Provider-Provisioned VPN Working Group
Internet Draft
Expiration Date: January 2002

Giles Heron
PacketExchange Ltd.

Rick Wilder
Masurgy

Juha Heinanen
Song Networks

Tom Soon
SBC Communications

Luca Martini
Level3 Communications

Vach Kompella
Joe Regan
Sunil Khandekar
TiMetra Networks

July 2001

Requirements for Virtual Private Switched Networks

draft-heron-ppvpn-vpsn-reqmts-00.txt

Status of this Memo

Abstract

This document defines requirements for provision of Virtual Private Switched Network services over tunneled packet switched networks. These requirements are common for both IP and MPLS networks.

Table of Contents

1. Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119

2. Placement of this Memo in Sub-IP Area

RELATED DOCUMENTS

draft-vkompella-ppvpn-vpsn-mpls-00.txt

WHERE DOES THIS FIT IN THE PICTURE OF THE SUB-IP WORK

This fits in the PPVPN box.

WHY IS IT TARGETED AT THIS WG

Because it specifies requirements for layer 2 VPNs which emulate a LAN segment (or "Virtual Private Switched Networks")

JUSTIFICATION

We believe the WG should consider this draft because it specifies requirements for a class of layer 2 VPN that has up to now not been sufficiently addressed in this WG.

3. Introduction

This document describes service provider requirements for providing Virtual Private Switched Networks over an IP or MPLS based network infrastructure. A Virtual Private Switched Network (or VPSN) is a class of VPN that allows the connection of multiple sites in a sin- gled bridged domain over a provider managed IP or MPLS network. All customer sites in the VPSN appear to be on the same LAN regardless of their location.

VPSN services (often known in this context as "Transparent LAN" ser- vices, or TLS) have traditionally been offered by service providers over an ATM infrastructure. The TLS service is provisioned using mesh of ATM PVCs between locations. Some customers prefer TLSs to point-to-point Frame Relay circuits since a TLS hides the complexity of designing and managing multiple Frame Relay circuits. While this reduces the complexity for the customer, the service provider must deal with managing and operating the ATM network in addition to edge Ethernet switches. This model does not scale well and is expensive

to maintain and manage for the service provider.

It should be noted that while VPSNs provide certain advantages to both customers and service providers, they have properties that do not scale over large numbers of sites or where there are large numbers of hosts in the VPSN. In addition the single metric that a routing protocol will assign to a layer 2 subnet emulated using VPSN makes this model unsuitable if the costs of connectivity between different sites in the VPSN differ. The most likely application of this model is to connect a few, or a few tens of, sites over a metropolitan or regional reach and with only a single customer router (or a few customer hosts) connected to the VPSN at each site.

The customer and provider should choose judiciously whether to implement VPSN or some other network connectivity model.

VPSN can be contrasted with two other service models. One is the Virtual Private Routed Network (VPRN) [1], in which the service provider participates in the customer's routing protocol. This involves minimal change to configuration on the customer edge (CE) router, but adds complexity in the provider edge (PE) router. The second model is a set of point-to-point circuits (L2VPN) [2], which involves configuration on the CE router, but simplifies the PE router operation. VPSN, on the other hand, requires very little CE or PE router configuration.

The scope of this document will be limited to supporting Ethernet as the access framing technology for VPSN implementation.

This document discusses the motivation and requirements of network-based, provider provisioned VPSN over a native IP or MPLS network. A proposal for the implementation of VPSN over MPLS is given in [3].

4. Virtual Private Switched Network

We define a Virtual Private Switched Network or VPSN as a service where the service provider provides an emulated layer-2 bridged network that connects multiple customer sites over an IP or MPLS network infrastructure.

To the customer, this appears as a single VLAN or a layer-2 switched LAN segment where a single connection into the LAN (VPSN) offers any-to-any connectivity between all sites. All traffic is switched based on MAC addresses but forwarded between all participating PE devices using IP or MPLS tunnels. Tunneling traffic, using either IP or MPLS, allows the service provider to use a single network for multiple services without requiring any overlay networks.

5. Attributes of VPSN

5.1. Simplicity

With VPSN, the CE has a single logical attachment to the PE. A mesh of point-to-point tunnels is built between the participating PE devices to carry multiple services. This offers any-to-any connectivity without requiring the customer to manage connectivity to every site. For the customer, this simplifies manageability by reducing the operational complexity of managing circuits to every site.

Each CE device requires minimal to no configuration since this is analogous to connecting a device into a VLAN. If the CE device is a router, it only needs to be configured with a network address and mask for the VPSN subnet and the customer specific routing protocol. For smaller locations, customers can simply connect a layer-2 device into the VPSN. This eliminates the need for trained internetworking personnel at each site.

5.2. Ease of Provisioning

Participating PE devices must be configured with knowledge of all other endpoints (PEs) for the service. When a new customer site is commissioned on the VPSN, the existing PE devices are told about the PE where the new site is connected. The service provider is not required to over-provision each PE in anticipation of new customer sites being added later.

Note that it is also possible for participating PE devices to automatically discover all other PE devices for a service. This is considered to be preferable to configuring each PE with knowledge of the other PEs, since it makes it easier to add new sites to a VPSN and also removes the risk of a VPSN being configured as a partial rather than as a full mesh of PEs.

5.3. Virtual Bridging

The service provider is responsible for providing a true layer-2 switched service between customer sites. The customer specific VPSN instance, in the PE, acts as a virtual bridge and performs normal bridging functions.

5.4. Protocol Independence

The customer is free to run any Ethernet encapsulated layer-3 protocol between multiple sites since the traffic is switched based on layer-2 MAC addresses. IP or MPLS based transport tunnels are used to carry traffic for multiple VPSNs between common PEs.

5.5. Routing Independence

As in case of L2VPN [2], the service provider does not participate in any customer routing. The customer is free to run any routing protocol. There is no interaction between the service provider and customer routing protocols.

Since the VPSN depends on the service provider routing protocol to provide a resilient network service it will generally be the case that this protocol will be tuned for faster reconvergence than the customer routing protocols that are likely to run over the service. Note that instability may occur if the customer routing protocol detects a failure and starts to reroute traffic before the service provider routing protocol has been able to do so.

5.6. Sub-Rate Services

Since Ethernet is the service interface, the access speeds available for the VPSN can match the standard Ethernet LAN speeds. At the same time, sub-rate access speeds can be offered in granular increments by leveraging the traffic management capabilities of the PE equipment. This enables service providers to customize access rates that best suit each customer.

5.7. Quality of Service

MPLS-TE tunnels that offer specific quality of service can be set up between PE devices. Customer Ethernet packets marked with IEEE 802.1p [4] bits can be mapped to tunnels that offer differentiated quality of service.

In addition the IEEE 802.1p bits may be mapped to MPLS EXP bits or IP DSCPs.

5.8. Scaling

Each customer VPSN instance on the PE router is required to maintain a Forwarding Information Base (FIB) of the MAC addresses that are part of the customer VPSN. This raises the question of scalability on the part of the PE routers that support the VPSN implementation.

Each PE router must deal with maintaining a MAC address FIB per VPSN instance. No interaction is required with the routing protocol on the PE router. The solution must scale to large numbers of VPSN instances per PE router.

Unlike a VPRN implementation, where the virtual router instance must deal with all the network routes behind CE devices at each site, the VPSN implementation requires the PE router to have knowledge of only the MAC addresses of a single broadcast domain.

However, it is not practical to scale a VPSN implementation over a large number of hosts in the emulated network. This is because each PE router has to have knowledge of all MAC addresses in the emulated network.

In addition it is not practical to scale VPSN implementation over a large number of sites. This will invariably lead to scaling issues associated with flooding, replication, aging and learning.

It is likely that the most common implementation of VPSN will have a customer-managed router connected into the VPSN at each site. Thus, the MAC addresses are limited to the number of sites connected to the VPSN.

6. VPSN Requirements

A network based, provider provisioned VPSN service MUST support the following features.

6.1. Network Infrastructure

The VPSN service MUST be provided transparently over a shared IP or MPLS based network infrastructure. The network core MUST NOT require any knowledge of layer-3 protocols addressing to support VPSN service. It MUST NOT be necessary to introduce any spanning tree state into the service provider network to support the VPSN service. Resiliency and fail-over capabilities for the VPSN MAY be offered using IP or MPLS techniques only. This does not preclude running a spanning tree protocol (STP) on the customer-facing network.

The network core MUST provide any-to-any connectivity between the PE devices.

The service provider SHOULD be able to offer different levels of service to its customers by building service specific tunnels or tunnels based on differing quality of service.

6.2. Network Transport

Customer packets belonging to different VPSN services may be carried over an IP network or an MPLS network using L2TP, GRE or MPLS tunneling techniques.

The tunneling techniques used SHOULD be those defined in the PWE3 Working Group.

Broadcast, multicast and unknown frames MUST either be replicated by the ingress PE router or by the use of pre-configured multicast tunnels. In the former case performance will be severely limited by the requirement to replicate packets at the ingress, and hence this architecture SHOULD not be used to support broadcast or multicast services.

6.3. Frame Size

The service SHOULD support the standard Ethernet frame size of between 64 and 1518 octets, as defined in [5].

In addition the service MAY offer support for larger frame sizes ("jumbo frames").

Any frame smaller than 64 octets or larger than the supported maximum frame size MUST be discarded at the ingress PE.

6.4. Data Delivery

Valid frames offered to the service MUST be delivered in order and without duplication. Frames MAY be discarded under network failure or under very high network load. Note that the responsibility for delivering frames in order and without duplication SHOULD be delegated to the underlying network.

Invalid frames (i.e. invalid frames as defined in [5], and short or long frames as defined above) MUST be discarded at the ingress PE.

Frames MUST NOT be corrupted in transit. The Ethernet FCS MAY be

carried across the network and optionally verified before the frame is forwarded by the egress PE. Alternatively the FCS MAY be stripped at the ingress PE and regenerated by the egress PE, in which case all links in the underlying network MUST utilise a 32 bit or better FCS.

6.5. Traffic Separation

Complete separation MUST be maintained between multiple VPSN instances. Traffic separation between different VPSN instances MUST be provided using IEEE 802.1Q tags [4] on the customer facing ports or by assigning a different Ethernet port for each VPSN instance. Traffic separation in the core MUST be provided by using an appropriate encapsulation (for example that defined in [6]) in the provider network.

In addition multiple VLANs MAY be provisioned over a single VPSN instance. In this case traffic separation between the different VLANs MUST be provided using IEEE 802.1Q tags [4] in the customer facing ports and in the provider network. Traffic separation between this and other VPSN instances MUST be provided by assigning a dedicated Ethernet port for this instance and by using an appropriate encapsulation (for example that defined in [6]) in the provider network.

6.6. Membership Integrity

The signalling used to establish membership of the VPSN MUST be secured to prevent any unauthorised participation in a VPSN.

The underlying tunneling protocol used to transport frames from one PE to another MUST be secured to prevent injection of unauthorised traffic into the VPSN.

6.7. Protocol Independence

The VPSN service MUST be able to support any Ethernet encapsulated layer-3 protocol, and MUST NOT rely on protocol specific features to enhance support for particular layer-3 protocols.

6.8. VPSN Instance

A VPSN instance per emulated network per PE MUST be supported. Each VPSN instance MUST be capable of supporting normal LAN bridging functions such as MAC learning and aging at the PE on customer and provider facing ports (learning tunnels) and replication of frames with broadcast, multicast and unknown MACs.

If multiple VLANs are being supported over a single VPSN, as described above, the VPSN instance MUST associate learned MAC addresses with the correct VLAN.

In addition, some form of loop detection SHOULD be provided at the PE. Loop prevention MAY be provided at the PE.

6.9. Duplicate MAC Addresses

The PE router MUST be able to support duplicate MAC addresses that are part of separate VPSN instances.

If multiple VLANs are being supported over a single VPSN, as described above, the PE router MUST be able to support duplicate MAC addresses that are part of the same VPSN instance but which are associated with different VLANs.

6.10. MAC Address Limiting

The PE SHOULD be able to limit the number of MAC addresses per VPSN instance. This will limit the amount of memory consumed by the VPSN FIB.

6.11. Any to Any Connectivity

A single physical or logical (802.1Q VLAN) customer connection from each site MUST be sufficient to provide any-to-any connectivity just like a LAN.

6.12. Provisioning

VPSN MUST require minimal or no configuration on the CE device, depending on the CE device that connects into the VPSN. In addition, service providers MUST be able to offer VPSN service without requiring substantial configuration at each participating PE. When additional sites are provisioned, minimal configuration MAY be required on the existing PEs that have CE devices connected in the same VPSN.

6.13. Network Management

Management of the underlying tunnels SHOULD be delegated to the management function of the underlying packet switched network, and management of the Ethernet layer SHOULD be delegated to the customer's network management function.

Each VPSN instance MUST maintain appropriate state information of other VPSN instances in the VPSN. This state information MUST include, but is not limited to, physical interface state and reachability from the local VPSN instance. While further OAM functionality, such as the ability to trigger remote network loopbacks or to verify that frames are successfully delivered to the intended remote VPSN instance, is desirable it is to be considered out of scope for this effort. Other groups are defining such functionality, for example the LSP-ping effort [7] and the MPLS OAM effort [8], and it may be possible to leverage this work in VPSN implementations.

Each VPSN instance MUST maintain counts of the number of frames transmitted to and received from each remote PE, as well as counts of the number of frames replicated to all available remote PEs for each of the three categories: broadcast, multicast and unknown.

7. Security Considerations

The traffic separation and membership integrity requirements described above MUST be adhered to in order for the solution to be considered minimally secure. It is recommended that any security measures over and above this level (for example data authentication or encryption) be applied by the VPSN customer on an edge-to-edge (i.e. CE router to CE router) or an end-to-end (i.e. application to application) basis.

8. Intellectual Property Disclaimer

This document is being submitted for use in IETF standards discussions.

9. References

[1] "A Framework for IP Based Virtual Private Networks", B. Gleeson, A. Lin, J. Heinanen, G. Armitage, A. Malis. RFC 2764. February 2000.

[2] "MPLS-Based Layer 2 VPNs", K. Kompella, M. Leelanivas, Q. Vohra, R. Bonica, E. Metz. draft-kompella-mpls-l2vpn-02.txt. Work in progress.

[3] "Virtual Private Switched Network Services over an MPLS Network", V. Kompella et al. draft-vkompella-ppvpn-vpsn-mpls-00.txt. Work in progress.

[4] IEEE STD 802.1Q-1998. December 1998.

[5] IEEE STD 802.3-2000. October 2000.

[6] "Encapsulation Methods for Transport of Layer 2 Frames Over MPLS", L. Martini et. al. draft-martini-l2circuit-encap-mpls-02.txt. Work in progress.

[7] "Detecting Data Plane Liveliness in RSVP-TE", P. Pan, N. Sheth, D. Cooper. draft-pan-lsp-ping-00.txt. Work in progress.

[8] "OAM Functionality for MPLS Networks", N. Harrison et al. draft-harrison-mpls-oam-00.txt. Work in progress.

10. Author Information

Giles Heron
PacketExchange Ltd.
The Truman Brewery
91 Brick Lane
LONDON E1 6QL
United Kingdom
Tel.: +44 7880 506185
Email: giles@packetexchange.net

Rick Wilder
Masergy Inc.
2901 Telestar Ct.
Falls Church, VA 22042

Juha Heinanen
Song Networks, Inc.

Tom S. C. Soon
SBC Technology Resources Inc.
4698 Willow Road
Pleasanton, CA 94588
Tel.: +1 (925) 598-1227
Email: sxsoon@tri.sbc.com

Luca Martini
Level 3 Communications, LLC.
1025 Eldorado Blvd.
Broomfield, CO, 80021
Email: luca@level3.net

Vach Kompella
TiMetra Networks
274 Ferguson Dr.
Mountain View, CA 94043
Tel.: +1 (650) 237-5152
Email: vkompella@timetra.com

Joe Regan
TiMetra Networks
274 Ferguson Dr.
Mountain View, CA 94043
Tel.: +1 (650) 237-5103
Email: jregan@timetra.com

Sunil Khandekar
TiMetra Networks
274 Ferguson Dr.
Mountain View, CA 94043
Tel.: +1 (650) 237-5105
Email: sunil@timetra.com