST always override parameters derived from template.

The example above defines two site templates :

- o Template-VPNsite-Customer1 that will be used to configure all the VPN sites for customer 1.
- o Template-VoiceCoS-Cust1 that will be used to configure some special CoS policy on some specific accesses of the VPN.

In the example below, all sites of VPN1 are inheriting basic configuration from template Template-VPNsite-Customer1. Some specific parameters like svc-input-bandwidth are also defined for each site. For Site 4 and 5 , specific QoS parameters are required, a new template Template-VoiceCoS-Cust1 is applied at service level for these two sites, overriding the service parameters from the Template-VPNsite-Customer1 template.


```
<sites>
  <site-id>Site1</site-id>
  <apply-template>Template-VPNsite-Customer1</apply-template>
  <service>
    <svc-input-bandwidth>5000000</svc-input-bandwidth>
    <svc-output-bandwidth>5000000</svc-output-bandwidth>
  </service>
</sites>
<sites>
  <site-id>Site2</site-id>
  <apply-template>Template-VPNsite-Customer1</apply-template>
  <service>
    <svc-input-bandwidth>20000000</svc-input-bandwidth>
    <svc-output-bandwidth>20000000</svc-output-bandwidth>
  </service>
</sites>
<sites>
  <site-id>Site3</site-id>
  <apply-template>Template-VPNsite-Customer1</apply-template>
  <service>
    <svc-input-bandwidth>30000000</svc-input-bandwidth>
    <svc-output-bandwidth>30000000</svc-output-bandwidth>
  </service>
</sites>

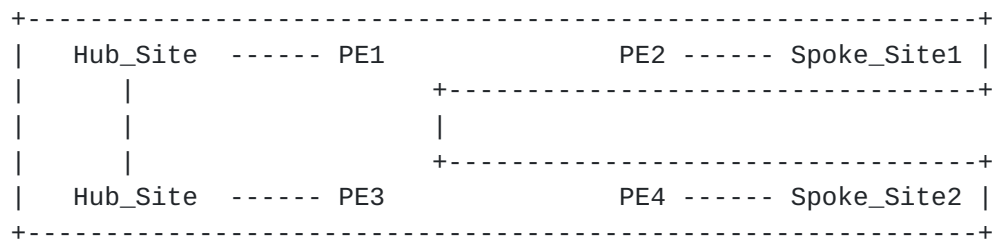
<sites>
  <site-id>Site4</site-id>
  <apply-template>Template-VPNsite-Customer1</apply-template>
  <service>
    <apply-template>Template-VoiceCoS-Cust1</apply-template>
    <svc-input-bandwidth>100000000</svc-input-bandwidth>
    <svc-output-bandwidth>100000000</svc-output-bandwidth>
  </service>
</sites>
<sites>
  <site-id>Site5</site-id>
  <apply-template>Template-VPNsite-Customer1</apply-template>
  <service>
    <apply-template>Template-VoiceCoS-Cust1</apply-template>
    <svc-input-bandwidth>450000000</svc-input-bandwidth>
    <svc-output-bandwidth>450000000</svc-output-bandwidth>
  </service>
</sites>
```


6. Service model usage example

As explained in [Section 4](#), this service model is intended to be instantiated at a management layer and is not intended to be used directly on network elements. The management system serves as a central point of configuration of the overall service.

This section provides an example on how a management system can use this model to configure an IPVPN service on network elements.

The example wants to achieve the provisioning of a VPN service for 3 sites using hub and spoke topology. One of the site will be dual homed and loadsharing is expected.



The following XML describes the overall simplified service configuration of this VPN.

```

<vpn-svc>
  <name>VPN1</name>
  <id>12456487</id>
  <customer-name>CUSTOMER1</customer-name>
  <topology>hub-spoke</topology>
</vpn-svc>

```

When receiving the request for provisioning the VPN service, the management system will internally (or through discussion with other OSS component) allocates VPN route-targets. In this specific case two RTs will be allocated (100:1 for Hub and 100:2 for Spoke). The output below describes the configuration of Spoke1.

```

<sites>
  <site-id>Spoke_Site1</site-id>
  <site-diversity>
    <type>pe-diverse</type>
    <site-group>100</site-group>
  </site-diversity>
  <location>
    <city-code>NY</city-code>
    <country-code>US</country-code>
  </location>

```



```
<attachment>
  <connection>
    <ipv4>
      <subnet-prefix>203.0.113.0/30</subnet-prefix>
    </ipv4>
    <routing-protocol>
      <type>bgp</type>
      <bgp>
        <address-family>ipv4-unicast</address-family>
      </bgp>
    </routing-protocol>
  </connection>
</attachment>
<management>
  <type>provider-managed</type>
  <management-transport>ipv4-unicast</management-transport>
  <address>10.46.1.1</address>
</management>
<service>
  <svc-input-bandwidth>450000000</svc-input-bandwidth>
  <svc-output-bandwidth>450000000</svc-output-bandwidth>
</service>
<vpn-policy>
  <entries>
    <id>1</id>
    <vpn>
      <vpn>VPN1</vpn>
      <site-role>spoke-role</site-role>
    </vpn>
  </entries>
</vpn-policy>
<customer-specific-information>
  <customer-lan-connection>
    <address>192.0.2.254</address>
    <lan-protocol>ipv4-unicast</lan-protocol>
  </customer-lan-connection>
  <cascaded-lan-prefixes>
    <ipv4-lan-prefixes>
      <lan>198.51.100.0/30</lan>
      <nexthop>192.0.2.253</nexthop>
    </ipv4-lan-prefixes>
    <ipv4-lan-prefixes>
      <lan>198.51.100.4/30</lan>
      <nexthop>192.0.2.253</nexthop>
    </ipv4-lan-prefixes>
    <ipv4-lan-prefixes>
      <lan>198.51.100.8/30</lan>
      <nexthop>192.0.2.253</nexthop>
    </ipv4-lan-prefixes>
```



```
</ipv4-lan-prefixes>
</cascaded-lan-prefixes>
</customer-specific-information>
</sites>
```

When receiving the request for provisioning Spoke1 site, the management system MUST allocate network resources for this site. It MUST first decide the target network elements to provision the access, and especially the PE router (and may be an aggregation switch). As described in [Section 5.2.1](#), the management system MAY use the location information and MUST use the site-diversity constraint to find the appropriate PE. In this case, we consider Spoke1 requires PE diversity with Hub and that management system allocate PEs based on lowest distance. Based on the location information, the management system finds the available PEs in the nearest area of the customer and picks one that fits the site-diversity constraint.

When the PE is chosen, management system needs to allocate interface resources on the node, one interface is so picked from the PE available pool. The management system can start provisioning the PE node by using any mean (Netconf, CLI, ...). The management system will check if a VRF is already present that fits the needs. If not, it will provision the VRF : Route distinguisher will come from internal allocation policy model, route-targets are coming from the vpn-policy configuration of the site (management system allocated some RTs for the VPN). As the site is a spoke site (site-role), the management system knows which RT must be imported and exported. As the site is provider managed, some management route-targets may also be added (100:5000). Standard provider VPN policies MAY also be added in the configuration.

Example of generated PE configuration :

```
ip vrf Customer1
  export-map STD-CUSTOMER-EXPORT          <---- Standard SP configuration
  route-distinguisher 100:3123234324
  route-target import 100:1
  route-target import 100:5000            <---- Standard SP configuration
  route-target export 100:2                for provider managed
!
```

When the VRF has been provisioned, the management system can start configuring the access on the PE using the allocated interface information. IP addressing is derived from the subnet-prefix of the connection. One address will be picked from the subnet for the PE, another will be used for the CE configuration. Routing protocols will also be configured on the PE, bgp will be used as requested in

the service model. Peering addresses will be derived from subnet-prefix. PE AS number is well known and as CE is provider managed, CE AS number can be automatically allocated by the management system. Some provider standard configuration templates may also be added.

Example of generated PE configuration :

```
interface Ethernet1/1/0.10
  encapsulation dot1q 10
  ip vrf forwarding Customer1
  ip address 203.0.113.1 255.255.255.252 <---- Comes from
                                          subnet-prefix
  ip access-group STD-PROTECT-IN      <---- Standard SP config
!
router bgp 100
  address-family ipv4 vrf Customer1
    neighbor 203.0.113.2 remote-as 65000    <---- Comes from
                                          subnet-prefix
                                          and allocated CE ASN
    neighbor 203.0.113.2 route-map STD in  <---- Standard SP config
    neighbor 203.0.113.2 filter-list 10 in <---- Standard SP config
!
ip route vrf Customer1 203.0.113.254 255.255.255.255 203.0.113.2
! Static route for provider administration of CE
!
```

As the CE router is not reachable at this stage, the management system can produce a complete CE configuration that can be uploaded to the node by manual operation before sending the CE to customer premise. The CE configuration will be built as for the PE. Based on the CE type (vendor/model) allocated to the customer and bearer information, the management system knows which interface must be configured on the CE. Interface parameters like IP addressing are derived from subnet-prefix taking into account how management system distributes addresses between PE and CE within the subnet. Routing protocol configuration is done as for the PE, in this example, bgp peering address is retrieved from subnet prefix and internal allocation policy, remote AS number is well known. LAN addresses will be added to the configuration and exported to BGP. This will allow to produce a plug'n'play configuration for the CE.

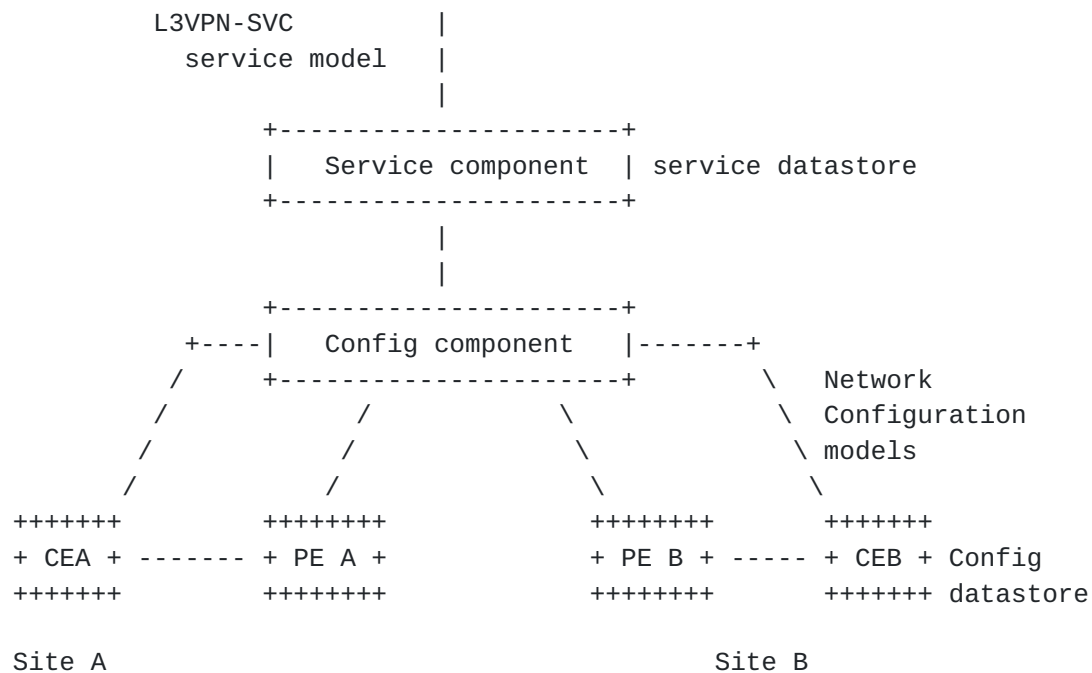
Example of generated CE configuration :

```
interface Loopback10
  description "Administration"
  ip address 203.0.113.254 255.255.255.255
!
interface FastEthernet10
  description "WAN"
  ip address 203.0.113.2 255.255.255.252 <---- Comes from
                                          subnet-prefix
!
interface FastEthernet11
  description "LAN"
  ip address 192.0.2.254 255.255.255.252 <---- Comes from
                                          customer-lan-connection
!
router bgp 65000
  redistribute static route-map STATIC2BGP <---- Standard SP
                                          configuration
  neighbor 203.0.113.1 remote-as 100      <---- Comes from
                                          subnet-prefix
                                          and allocated CE ASN
!
route-map STATIC2BGP permit 10
  match tag 10
!
ip route 198.51.100.0 255.255.255.252 192.0.2.253 tag 10
ip route 198.51.100.4 255.255.255.252 192.0.2.253 tag 10
ip route 198.51.100.8 255.255.255.252 192.0.2.253 tag 10
```

7. Interaction with Other YANG Modules

As expressed in [Section 4](#), this service module is intended to be instantiated in management system and not directly on network elements.

It will be the role of the management system to configure the network elements. The management system MAY be modular, so the component instantiating the service model (let's call it service component) and the component responsible for network element configuration (let's call it configuration component) MAY be different.



In the previous sections, we provided some example of translation of service provisioning request to router configuration lines as illustration. In the NetConf/Yang ecosystem, it will be expected NetConf/YANG to be used between configuration component and network elements to configure the requested service on these elements.

In this framework, it is expected from standardization to also work on specific configuration YANG modelization of service components on network elements. There will be so a strong relation between the abstracted view provided by this service model and the detailed configuration view that will be provided by specific configuration models for network elements.

Authors of this document are expecting definition of YANG models for network elements on this non exhaustive list of items :

- o VRF definition including VPN policy expression.
- o Physical interface.
- o IP layer (IPv4, IPv6).
- o QoS : classification, profiles...
- o Routing protocols : support of configuration of all protocols listed in the document, as well as routing policies associated with these protocols.

- o Multicast VPN.
- o Network Address Translation.
- o ...

Example of VPN site request at service level using this model :

```
<sites>
  <site-id>Site A</site-id>
  <attachment>
    <connection>
      <ipv4>
        <address-allocation-type>static-address</address-allocation-type>
        <subnet-prefix>203.0.113.0/30</subnet-prefix>
      </ipv4>
      <routing-protocol>
        <name>static</name>
        <type>static</type>
      </routing-protocol>
    </connection>
  </attachment>
  <vpn-policy>
    <entries>
      <id>1</id>
      <vpn>
        <vpn>VPN1</vpn>
        <site-role>any-to-any-role</site-role>
      </vpn>
    </entries>
  </vpn-policy>
  <customer-specific-information>
    <cascaded-lan-prefixes>
      <ipv4-lan-prefixes>
        <lan>198.51.100.0/30</lan>
        <nexthop>192.0.2.253</nexthop>
      </ipv4-lan-prefixes>
    </cascaded-lan-prefixes>
  </customer-specific-information>
</sites>
```

In the service example above, it is expected that the service component requests to the configuration component of the management system the configuration of the service elements. If we consider that service component selected a PE (PE A) as target PE for the site, the configuration component will need to push the configuration to PE A. The configuration component will use several YANG data

models to define the configuration to be applied to PE A. The XML configuration of PE-A may look like this :

```
<if:interfaces>
  <if:interface>
    <if:name>eth0</if:name>
    <if:type>ianaift:ethernetCsmacd</if:type>
    <if:description>
      Link to CEA.
    </if:description>
    <ip:ipv4>
      <ip:address>
        <ip:ip>203.0.113.1</ip:ip>
        <ip:prefix-length>30</ip:prefix-length>
      </ip:address>
      <ip:forwarding>true</ip:forwarding>
    </ip:ipv4>
  </if:interface>
</if:interfaces>
<rt:routing>
  <rt:routing-instance>
    <rt:name>VRF_CustA</rt:name>
    <rt:type>l3vpn:vrf</rt:type>
    <rt:description>VRF for CustomerA</rt:description>
    <l3vpn:route-distinguisher>
      100:1546542343
    </l3vpn:route-distinguisher>
    <l3vpn:import-rt>100:1</l3vpn:import-rt>
    <l3vpn:export-rt>100:1</l3vpn:export-rt>
    <rt:interfaces>
      <rt:interface>
        <rt:name>eth0</rt:name>
      </rt:interface>
    </rt:interfaces>
    <rt:routing-protocols>
      <rt:routing-protocol>
        <rt:type>rt:static</rt:type>
        <rt:name>st0</rt:name>
        <rt:static-routes>
          <v4ur:ipv4>
            <v4ur:route>
              <v4ur:destination-prefix>
                198.51.100.0/30
              </v4ur:destination-prefix>
              <v4ur:next-hop>
                <v4ur:next-hop-address>
                  203.0.113.2
                </v4ur:next-hop-address>
              </v4ur:route>
            </v4ur:ipv4>
          </rt:static-routes>
        </rt:routing-protocol>
      </rt:routing-protocols>
    </rt:routing-instance>
  </rt:routing>
</rt:routing>
```



```
        </v4ur:next-hop>
      </v4ur:route>
    </v4ur:ipv4>
  </rt:static-routes>
</rt:routing-protocol>
</rt:routing-protocols>
</rt:routing-instance>
</rt:routing>
```

8. YANG Module

<CODE BEGINS> file "ietf-l3vpn-svc@2015-08-03.yang"

```
module ietf-l3vpn-svc {
  namespace "urn:ietf:params:xml:ns:yang:ietf-l3vpn-svc";

  prefix l3vpn-svc;

  import ietf-routing {
    prefix "rt";
  }

  import ietf-inet-types {
    prefix inet;
  }

  import ietf-yang-types {
    prefix yang;
  }

  organization
    "IETF L3SM Working Group";

  contact
    "WG List:  <mailto:l3sm@ietf.org>;

    Editor:

    ";

  description
    "The YANG module defines a generic service configuration
    model for Layer 3 VPN common across all of the vendor
    implementations.";

  revision 2015-08-03 {
    description
```



```
" * Creating multiple reusable groupings
* Added mpls leaf in vpn-svc for carrier's carrier case
* Modify identity single to single-site
* Modify site-type to site-role and also child identities.
* Creating OAM container under site and moved BFD in.
* Creating flow-definition grouping to be reused
in ACL, QoS ...
* Simplified VPN policy.
* Adding multicast static group to RP mappings.
* Removed native-vpn and site-role from global site
cfg, now managed within the VPN policy.
* Creating a separate list for site templates.
";
reference "";
}
revision 2015-07-02 {
  reference "draft-ietf-l3sm-l3vpn-service-yang-00";
}
revision 2015-04-24 {
  description "
    * Add encryption parameters
    * Adding holdtime for BFD.
    * Add postal address in location
  ";
  reference "draft-lstd-l3sm-l3vpn-service-yang-00";
}
revision 2015-02-05 {
  description "Initial revision.";
  reference "draft-l3vpn-service-yang-00";
}

identity management {
  description
    "Base identity for site management scheme.";
}
identity comanaged {
  base management;
  description
    "Base identity for comanaged site.";
}
identity customer-managed {
  base management;
  description
    "Base identity for customer managed site.";
}
identity provider-managed {
  base management;
  description
```



```
    "Base identity for provider managed site.";
}

identity address-allocation-type {
    description
        "Base identity for address-allocation-type
        for PE-CE link.";
}
identity pe-dhcp {
    base address-allocation-type;
    description
        "PE router provides DHCP service to CE.";
}
identity static-address {
    base address-allocation-type;
    description
        "PE-CE addressing is static.";
}
identity slaac {
    base address-allocation-type;
    description
        "Use IPv6 SLAAC.";
}

identity site-availability {
    description
        "Base identity for site availability.";
}
identity loadsharing {
    base site-availability;
    description
        "Identity for loadsharing site.";
}
identity loadsharing-ibgp {
    base loadsharing;
    description
        "Identity for ECMP ibgp based loadsharing site.";
}
identity loadsharing-eibgp {
    base loadsharing;
    description
        "Identity for ECMP eibgp based loadsharing site.";
}
identity primary-backup {
    base site-availability;
    description
        "Identity for primary-backup site.";
}
```



```
identity single-site {
    base site-availability;
    description
        "Identity for single site.";
}
identity access-availability-type {
    description
        "base identity for access-availability-type";
}
identity primary-access {
    base access-availability-type;
    description
        "Identity for primary access type.";
}
identity backup-access {
    base access-availability-type;
    description
        "Identity for backup access type.";
}
identity single-access {
    base access-availability-type;
    description
        "Identity for single access type.";
}
identity loadsharing-access {
    base access-availability-type;
    description
        "Identity for loadsharing access type.";
}

identity site-role {
    description
        "Base identity for site type.";
}
identity any-to-any-role {
    base site-role;
    description
        "Site in a any to any IPVPN.";
}
identity spoke-role {
    base site-role;
    description
        "Spoke Site in a Hub & Spoke IPVPN.";
}
identity hub-role {
    base site-role;
    description
```



```
    "Hub Site in a Hub & Spoke IPVPN.";
}

identity vpn-topology {
    description
        "Base identity for VPN topology.";
}
identity any-to-any {
    base vpn-topology;
    description
        "Identity for any to any VPN topology.";
}
identity hub-spoke {
    base vpn-topology;
    description
        "Identity for Hub'n'Spoke VPN topology.";
}
identity hub-spoke-disjoint {
    base vpn-topology;
    description
        "Identity for Hub'n'Spoke VPN topology
        where Hubs cannot talk between each other.";
}

identity multicast-tree-type {
    description
        "Base identity for multicast tree type.";
}

identity ssm-tree-type {
    base multicast-tree-type;
    description
        "Identity for SSM tree type.";
}
identity asm-tree-type {
    base multicast-tree-type;
    description
        "Identity for ASM tree type.";
}
identity bidir-tree-type {
    base multicast-tree-type;
    description
        "Identity for BiDir tree type.";
}

identity multicast-rp-discovery-type {
    description
```



```
        "Base identity for rp discovery type.";
    }
    identity auto-rp {
        base multicast-rp-discovery-type;
        description
            "Base identity for auto-rp discovery type.";
    }
    identity static-rp {
        base multicast-rp-discovery-type;
        description
            "Base identity for static type.";
    }
    identity anycast-rp {
        base multicast-rp-discovery-type;
        description
            "Base identity for anycast rp type.";
    }
    identity bsr-rp {
        base multicast-rp-discovery-type;
        description
            "Base identity for BDR discovery type.";
    }

    identity routing-protocol-type {
        description
            "Base identity for routing-protocol type.";
    }

    identity ospf {
        base routing-protocol-type;
        description
            "Identity for OSPF protocol type.";
    }

    identity bgp {
        base routing-protocol-type;
        description
            "Identity for BGP protocol type.";
    }

    identity static {
        base routing-protocol-type;
        description
            "Identity for static routing protocol type.";
    }

    identity rip {
        base routing-protocol-type;
```



```
    description
      "Identity for RIP protocol type.";
  }

  identity rip-ng {
    base routing-protocol-type;
    description
      "Identity for RIPng protocol type.";
  }

  identity vrrp {
    base routing-protocol-type;
    description
      "Identity for VRRP protocol type.
      This is to be used when LAN are directly connected
      to provider Edge routers.";
  }

  identity direct {
    base routing-protocol-type;
    description
      "Identity for direct protocol type.
      .";
  }

  identity protocol-type {
    description
      "Base identity for protocol field type.";
  }

  identity tcp {
    base protocol-type;
    description
      "TCP protocol type.";
  }

  identity udp {
    base protocol-type;
    description
      "UDP protocol type.";
  }

  identity icmp {
    base protocol-type;
    description
      "icmp protocol type.";
  }

  identity icmp6 {
    base protocol-type;
    description
```



```
        "icmp v6 protocol type.";
    }
    identity gre {
        base protocol-type;
        description
            "GRE protocol type.";
    }
    identity ipip {
        base protocol-type;
        description
            "IPinIP protocol type.";
    }
    identity hop-by-hop {
        base protocol-type;
        description
            "Hop by Hop IPv6 header type.";
    }
    identity routing {
        base protocol-type;
        description
            "Routing IPv6 header type.";
    }
    identity esp {
        base protocol-type;
        description
            "ESP header type.";
    }
    identity ah {
        base protocol-type;
        description
            "AH header type.";
    }
}

/* Groupings */

grouping vpn-svc-cfg {
    description
        "Configuration block for l3vpn svc.";

    leaf name {
        type string;
        description
            "Name of the VPN.";
    }
    leaf id {
        type uint32;
        description
```



```
        "VPN identifier. Local administration meaning.";
    }
    leaf customer-name {
        type string;
        description
            "Name of the customer.";
    }
    leaf topology {
        type identityref {
            base vpn-topology;
        }
        description
            "VPN topology.";
    }
}

list cloud-access {
    key cloud-identifier;

    leaf cloud-identifier {
        type string;
        description
            "Identification of cloud service. Local
            admin meaning.";
    }
    list authorized-sites {
        key site-id;

        leaf site-id {
            type leafref {
                path "../../../../../sites/site-id";
            }
            description
                "Site ID.";
        }
        description
            "List of authorized sites.";
    }
}

list denied-sites {
    key site-id;

    leaf site-id {
        type leafref {
            path "../../../../../sites/site-id";
        }
        description
            "Site ID.";
    }
    description
```



```
        "List of denied sites.";
    }
    leaf nat-enabled {
        type boolean;
        description
            "Control if NAT is required or not.";
    }
    leaf customer-nat-address {
        type inet:ipv4-address;
        description
            "NAT address to be used in case of public
            or shared cloud.
            This is to be used in case customer is providing
            the public address.";
    }
    description
        "Cloud access configuration.";
}

leaf mpls {
    type boolean;
    default false;
    description
        "The VPN is using Carrier Supporting Carrier,
        and so MPLS is required.";
}

container multicast {
    leaf-list tree-flavor {
        type identityref {
            base multicast-tree-type;
        }
        description
            "Type of tree to be used.";
    }
}

container rp {
    list static-rps {
        key "rp-address group";

        leaf rp-address {
            type union {
                type inet:ipv4-address;
                type inet:ipv6-address;
            }
            description
                "Defines the address of the RendezvousPoint.";
        }
        leaf group {
```



```
        type union {
            type inet:ipv4-address;
            type inet:ipv6-address;
        }
        description
        "Defines the address of the multicast group
        handled by the RP.";
    }
    description
    "List of RP to group mappings.";
}
leaf rp-discovery {
    type identityref {
        base multicast-rp-discovery-type;
    }
    description
    "Type of RP discovery used.";
}
leaf-list anycast-rp-location {
    type string;
    description
    "Encodes the location of the RPs in anycast RP
    deployment.";
}
description
"RendezvousPoint parameters.";
}
description
"Multicast global parameters for the VPN service.";
}
}

grouping customer-location-info {
    container location {
        leaf address {
            type string;
            description
            "Address (number and street)
            of the site.";
        }
    }
    leaf zip-code {
        type string;
        description
        "ZIP code of the site.";
    }
    leaf city {
        type string;
```



```
        description
            "City of the site.";
    }
    leaf country-code {
        type string;
        description
            "Country of the site.";
    }
    description
        "Location of the site.";
}
description
    "This grouping defines customer location
    parameters";
}

grouping site-diversity {
    container site-diversity {
        leaf type {
            type enumeration {
                enum "pop-diverse" {
                    description
                        "The access must use another PoP
                        compared
                        to other accesses in the same group.";
                }
                enum "pe-diverse" {
                    description
                        "The access must use another PE
                        compared
                        to other accesses in the same group.";
                }
            }
        }
        description
            "Diversity constraint type.";
    }

    leaf-list site-group {
        type uint32;
        description
            "IDs of group.";
    }
    description
        "Diversity constraint type.";
}
description
    "This grouping defines site diversity
    parameters";
```



```
}

grouping operational-requirements {
  leaf requested-site-start {
    type yang:date-and-time;
    description
      "Optional leaf indicating requested date
      and time
      when the service at a particular site is
      expected
      to start";
  }

  leaf requested-site-stop {
    type yang:date-and-time;
    description
      "Optional leaf indicating requested date
      and time
      when the service at a particular site is
      expected
      to stop";
  }

  leaf actual-site-start {
    type yang:date-and-time;
    description
      "Optional leaf indicating actual date
      and time
      when the service at a particular site
      actually
      started";
  }

  leaf actual-site-stop {
    type yang:date-and-time;
    description
      "Optional leaf indicating actual date
      and time
      when the service at a particular site
      actually
      stopped";
  }
  description
    "This grouping defines some operational parameters
    parameters";
}

grouping site-availability {
  container availability {
```



```
    container availability-type {
      leaf service-type {
        type identityref {
          base site-availability;
        }
        description
          "Type of availability,
          ex : single, backup, loadsharing";
      }
      leaf access-type {
        type identityref {
          base access-availability-type;
        }
        description
          "Access type within the service.";
      }
      description
        "Availability solution parameters";
    }
    leaf loadsharing-type {
      type identityref {
        base loadsharing;
      }
      description
        "Loadsharing flavor.";
    }
    description
      "Site availability parameters.";
  }
description
  "This grouping defines site availability
  parameters";
}

grouping flow-definition {
  container match-flow {
    leaf ipv4-src-prefix {
      type inet:ipv4-prefix;
      description
        "Match on IPv4 src address.";
    }
    leaf ipv6-src-prefix {
      type inet:ipv6-prefix;
      description
        "Match on IPv6 src address.";
    }
    leaf ipv4-dst-prefix {
      type inet:ipv4-prefix;
```



```
        description
            "Match on IPv4 dst address.";
    }
    leaf ipv6-dst-prefix {
        type inet:ipv6-prefix;
        description
            "Match on IPv6 dst address.";
    }
    leaf l4-src-port {
        type uint16;
        description
            "Match on layer 4 src port.";
    }
    leaf l4-dst-port {
        type uint16;
        description
            "Match on layer 4 dst port.";
    }
    leaf protocol-field {
        type union {
            type uint8;
            type identityref {
                base protocol-type;
            }
        }
        description
            "Match on IPv4 protocol or
            Ipv6 Next Header
            field.";
    }
}

description
    "Describe flow matching
    criterions.";
}
description
    "Flow definition based on criteria.";
}

grouping site-service-basic {
    leaf svc-input-bandwidth {
        type uint32;
        units bps;
        description
            "From the PE perspective, the service input
            bandwidth of the connection.";
    }
    leaf svc-output-bandwidth {
        type uint32;
```



```
        units bps;
        description
            "From the PE perspective, the service output
            bandwidth of the connection.";
    }
    leaf svc-mtu {
        type uint16;
        units bytes;
        description
            "MTU at service level.
            If the service is IP,
            it refers to the IP MTU.";
    }
    description
        "Defines basic service parameters for a site.";
}
grouping site-service-protection {
    container traffic-protection {
        leaf link-local-protection {
            type boolean;
            description
                "Enables protection of PE-CE link.";
        }
        leaf node-local-protection {
            type boolean;
            description
                "Enables protection against local
                PE node failure.";
        }
        leaf node-global-protection {
            type boolean;
            description
                "Enables protection against remote
                PE node failure.
                This is also called PIC Edge.
                ";
        }
        description
            "Fast reroute service parameters
            for the site.";
    }
    description
        "Defines protection service parameters for a site.";
}
grouping site-service-mpls {
    container mpls {
        leaf signalling-type {
            type enumeration {
```



```
        enum "ldp" {
            description
                "Use LDP as signalling
                 protocol between PE and CE.";
        }
        enum "bgp" {
            description
                "Use BGP 3107 as signalling
                 protocol between PE and CE.
                 In this case, bgp must be also
                 configured
                 as routing-protocol.
                 ";
        }
    }
    description
        "MPLS signalling type.";
}
description
    "This container is used when customer provides
     MPLS based services.
     This is used in case of Carrier
     Supporting Carrier.";
}
description
    "Defines MPLS service parameters for a site.";
}
grouping site-service-qos-profile {
    container qos {
        container qos-classification-policy {
            list rules {
                key id;

                leaf id {
                    type uint16;
                    description
                        "ID of the rule.";
                }
                uses flow-definition;

                leaf target-class-id {
                    type string;
                    description
                        "Identification of the
                         class of service.
                         This identifier is internal to
                         the administration.";
                }
            }
        }
    }
}
```



```
        description
        "List of marking rules.";
    }
    description
    "Need to express marking rules ...";
}
leaf std-qos-profile {
    type string;
    description
    "QoS profile to be used";
}
container custom-qos-profile {
    list class {
        key class-id;

        leaf class-id {
            type string;
            description
            "Identification of the
            class of service.
            This identifier is internal to
            the administration.";
        }
        leaf rate-limit {
            type uint8;
            units percent;
            description
            "To be used if class must
            be rate
            limited. Expressed as
            percentage of the svc-bw.";
        }
        leaf priority-level {
            type uint8;
            description
            "Defines the level of the
            class in
            term of priority queueing.
            The higher the level is the
            higher
            is the priority.";
        }
        leaf guaranteed-bw-percent {
            type uint8;
            units percent;
            description
            "To be used to define the
            guaranteed
```



```
        BW in percent of the svc-bw
        available at the priority-level.";
    }
    description
        "List of class of services.";
    }
    description
        "Custom qos profile.";
    }
    description
        "QoS configuration.";
    }
    description
        "This grouping defines QoS parameters
        for a site";
    }

    grouping site-security-authentication {
        container authentication {
            description
                "Authentication parameters";
        }
        description
            "This grouping defines authentication
            parameters
            for a site";
    }

    grouping site-security-encryption {
        container encryption {
            leaf enabled {
                type boolean;
                description
                    "If true, access encryption is required.";
            }
            leaf layer {
                type enumeration {
                    enum layer2 {
                        description
                            "Encryption will occur at layer2.";
                    }
                    enum layer3 {
                        description
                            "IPSec is requested.";
                    }
                }
            }
            description
```



```
        "Layer on which encryption is applied.";
    }
    container encryption-profile {
        choice profile {
            case provider-profile {
                leaf profile-name {
                    type string;
                    description
                        "Name of the SP profile
                        to be applied.";
                }
            }
            case customer-profile {
                leaf algorithm {
                    type string;
                    description
                        "Encryption algorithm to
                        be used.";
                }
                choice key-type {
                    case psk {
                        leaf preshared-key {
                            type string;
                            description
                                "Key coming from
                                customer.";
                        }
                    }
                    case pki {
                    }
                    description
                        "Type of keys to be used.";
                }
            }
        }
        description
            "Choice of profile.";
    }
    description
        "Profile of encryption to be applied.";
}
description
    "Encryption parameters.";
}
description
    "This grouping defines encryption parameters
    for a site";
}
```



```
grouping site-security-acl {
  container access-control-list {
    description
      "Access control list.";
  }

  description
    "This grouping defines ACL security parameters
    for a site";
}

grouping site-attachment-bearer {
  container bearer {
    leaf type {
      type string;
      description
        "Type of bearer Ethernet, DSL, Wireless ...
        Operator specific.";
    }
    leaf bearer-reference {
      type string;
      description
        "This is an internal reference for the
        service provider.";
    }
    description
      "Bearer specific parameters.
      To be augmented.";
  }
  description
    "Defines physical properties of
    a site attachment.";
}

grouping site-attachment-ip-connection {
  container ip-connection {
    container ipv4 {
      leaf address-allocation-type {
        type identityref {
          base address-allocation-type;
        }
        description
          "Defines how addresses are allocated.
          Need to be detailed further.";
      }
      leaf subnet-prefix {
        type inet:ipv4-prefix;
        description
```



```
        "Interco subnet.";
    }
    description
        "IPv4 specific parameters";
}
container ipv6 {
    leaf address-allocation-type {
        type string;
        description
            "Defines how addresses are allocated.
            Need to be detaillled further.";
    }
    leaf subnet-prefix {
        type inet:ipv6-prefix;
        description
            "Interco subnet.";
    }
    description
        "IPv6 specific parameters";
}
container oam {
    container bfd {
        leaf bfd-enabled {
            type boolean;
            description
                "BFD activation";
        }

        choice holdtime {
            case profile {
                leaf profile-name {
                    type string;
                    description
                        "Service provider well
                        known profile.";
                }
            }
            description
                "Service provider well
                known profile.";
        }
        case fixed {
            leaf fixed-value {
                type uint32;
                units msec;
                description
                    "Expected holdtime
```



```
        expressed
        in msec.";
    }
}
description
    "Choice for holdtime flavor.";
}
description
    "Container for BFD.";
}
description
    "Define the OAM used on the connection.";
}
list routing-protocols {
    key type;

    leaf type {
        type identityref {
            base routing-protocol-type;
        }
        description
            "Type of routing protocol.";
    }
}

container ospf {
    when "type = 'ospf'" {
        description
            "Only applies
            when protocol is OSPF.";
    }
    leaf-list address-family {
        type identityref {
            base rt:address-family;
        }
        description
            "Address family to be activated.";
    }
    leaf area-address {
        type yang:dotted-quad;
        description
            "Area address.";
    }
    leaf metric {
        type uint16;
        description
            "Metric of PE-CE link.";
    }
}
```



```
list sham-link {
  key target-site;

  leaf target-site {
    type leafref {
      path "../../../../../.."
      +"../sites/site-id";
    }
    description
      "Target site for the sham link
      connection.";
  }
  leaf metric {
    type uint16;
    description
      "Metric of the sham link.";
  }
  description
    "Creates a shamlink with another
    site";
}
description
  "OSPF specific configuration.";
}

container bgp {
  when "type = 'bgp'" {
    description
      "Only applies when
      protocol is BGP.";
  }
  leaf-list address-family {
    type identityref {
      base rt:address-family;
    }
    description
      "Address family to be activated.";
  }
  description
    "BGP specific configuration.";
}

container static {
  when "type = 'static'" {
    description
      "Only applies when protocol
      is static.";
  }
  leaf-list address-family {
```



```
        type identityref {
            base rt:address-family;
        }
        description
            "Address family to be activated.";
    }
    description
        "Static routing
        specific configuration.";
}
container rip {
    when "type = 'rip'" {
        description
            "Only applies when
            protocol is RIP.";
    }
    leaf-list address-family {
        type identityref {
            base rt:address-family;
        }
        description
            "Address family to be
            activated.";
    }

    description
        "RIP routing specific
        configuration.";
}

container vrrp {
    when "type = 'vrrp'" {
        description
            "Only applies when
            protocol is VRRP.";
    }
    leaf-list address-family {
        type identityref {
            base rt:address-family;
        }
        description
            "Address family to be activated.";
    }
    description
        "VRRP routing specific configuration.";
}
```



```
        description
        "List of routing protocols used
        on the site.
        Need to be augmented.";
    }
    description
    "Defines connection parameters.";
}
description
"This grouping defines IP connection parameters.";
}

grouping site-service-multicast {
    container multicast {
        leaf multicast-site-type {
            type enumeration {
                enum receiver-only {
                    description
                    "The site has only receivers.";
                }
                enum source-only {
                    description
                    "The site has only sources.";
                }
                enum source-receiver {
                    description
                    "The site has both
                    sources & receivers.";
                }
            }
        }
        default "source-receiver";
        description
        "Type of multicast site.";
    }
    container multicast-transport-protocol {
        leaf ipv4 {
            type boolean;
            default true;
            description
            "Enables ipv4 multicast transport";
        }
        leaf ipv6 {
            type boolean;
            default false;
            description
            "Enables ipv6 multicast transport";
        }
    }
    description
```



```
        "Defines protocol to transport multicast.";
    }
    leaf protocol-type {
        type enumeration {
            enum host {
                description
                "
                Hosts are directly connected
                to the provider network.
                Host protocols like IGMP, MLD
                are required.
                ";
            }
            enum router {
                description
                "
                Hosts are behind a customer router.
                PIM will be implemented.
                ";
            }
            enum both {
                description
                "Some Hosts are behind a customer
                router and some others are directly
                connected to the provider network.
                Both host and routing protocols must be
                used. Typically IGMP and PIM will be
                implemented.
                ";
            }
        }
        default "both";
        description
        "Multicast protocol type to be used
        with the customer site.";
    }

    description
    "Multicast parameters for the site.";
}
description
"Multicast parameters for the site.";
}

grouping site-management {
    container management {
        leaf type {
            type identityref {
```



```
        base management;
    }
    description
    "Management type of the connection.";
}
leaf management-transport {
    type identityref {
        base rt:address-family;
    }
    description
    "Transport protocol used for management.";
}
leaf address {
    type union {
        type inet:ipv4-address;
        type inet:ipv6-address;
    }
    description
    "Management address";
}

description
    "Management configuration";
}
description
    "Management parameters for the site.";
}

grouping site-vpn-policy {
    container vpn-policy {
        list entries {
            key id;

            leaf id {
                type uint32;
                description
                "Unique identifier for
                the policy entry.";
            }
            container filter {
                container lan-prefixes {
                    list ipv4-lan-prefixes {
                        key lan;

                        leaf lan {
                            type inet:ipv4-prefix;
                            description
                            "Lan prefixes.";
                        }
                    }
                }
            }
        }
    }
}
```



```
    }
    description "
        List of LAN prefixes
        for the site.
    ";
}
list ipv6-lan-prefixes {
    key lan;

    leaf lan {
        type inet:ipv6-prefix;
        description
            "Lan prefixes.";
    }
    description "
        List of LAN prefixes
        for the site.
    ";
}
description
    "LAN prefixes from the customer.";
}
leaf-list lan-tag {
    type string;
    description
        "List of lan-tags to be matched.";
}
description
    "If used, it permit to split site LANs
    among multiple VPNs.
    If no filter used, all the LANs will be
    part of the same VPNs with the same
    role.";
}
container vpn {
    leaf vpn {
        type leafref {
            path "../.../.../vpn-svc/name";
        }
        description
            "Reference to an IPVPN.";
    }
    leaf site-role {
        type identityref {
            base site-role;
        }
        description
            "Role of the site in the IPVPN.";
    }
}
```



```
        }
        description
            "List of VPNs the LAN is associated to.";
    }
    description
        "List of entries for export policy.";
}
description
    "VPN policy.";
}
description
    "VPN policy parameters for the site.";
}

grouping site-customer-routing-specific-information {
    container customer-specific-information {
        leaf name {
            type string;
            description
                "Name of the customer router.";
        }
        leaf autonomous-system {
            type uint32;
            description
                "AS number.";
        }
        leaf interface {
            type string;
            description
                "Interface reference of the access.";
        }
        list customer-lan-connection {
            key "address";
            leaf address {
                type union {
                    type inet:ipv4-address;
                    type inet:ipv6-address;
                }
                description
                    "Address given by the customer on its LAN
                    for the SP router.";
            }
            leaf lan-protocol {
                type identityref {
                    base rt:address-family;
                }
                description
                    "Transport protocol used on LAN.";
            }
        }
    }
}
```



```
    }
    description
      "List of customer LAN to be connected
      directly on the CE.";
  }
  container cascaded-lan-prefixes {
    list ipv4-lan-prefixes {
      key lan;

      leaf lan {
        type inet:ipv4-prefix;
        description
          "Lan prefixes.";
      }
      leaf lan-tag {
        type string;
        description
          "Internal tag to be used in vpn
          policies.";
      }
      leaf next-hop {
        type inet:ipv4-address;
        description
          "Nexthop address to use at customer
          side.";
      }
      description "
        List of LAN prefixes for
        the site.
        ";
    }
    list ipv6-lan-prefixes {
      key lan;

      leaf lan {
        type inet:ipv6-prefix;
        description
          "Lan prefixes.";
      }
      leaf lan-tag {
        type string;
        description
          "Internal tag to be used
          in vpn policies.";
      }
      leaf next-hop {
        type inet:ipv6-address;
        description
```



```
        "Nexthop address to use at
        customer side.";
    }
    description "
        List of LAN prefixes for the site.
        ";
    }
    description
        "LAN prefixes from the customer.";
}

description
    "Customer premise configuration.";
}
description
    "Customer routing specific information.";
}

grouping site-maximum-routes {
    container maximum-routes {
        list address-family {
            key af;

            leaf af {
                type identityref {
                    base rt:address-family;
                }
                description
                    "Address-family.";
            }
            leaf maximum-routes {
                type uint32;
                description
                    "Maximum prefixes the VRF can
                    accept for this
                    address-family.";
            }
            description
                "List of address families.";
        }

        description
            "Define maximum-routes for the VRF.";
    }
    description
        "Define maximum-routes for the site.";
}
```



```
/* Main blocks */

container l3vpn-svc {
    list vpn-svc {
        key name;

        uses vpn-svc-cfg;

        description "
            List of VPN services.
            ";
    }

    list sites {
        key site-id;

        leaf site-id {
            type string;
            description
                "Identifier of the site.";
        }

        leaf apply-template {
            type leafref {
                path "../..sites/site-id";
            }
            description
                "It we would be good to ensure that
                site-id is template type";
        }

        uses operational-requirements;
        uses customer-location-info;
        uses site-diversity;
        uses site-availability;
        uses site-management;
        uses site-vpn-policy;
        uses site-maximum-routes;
        uses site-customer-routing-specific-information;

        container security {
            leaf apply-template {
                type leafref {
                    path "../..sites/site-id";
                }
                description
                    "It we would be good to ensure that site-id is
```



```
        template type";
    }
    uses site-security-authentication;
    uses site-security-encryption;
    uses site-security-acl;

    description
        "Site specific security parameters.";
}

container attachment {
    leaf apply-template {
        type leafref {
            path "../../sites/site-id";
        }
        description
            "It we would be good to ensure
            that site-id is
            template type";
    }
    uses site-attachment-bearer;
    uses site-attachment-ip-connection;

    description
        "Parameters of the attachement.";
}

container service {
    leaf apply-template {
        type leafref {
            path "../../sites/site-id";
        }
        description
            "It we would be good to ensure that site-id is
            template type";
    }
    uses site-service-basic;
    uses site-service-qos-profile;
    uses site-service-protection;
    uses site-service-mpls;
    uses site-service-multicast;

    description
        "Service parameters on the attachement.";
}

description "List of sites.";
}
```



```
list site-templates {
  key site-template-id;

  leaf site-template-id {
    type string;
    description
      "Identifier of the site.";
  }

  uses operational-requirements;
  uses customer-location-info;
  uses site-diversity;
  uses site-availability;
  uses site-management;
  uses site-vpn-policy;
  uses site-maximum-routes;
  uses site-customer-routing-specific-information;

  container security {
    uses site-security-authentication;
    uses site-security-encryption;
    uses site-security-acl;

    description
      "Site specific security parameters.";
  }

  container attachment {
    uses site-attachment-bearer;
    uses site-attachment-ip-connection;

    description
      "Parameters of the attachement.";
  }

  container service {
    uses site-service-basic;
    uses site-service-qos-profile;
    uses site-service-protection;
    uses site-service-mpls;
    uses site-service-multicast;

    description
      "Service parameters on the attachement.";
  }

  description "List of sites.";
}
```



```
        description
        "Main container for L3VPN service configuration.";
    }

}
<CODE ENDS>
```

9. To do list / Open questions / Open issues

Open issues :

- o Multicast modeling : fix RP definition
- o VPN policy modeling : does this new version correctly address it ?
- o Operational states : do we need more ?

TODO :

- o Security items
- o Bearer details
- o Generalization to CE to CE IPVPN
- o Need to look at interAS case
- o Need to look at Hybrid VPNs
- o Need to look at OAM (only BFD now)
- o ...

10. Security Considerations

TBD.

11. Contributors

12. Acknowledgements

Thanks to Qin Wu, Maxim Klyus, Luis Miguel Contreras, Gregory Mirsky, Zitao Wang, Jing Zhao, Kireeti Kompella, Eric Rosen, Aijun Wang and Andrew Leu for the contributions to the document.

13. IANA Considerations

TBD.

14. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4110] Callon, R. and M. Suzuki, "A Framework for Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", [RFC 4110](#), DOI 10.17487/RFC4110, July 2005, <<http://www.rfc-editor.org/info/rfc4110>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), DOI 10.17487/RFC4364, February 2006, <<http://www.rfc-editor.org/info/rfc4364>>.
- [RFC4577] Rosen, E., Psenak, P., and P. Pillay-Esnault, "OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4577](#), DOI 10.17487/RFC4577, June 2006, <<http://www.rfc-editor.org/info/rfc4577>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), DOI 10.17487/[RFC4862](#), September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", [RFC 6020](#), DOI 10.17487/RFC6020, October 2010, <<http://www.rfc-editor.org/info/rfc6020>>.
- [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, <<http://www.rfc-editor.org/info/rfc6241>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", [RFC 6513](#), DOI 10.17487/RFC6513, February 2012, <<http://www.rfc-editor.org/info/rfc6513>>.

Appendix A. Example: NETCONF <get> Reply

This section gives an example of a reply to the NETCONF <get> request for a device that implements the data model defined in this document. The example is written in XML.

Authors' Addresses

Stephane Litkowski
Orange Business Service

Email: stephane.litkowski@orange.com

Rob Shakir
BT

Email: rob.shakir@bt.com

Luis Tomotaki
Verizon

Email: luis.tomotaki@verizon.com

Kevin D'Souza
ATT

Email: kd6913@att.com

