

Expires: April 2017

October 31, 2016

Client Defined Private Networks laid over Thin CPEs
draft-dunbar-opsawg-private-networks-over-thin-cpe-01

Abstract

This document specifies a type of private networks that interconnect thin CPEs at multiple client sites by IP tunnels, or more specifically, lay over multiple client sites' Thin CPEs via IP tunnels. Those private overlay networks not only interconnect those sites by secure IP tunnels but can also enforce the client specified policies to govern how applications or hosts within those sites communicate and how to access public internet.

Hosts or applications in those sites can be interconnected by Layer 2 networks or/and by Layer 3 networks. The network that the IP tunnels are traversing can be IPv4 or IPv6 networks. This document describes the special properties of the client defined networks over Thin CPEs.

A separate draft will describes the special features that those IP tunnels need to have in order to interconnect multiple sites as if those sites are directly connected by wires and how communication policies are enforced.

Status of This Document

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 April 31, 2017.

Copyright Notice

Copyright (c) 2016 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.....	4
2. Terminology.....	4
2.1. Requirements Language.....	4
2.2. Terms defined in this document.....	4
3. Brief Description of the Private networks laid over Thin CPEs..	6
4. Overlay Private Network Configuration from Client Perspective..	8
4.1. Client Defined Overlay Private Networks.....	8
4.2. Client's site Configuration.....	8
4.3. Internet Gateway for each Site.....	9
4.4. Overlay-VPN Gateway.....	9
4.5. Interconnection among Sites.....	9
5. Protocols needed for the Client Defined Overlay Private Networks	
.....	10
5.1. Thin CPE Auto Instantiation.....	10
5.2. Network agnostic interworking.....	10
5.3. Gateway Anchor Auto-Selection.....	10
5.4. Middle boxes auto-creation and rules exchanges.....	10
5.5. Thin CPE on Third Party location.....	11
5.6. Client Defined Policies for traffic to/from client sites..	11
5.7. QoS policies.....	11
5.8. Explicit Service functions chain specified by clients....	11
5.9. Thin CPE monitoring.....	11

5.10. Alarm & Events via Thin CPE.....	11
5.11. Resource management via Thin CPE instantiated in Remote Locations.....	11
5.12. Client traffic flows management, monitoring, and reporting	11
6. Networks carried by IP tunnels in conjunction with existing L2VPN/L3VPN.....	12
7. IANA Considerations.....	12
8. Security Considerations.....	12
9. References.....	12
9.1. Normative References.....	12
9.2. Informative Reference.....	12
10. Authors' Addresses.....	12
11. Contributors Addresses.....	13

1. Introduction

This document specifies a type of private networks that interconnect thin CPEs at multiple client sites by IP tunnels, or more specifically, lay over multiple client sites' Thin CPEs via IP tunnels. Those private overlay networks not only interconnect those sites by secure IP tunnels but can also enforce the client specified policies to govern how applications or hosts within those sites communicate and how to access public internet.

Hosts or applications in those sites can be interconnected by Layer 2 networks or/and by Layer 3 networks. The network that the IP tunnels are traversing can be IPv4 or IPv6 networks. This document describes the special properties of the client defined networks over Thin CPEs.

For ease of description, the "Client Defined Private Overlay Network" is also called the client's "Overlay Private Network" or "Overlay Virtual Private Network (Overlay-VPN)" throughout this document.

A separate draft will describes the special features that those IP tunnels need to have in order to interconnect multiple sites as if those sites are directly connected by wires and how communication policies are enforced.

2. Terminology

2.1. 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].

2.2. Terms defined in this document

Internet Gateway: a network function, which can be a physical device in the provider site or a virtual function instantiated to connect client site traffic to the public internet, and can enforce client specified policies.

Overlay Private Network: private network over a set of thin CPEs at multiple sites created by clients or users, who don't need to worry

about how thin CPEs are connected nor the protocol setting at network side. The "Overlay Private Network" not only interconnects multiple sites by (secure) IP tunnels but can also enforce the client specified policies to govern how applications or hosts within those sites communicate and how to access public internet.

Overlay-VPN: Overlay Private Network.

Provider site: the location where the provider have access to the devices or equipment.

Site: A place that contains switches, routers, services, appliances and these devices are configured to form L2 domain (s) or L3 domain. For example an Enterprise company data center, a college campus network center. For L3 subnets, either private IPv4 or IPv6 address or public IPv4 or Ipv6 address can be used.

SITE: Site Interconnection Tunnel Encapsulation Protocol

Thin CPE: a simple device at a customer premise that maps the site local traffic to either the IP tunnels connected to the Internet Gateway, or the IP tunnels connected to the VPN Gateway.

Overlay-VPN Gateway: the function (which can be virtual) that establish private (secure) connections to other sites belonging to the same client.

3. Brief Description of the Private networks laid over Thin CPEs

The following figure depicts multiple overlay private networks that interconnect the client's various sites. Note, the Overlay Private Network is marked as "Overlay" in the figure. The client can create multiple overlay private networks and then assign each site to specific overlay private networks. The client also specify the policies on what traffic to/from the clients can be exchanged with external network, which are enforced by the "Internet gateways" created by the provider.

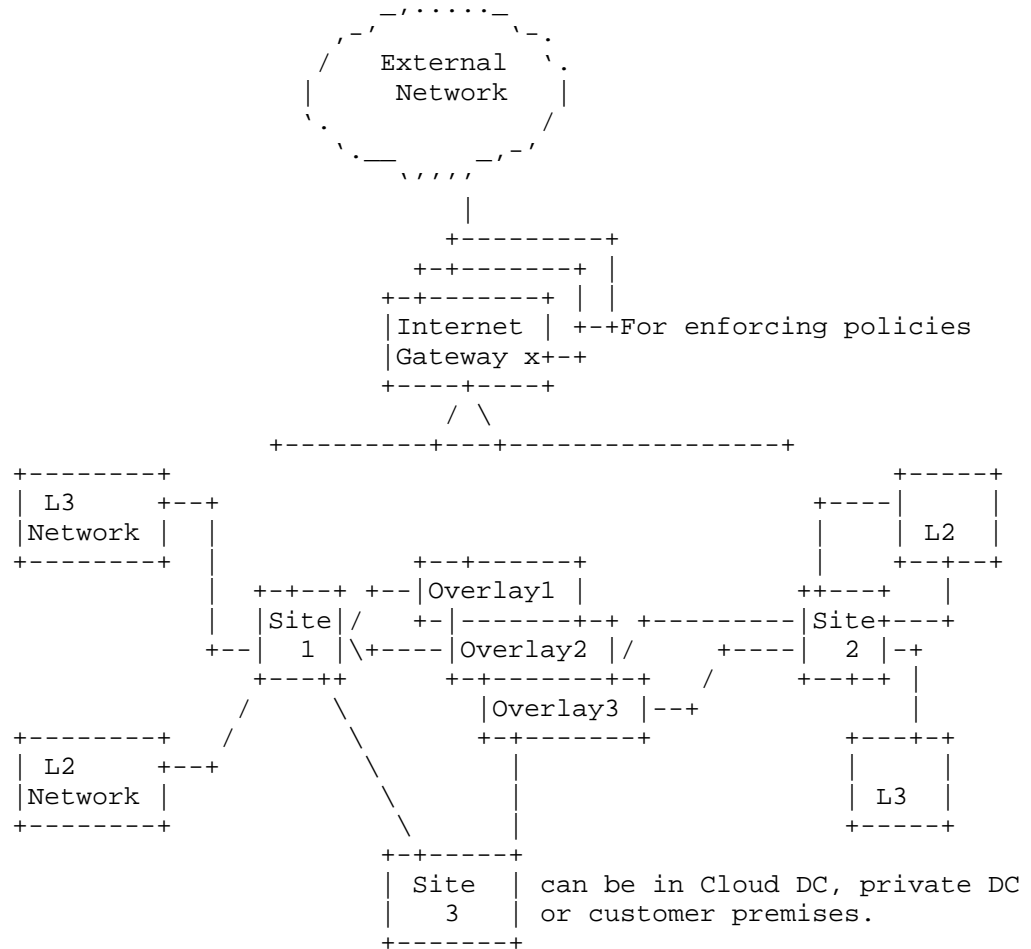


Figure 1 Overlay Private Networks interconnecting sites

Here are some key properties of Client defined Overlay Private Networks:

- Each client "Site" has a Thin CPE that is connected to a VPN gateway which is hosted in the provider site via IP Tunnel (which can be secured per customer request). The Thin CPE can be software image instantiated on virtual machines, physical CPE, or other form factors.




Figure 2 site Thin CPE connect to Overlay GW via IP Tunnel

- Each Thin CPE is connected to an "Internet Gateway" via IP Tunnel (that is automatically created by provider). The "Internet Gateway", virtual or physical, can be located anywhere. An IP Tunnel is created automatically between the Thin CPE and the "Internet Gateway".
- When the provider don't own the infrastructure to interconnect multiple sites, (secure) IP Tunnels are created among each site's VPN Gateway, so that each site's local networks (L2 or L3) attached to the Thin CPEs are interconnected as if those networks are directly connected by physical wire.
- Some traffic between Thin CPE have to go through secure tunnel, e.g. IPSec. Clients can specify what traffic to go through secure tunnels without specifically worrying about how to establish or maintain the secure tunnels. The client traffic can be carried by VxLAN (for interconnecting layer 2 traffic) or GRE (for L3 traffic) over the IPSEC tunnel.
- Client specifies the policies on how/what/when hosts from the interconnected sites can communicate with external peers; E.g. Hosts in one Layer 2 domain from one site may communicate with hosts in different Layer 2 domains in different sites.

The Client Defined Overlay Networks can be viewed by client as their own private networks. For ease of description, the terminology "Overlay Private Network" or "Overlay-VPN" is used throughout this document to refer to this kind of client defined overlay network over Thin CPEs.

"Overlay Private Network" is different from the IETF's L2VPN or L3VPN for the following reasons:

- Overlay-Private-Network is built upon IP network (whereas L2VPN/L3VPN is built upon MPLS network),
- Traffic originated from a client's site (where Thin CPE is instantiated) not only can communicate with hosts in other sites of the client via IP tunnels, but also can communicate with public internet (governed by the policies specified by the client),
- Client's site Thin CPE don't participate in IGP or BGP routing with provider side. Client can specify the prefixes and/or VLANs for each site so that they can be reached by external hosts,
- IP tunnel is automatically created between a Thin CPE and provider site where VPN gateway and internet gateway are instantiated and maintained.

4. Overlay Private Network Configuration from Client Perspective

4.1. Client Defined Overlay Private Networks

The client can specify multiple overlay private networks (a.k.a. Overlay-VPNs). Client can specify which sites connect to which Overlay-VPNs. Each Site can connect to multiple Overlay-VPNs.

As features on Thin CPE are very limited, each Overlay-VPN has its own Overlay VPN gateway in provider site to connect to Thin CPE via IP tunnel, as depicted in Figure 2 above.

4.2. Client's site Configuration

For each site, the client needs to specify:

- Site Identifier (include unique system Identifier, name, etc.)
- VLANs enabled on the site (i.e. the VLANs enabled on the client facing ports of the Thin CPE).
- Subnets from the site (i.e. the subnets enabled on the client facing ports of the Thin CPE)

- IP address for the Overlay-VPN Gateway that connect other sites belonging to the client
- IP address for the Internet Gateway

The configuration on the site is mainly for the Thin CPE instantiated on the site. Therefore, the client also needs to specify which VLANs/subnets are enabled on the ports of the Thin CPE facing the local network on the site.

4.3. Internet Gateway for each Site

Each site is associated with an Internet Gateway, which is automatically created by the provider. The Interconnect gateway can be a physical device on the provider site or a virtual function, to connect client site traffic to the public internet, and can enforce client specified policies.

Considering one client can have multiple sites in different geographic locations, the client can specify different policies for traffic to/from each site.

4.4. Overlay-VPN Gateway

The Overlay-VPN Gateway is on the provider site, connected to Thin CPE via IP tunnel. The purpose of the Overlay-VPN Gateway is to connect a site to its specified Overlay VPNs. Each site can be connected to multiple Overlay VPNs.

For each Overlay-VPN gateway, the client needs to specify:

- Identifier
- Which VPN is the Gateway connected to
- Upstream bandwidth from Thin CPE to the Overlay VPN GW
- Downstream bandwidth from the Overlay VPN GW to the Thin CPE

4.5. Interconnection among Sites

For each Overlay VPN, the Client can choose which sites are connected by specifying the VPN Gateway associated with each site.

5. Protocols needed for the Client Defined Overlay Private Networks

5.1. Thin CPE Auto Instantiation

Thin CPE is a simple device that maps the site local traffic to either the IP tunnels connected to the Internet Gateway, or the IP tunnels connected to the VPN Gateway.

5.2. Network agnostic interworking

IP tunnels are automatically created between Thin CPE and (Internet/VPN) gateways based on the traffic to the access network.

For Layer 2 traffic from the client local site, VxLAN is used to build the IP Tunnels to the site's Internet gateway or VPN gateway respectively.

For Layer 3 traffic from the client local site, GRE is used to build the IP Tunnels to the site's Internet gateway or VPN gateway respectively.

If the client specifies secure connection to other sites, IPsec is added to the tunnels between the Thin CPE and the VPN Gateway.

5.3. Gateway Anchor Auto-Selection

For each client site, internet gateway and VPN gateway will be automatically instantiated.

There will be protocol extension needed for the creation/deletion process and how NAT is used for client traffic from each site.

5.4. Middle boxes auto-creation and rules exchanges

To be added

5.5. Thin CPE on Third Party location

Thin CPEs can also be instantiated third party premises, such as cloud data centers. The instantiated Thin CPE can establish IP tunnels with the client's Internet Gateway or VPN Gateway.

5.6. Client Defined Policies for traffic to/from client sites

Depending on the policies specified by the clients, the Thin CPE jointly with the virtual GW will select the appropriate network security functions, i.e. (virtual) FW, IPS, IDS, or others to enforce the policies specified by the clients.

The policies specified by the clients will be more expressed in clients' oriented language, e.g. using client Identifier or virtual addresses (instead of IP addresses of the actual packets traverse the FW). Those policies will be translated to the implementable rules to the chosen network security functions, such as FW.

5.7. QoS policies

To be added

5.8. Explicit Service functions chain specified by clients

Clients can query network service functions available to them and the capabilities of those functions. Then, the client can choose a set of them, either in strict sequence or simply as a set to apply to their traffic.

The policies to service functions can follow the guideline specified by [I2NSF-framework].

5.9. Thin CPE monitoring

5.10. Alarm & Events via Thin CPE

To be added

5.11. Resource management via Thin CPE instantiated in Remote Locations

To be added

5.12. Client traffic flows management, monitoring, and reporting

To be added

6. Networks carried by IP tunnels in conjunction with existing
L2VPN/L3VPN

7. IANA Considerations

To be added

8. Security Considerations

To be added.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC2119, March 1997.

9.2. Informative Reference

[I2NSF-Framework] Lopez, D, et al, "Framework for Interface to
Network security functions", draft-ietf-i2nsf-framework-04,
Oct 2016

10. Authors' Addresses

Linda Dunbar
Huawei Technologies
Email: linda.dunbar@huawei.com

Lucy Yong
Huawei Technologies
Email: lucy.yong@huawei.com

Song Xiao Li
Huawei Technologies
Email: sxlin@huawei.com

11. Contributors Addresses

Xuan Ming fu
Huawei Technologies
xuanmingfu@huawei.com

