Nested VPN Technology White Paper

This document provides generic technical information, some of which might not be applicable to your products.

The information in this document is subject to change without notice.

Contents

Overview

Background

In an MPLS L3VPN network, generally a service provider runs the MPLS L3VPN backbone and provides VPN services through provider edge (PE) devices. VPN users access the MPLS L3VPN network through customer edge (CE) devices, which are connected with PEs to allow VPN users at different sites to communicate with each other.

In the scenario described above, users are on ordinary IP networks and all VPN users at a site belong to the same VPN. However, in actual applications, VPN user networks can be dramatically different in form and complexity. For example, a VPN user may have a large number of devices that reside in different management domains, and may wish to further divide the VPN into several smaller VPNs based on management domains.

In some other cases, a VPN user network itself or one of its sites might be an MPLS L3VPN network. For users that belong to the same VPN but reside at different sites to communicate, the service provider's MPLS L3VPN network must be able to transfer users' internal VPN attributes to different sites transparently.

To summarize, the requirements for an MPLS L3VPN network include:

- All sites of a user as a whole are regarded as a VPN user on the service provider's MPLS L3VPN network.
- A user may have internal VPNs, with all sites or only some of the sites having internal VPNs.
- Users of an internal VPN can communicate with each other, whether they reside at the same site or not. Users of different internal VPNs are isolated from each other.
- Because of cost and management concerns, users do not want to deploy all internal VPNs on the service provider's PEs directly.
- The service provider manages only VPNs on PEs, leaving the management of internal VPNs to their respective owners.

The nested VPN technology answers the requirements. Nested VPN allows a customer to manage its own internal VPNs. In this way, VPNs can be divided into two levels: large VPNs and small VPNs. Mutual access is allowed between a large VPN and a small VPN, but not between two small VPNs. This approach works well for scenarios where users have different access privileges.

Benefits

The main features of the nested VPN technology are:

- The nested VPN can implement VPN aggregation, so that a customer's multiple internal VPNs can be aggregated into one VPN on the service provider's MPLS L3VPN network.
- Nested VPN supports both symmetric networking and asymmetric networking. That is, sites of a customer network can have the same number or different number of internal VPNs.
- Nested VPN allows for two-level VPNs and flexible access between them. In addition, nested VPN supports multiple levels of nesting for internal VPNs of a customer. Sites of a user can access each other remotely in the same VPN.
- As a network solution for multilevel access to the service provider network, nested VPN features simple networking and flexible applications.

Nested VPN implementation

Concepts

Provider edge router (PE)

A PE device resides at the edge of a service provider network and is connected to one or more CEs. All MPLS VPN services are processed on PEs.

Customer edge router (CE)

A CE device resides on a customer network and has one or more interfaces directly connected to a service provider network. It does not support MPLS.

Provider router (P)

A P device is a core device on a service provider network. It is not directly connected to any CEs. A P device has only basic MPLS forwarding capability and does not handle VPN routing information.

VPN instance

MPLS L3VPN uses VPN instances for route isolation and data independence and security between VPNs. A VPN instance is also called a virtual routing and forwarding (VRF) instance.

A VPN instance has the following components:

- A separate Label Forwarding Information Base (LFIB).
- An IP routing table.
- Interfaces bound to the VPN instance.
- VPN instance administration information, including route distinguishers (RDs), route targets (RTs), and route filtering policies.

To associate a site with a VPN instance, bind the VPN instance to the PE's interface connected to the site. A site can be associated with only one VPN instance, and different sites can be associated with the same VPN instance. A VPN instance contains the VPN membership and routing rules of associated sites.

Site

A site for VPN users has the following features:

- A site is a group of IP systems with IP connectivity that does not rely on any service provider networks.
- The classification of a site depends on the topology relationship of the devices, rather than the geographical positions. However, the devices at a site are, in most cases, adjacent to each other geographically.
- The devices at a site can belong to multiple VPNs, which means that a site can belong to multiple VPNs.
- A site is connected to a provider network through one or more CEs. A site can contain multiple CEs, but a CE can belong to only one site.

Sites connected to the same provider network can be classified into different sets by policies. Only the sites in the same set can access each other through the provider network. Such a set is called a VPN.

Mechanism

In an ordinary MPLS L3VPN network, PEs supports only IPv4 route exchange with private networks. The nested VPN technology supports exchange of VPNv4 routes with private networks, so it allows a customer to manage its own internal VPNs.

One user site has its internal VPNs

A VPN customer can have a network which by itself is also an MPLS L3VPN network. In the simplest case, the VPN customer has only one PE and creates multiple internal VPNs on the PE to segregate local services. The network model can be depicted as follows:

- The customer network connects to the service provider network through a private network interface on the service provider's PE.
- The customer's PE exchanges VPNv4 routes with the service provider's PE directly.
- The service provider's PE treats the customer network as an ordinary VPN user, and regards the customer's PE as its CE.
- The customer network treats the service provider's PE as an ordinary PE of itself.

Figure 1 One user site having its internal VPNs

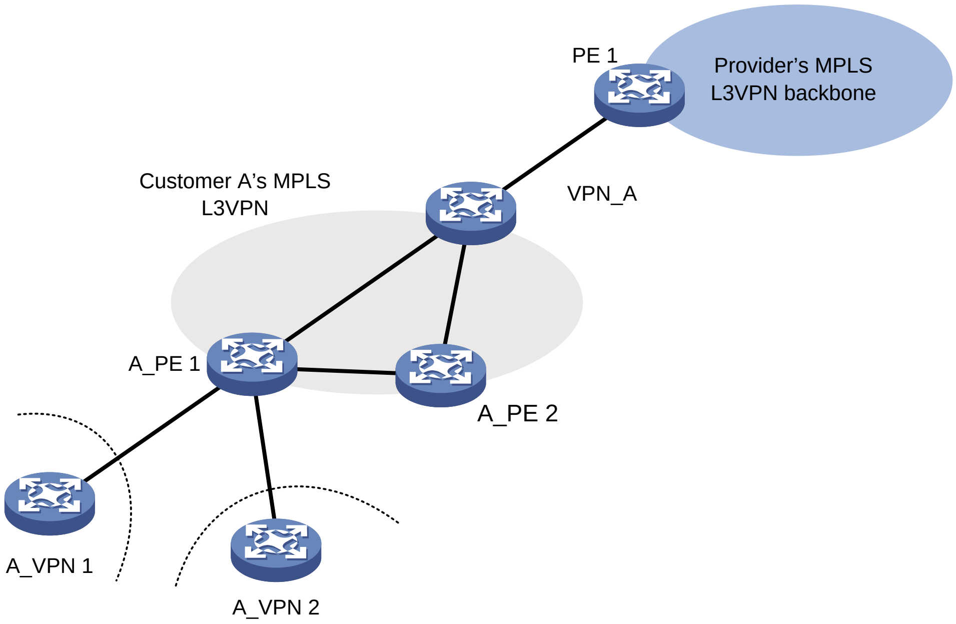

As shown in Figure 1, PE 1 of the service provider's MPLS L3VPN network connects a customer VPN called VPN_A, and VPN_A has two internal VPNs: A_VPN 1 and A_VPN 2.

A_PE 1 and A_PE 2 are PEs of the customer network. (There might be P devices on the customer network). A-PE 1 has two internal VPNs connected to it.

VPN_A treats PE 1 of the service provider as an ordinary PE on its network. PE 1 treats A-PE1 and A_PE2 of VPN_A as its CEs. VPNv4 routes are exchanged among A-PE 1, A_PE 2 and PE 1, which can be done by a route reflector.

For the service provider, there is only one customer VPN: VPN_A.

Multiple user sites have internal VPNs

In most cases, a user VPN has multiple sites connected with the service provider's MPLS L3VPN network. Some of the sites might have the same or similar internal VPN structure. Some sites might have their internal VPNs connected to the service provider's PE devices directly. Some sites each might have only one large VPN.

For mutual access between user's internal VPNs and between an internal VPN and the user's large VPN, the local VPN information of a site must be transmitted to other sites of the VPN user through the service provider's MPLS VPN.

Therefore, in the nested VPN solution, the service provider's PE needs to exchange VPNv4 routes directly with users' internal PEs.

Figure 2 Multiple user sites having internal VPNs

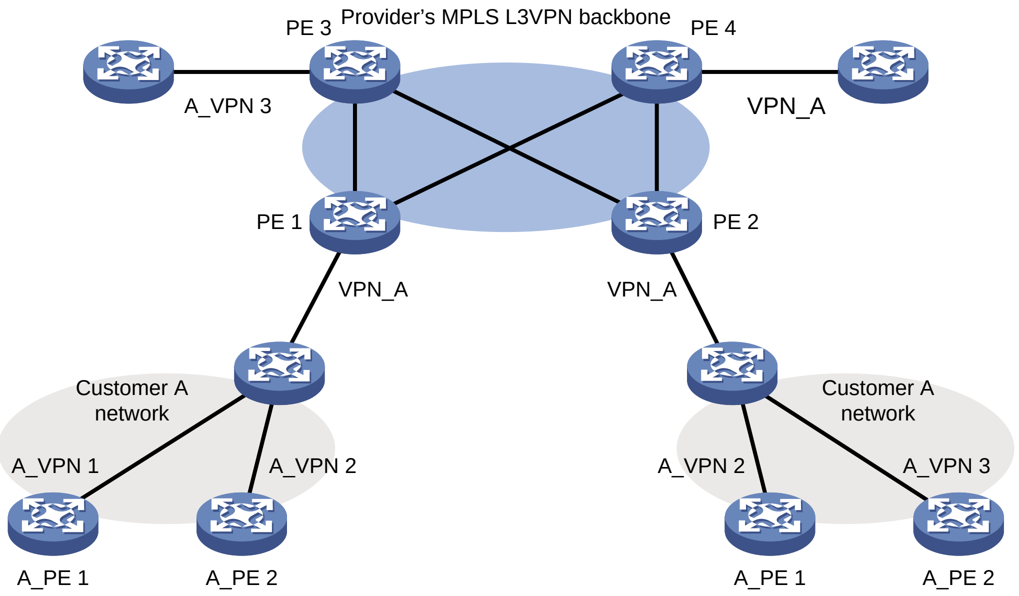

In a nested VPN network, routing information is propagated by using the following process:

1. After receiving VPN routes from customer CEs, a customer PE advertises VPNv4 routes to the provider CEs through MP-BGP.
2. The provider CEs advertise the VPNv4 routes to a provider PE through MP-BGP.
3. After receiving a VPNv4 route, the provider PE keeps the customer's internal VPN information, and appends the customer's MPLS VPN attributes on the service provider network. It replaces the RD of the VPNv4 route with the RD of the customer's MPLS VPN on the service provider network. It also adds the export route-target (ERT) attribute of the customer's MPLS VPN on the service provider network to the extended community attribute list of the route. The internal VPN information for the customer is maintained on the provider PE.
4. The provider PE advertises VPNv4 routes carrying the comprehensive VPN information to the other PEs of the service provider.
5. After another provider PE receives the VPNv4 routes, it matches the VPNv4 routes to the import targets of its local VPNs. Each local VPN accepts routes of its own and advertises them to provider CEs. If a provider CE is connected to a provider PE through an IPv4 connection, the PE advertises IPv4 routes to the CE. If it is a VPNv4 connection (a customer MPLS VPN network), the PE advertises VPNv4 routes to the CE.
6. After receiving VPNv4 routes from the provider CE, a customer PE matches those routes to local import targets. Each customer VPN accepts only its own routes and advertises them to connected customer CEs.

In this way, the service provider network can correctly process the following VPN relationships:

- Relationship between the VPNs for customers on the service provider network.
- Relationship between the internal VPNs of each customer.
- Relationship between an internal VPN of a customer and the VPN that the customer belongs to.

Restrictions

The internal VPNs of each customer and the customer VPNs of the service provider must have matching route target configuration.

Make sure the CEs of the service provider do not filter received BGP VPNv4 routes based on route targets.

Application scenarios

Symmetric networking of nested VPN

Symmetric networking means that each site of a customer has its own internal VPNs and they all use the nested VPN technology to connect to the service provider's PEs. For example, a university campus network resides on a large MPLS L3VPN network as one VPN. The campus network has several districts, and each district provides multiple majors, such as art, literature, physical education, and mathematics. The university intends to build a VPN network for each major. As shown in Figure 3, each district creates a standard MPLS VPN network with a separate VPN for each major. Then, each district uses the nested VPN technology to connect itself as a CE to the service provider's PE. In this way, the majors in different districts can communicate within their respective internal VPNs.

Figure 3 Symmetric networking of nested VPN

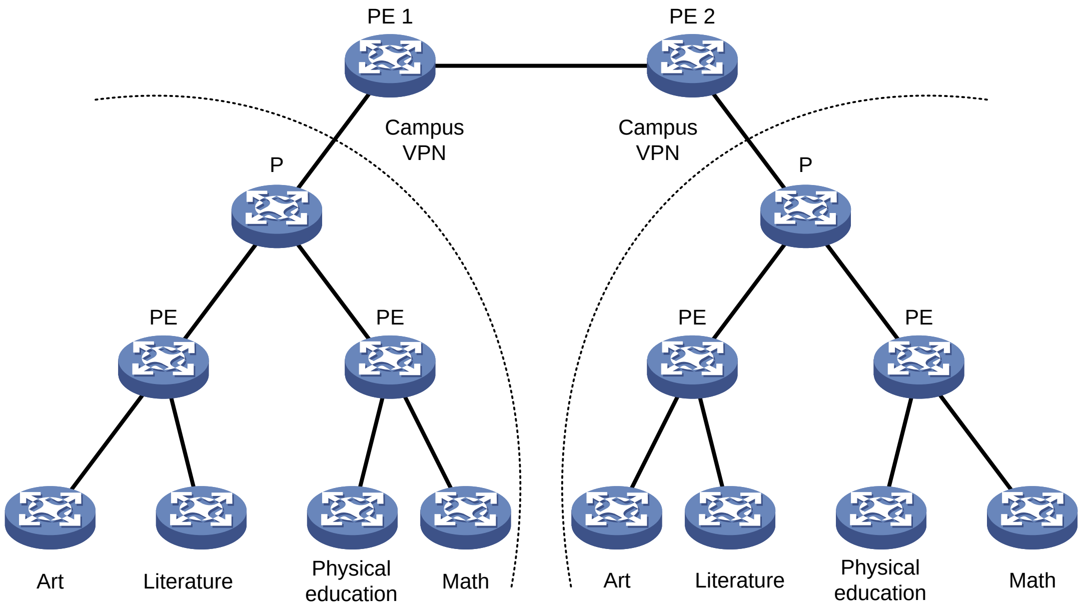

Asymmetric networking of nested VPN

If a site has many internal VPNs, it can use the nested VPN technology to aggregate the internal VPNs on its internal PE to form a large VPN before connecting to the service provider's PE, as site A does in Figure 4. Sites that do not have many internal VPNs can connect to the service provider's PEs directly, as site B and site C do in Figure 4.

This kind of networking is called asymmetric networking.

The asymmetric networking scheme can also be used to protect a VPN access device which does not have powerful routing or MPLS forwarding capabilities by hiding it behind a backbone router.

Generally, devices that form the MPLS VPN backbone should have powerful forwarding and routing capabilities. VPN access devices can reside behind the backbone devices to avoid being overwhelmed by large quantities of MPLS VPN routes or tunnels on the backbone. This networking method can also separate VPN access networks from the MPLS VPN backbone. For example, with nested VPN, the VPNs on the BRAS device are regarded as internal VPNs and are aggregated into a large VPN to connect to the MPLS VPN backbone.

Figure 4 Asymmetric networking of nested VPN

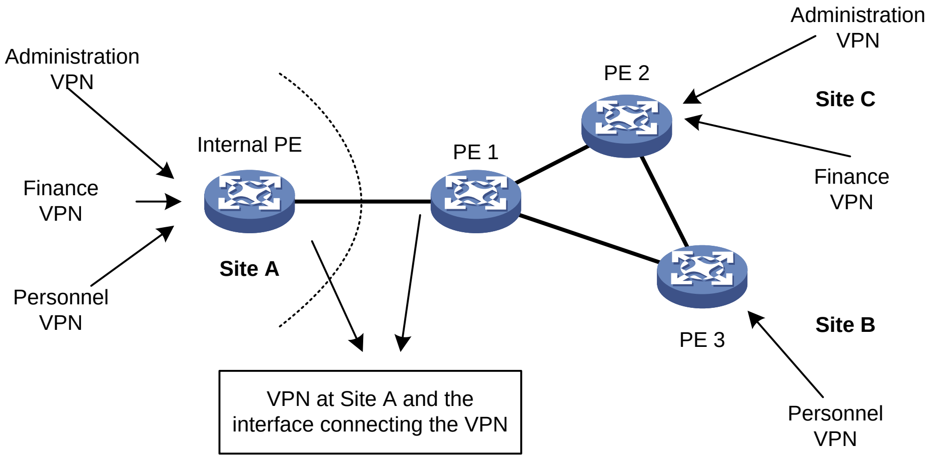

Mutually accessible two-level VPN using nested VPN

After adopting the nested VPN technology, there will be two levels of mutually accessible VPNs: customer-level VPN and customer internal VPN, which are also called large VPN and small VPN, respectively.

This two-level VPN application is widely adopted in actual networking applications. For example, devices like servers which need to be accessed by all small VPNs can be directly connected to the service provider's PE as CEs of the backbone. Hosts for special purposes, such as the administrative host, can also be connected to the service provider's PE as the large VPN's user, ensuring that these hosts can access all internal networks.

As shown in Figure 5, VPN_A is a large VPN, and A_VPN 1 and A_VPN 2 are small VPNs. The server cluster is directly connected to the service provider's PE 1 as a CE of the large VPN, and can be accessed by all small VPNs. A host for special purpose, the administrative host in this example, is connected to the service provider's PE 3 as the user of the large VPN, ensuring that it can access all internal networks.

Figure 5 Two-level VPN using nested VPN

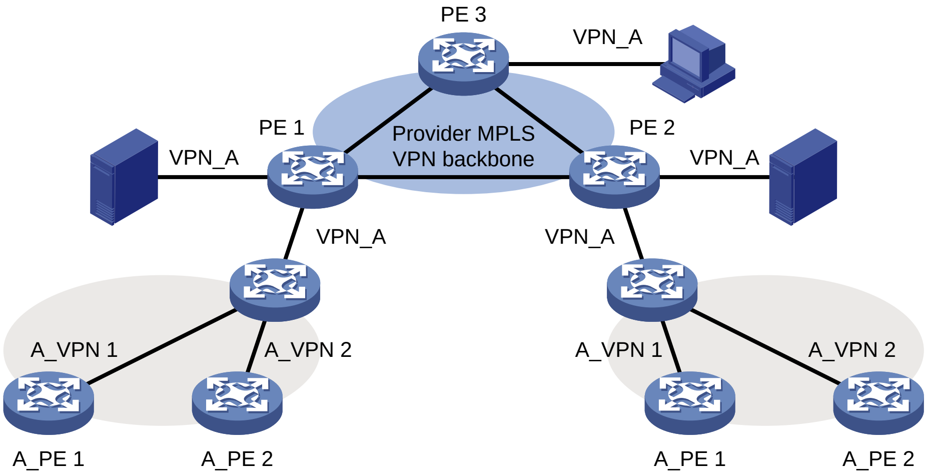
PE 3
VPN_A
PE 1
Provider MPLS
VPN backbone
PE 2
VPN_A
VPN_A
VPN_A
VPN_A
A_VPN 1
A_VPN 2
A_VPN 1
A_VPN 2
A_PE 1
A_PE 2
A_PE 1
A_PE 2