Network Working Group
Internet Draft
Document: draft-tsenevir-l2-req-01.txt
Category: Informational

Tissa Senevirathne
(Force10)

Waldemar Augustyn
(Nortel)

Loa Anderson
Tove Madsen
(Utfors Bredband AB)

Pascal Menazes
(TeraBeam)

July, 2001

Requirements for Network Based Layer 2 VPN

Status of this Memo

1. Abstract

Requirements that needed to be considered when implementing or providing Layer 2 NBVPN services are presented. This document does not suggest any specific solution, instead outline the requirements that need to be satisfied.

2. Conventions used in this document

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].

Table of Content

3. Introduction

Traditionally, the typical connectivity between a service provider and a customer is a WAN link with some type of a point-to-point protocol. This arrangement was borne out of the necessity to traverse TDM circuits originally designed for voice traffic. The introduction of WAN links to network architecture significantly increased the complexity of the network topologies and required high skilled personnel to manage and maintain the network.

The L2 protocols running over WAN networked served the sole purpose of bringing customer traffic to the network core. These protocols did not find any other use in a typical customer network and only burdened the customer with the necessity to acquire knowledge skills and maintenance ability.

With the explosive growth of the Internet, the metro and wide area networks change and now offer data oriented connectivity. The lower cost of fiber and the popularity of Optical Ethernet brings the opportunity to eliminate unnecessary complexity of WAN networks. More and more, the customers have the opportunity to extend their networks via a service provider by running native L2 protocols, such as Ethernet, without having to go through a complex and unnecessary transformation to WAN protocols.

It is anticipated that Network Based Layer 2 VPN services may be more transparent to providers, easier to manage, and less expensive to maintain. It is also anticipated that networks based

Layer 2 VPN may be a key service in future metro and wide area service infrastructure.

In this document we attempt to outline the requirements for network based Layer 2 VPN services. The work presented in this document intends to serve as a guideline for equipment vendors and service providers. This document does not discuss any specific protocol, however it does specifies requirements that needed to be satisfied by a given protocol.

3.1 Background

Typically Layer 2 NBVPN may be used to connect a local LAN segment to multiple remote LAN segments. These LAN segments may contain one or more Work stations. A pictorial representation of a typical deployment is depicted below.

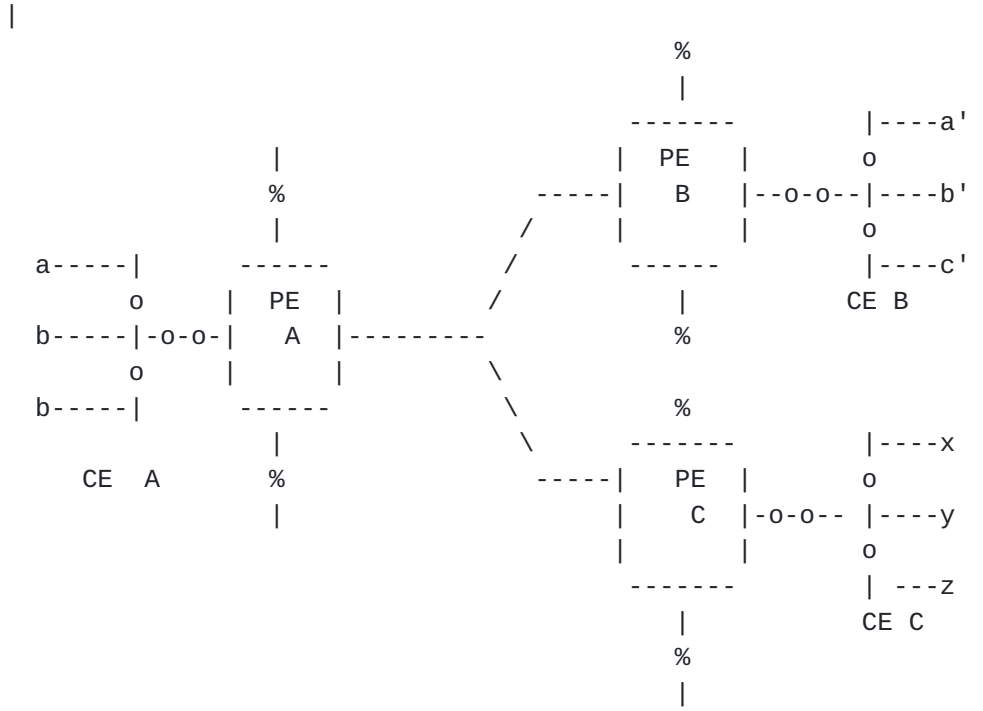

CE - Customer Edge (Represent one more Devices in a LAN)

PE - Provider Edge

-o- Physical Connection

--- Logical Connection

-%- Partition of Functionality

Fig: Typical Layer 2 NBVPN deployment

In general Layer 2 NBVPN contain multiple end points (sites). These end points may or may not be within the same service provider infrastructure. Within a given service provider infrastructure there exists multiple Layer 2 NBVPN planes that belongs to different customers. [3] Has logically represented this as a tower of disks.

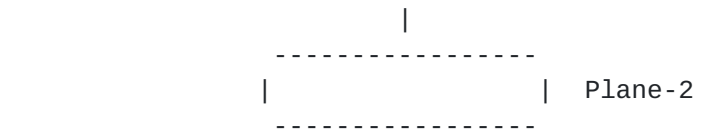

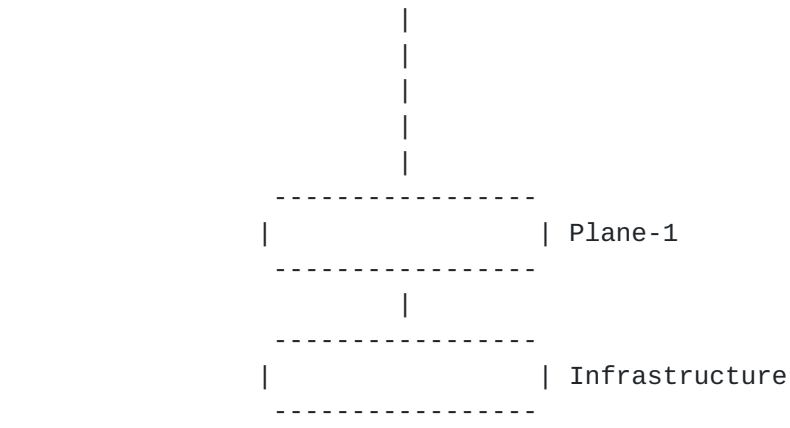

Fig: Logical representation of Layer 2 NBVPN

Within each plane, multiple broadcast domains exist based on the VLAN usage by the customers. A given VLAN in a given Layer 2 NBVPN plane may span across a set of given end points.

4. Requirements for Network Based Layer 2 VPN

In this section we present key requirements in Network Based Layer 2 VPN services.

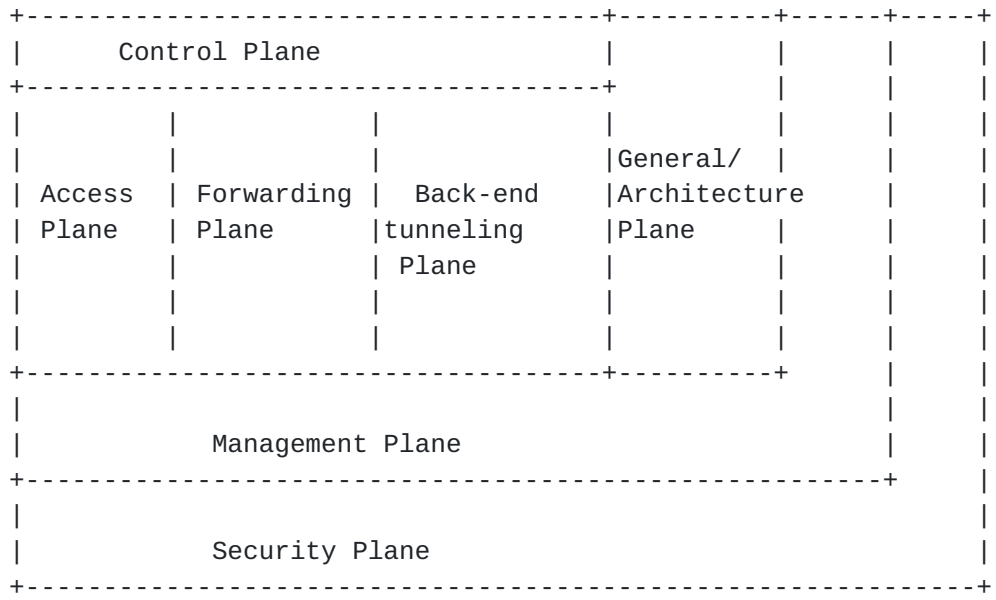

Fig: Placement of Layer 2 NBVPN requirements.

The above diagram depicts placement of key requirements of Layer 2 VPN services. In this document we identify six major areas. Underlying Layer 2 NBVPN requirements can be classified into one or more of these areas. Classification of underlying requirements into these major areas allows readers to co-relate stated

requirements into different components of the product or service infrastructure, easily. On the other hand it allows readers to properly evaluate requirements, technically.

General/Architecture Requirement Plane (G)

The requirements in this plane represent global or general requirements that covers all other planes or requirements such as concepts that specify the entire Layer 2 NBVPN infrastructure.

Control Plane (C)

Requirements in this plane represent signaling and control functionality.

Access Plane (A)

Requirements in this plane specify the CPE-PE Access.

Forwarding Plane (F)

Requirements in this plane specify the Forwarding behavior and architecture required to implement Layer 2 NBVPN.

Back-end tunneling Plane (B)

Requirements in this plane specify the requirements for back-end tunneling protocols to properly implement Layer 2 NBVPN services.

Management Plane (M)

Management plane specifies management requirements for Layer 2 NBVPN services. Management plane covers all other planes in the Layer 2 NBVPN infrastructure.

Security Plane (S)

Security plane specifies security aspects for Layer 2 VPN services. Security plane covers all other planes in the Layer 2 VPN infrastructure.

NOTE: Some requirements may qualify for more than one plane such requirements are identified by specifying the requirement scope within (x). As an example a requirement that qualify for both Access and Control plane are denoted with a qualifier of (A,C).

5. Control Plane Requirements

5.1 Customer End-point and VLAN discovery (C)

A given customer may have several VLANs spanning multiple sets of sites. The L2 NBVPN infrastructure MUST NOT require all VLANs span the same set of sites.

For the purpose of packet forwarding, address learning and tunnel construction, the infrastructure MUST employ methods to discover the end points that belong to different Layer 2 VLAN domains within a single L2 NBVPN plane.

In simpler small-scale implementations, providers may configure the remote end-points and VLAN span manually. Though such methods works for smaller deployments, ability to discover and bind remote sites is essential in larger deployments. Hence it is required to have methods in place to distribute VLAN and domain information.

Automatic VLAN distribution methods are required to have ability to limit the scope of distribution of customers VLAN information.

5.1.1 End-point-VLAN policies

End-point-VLAN policies are broadly classified in to two categories; Announce Polices and Bind polices

Bind Policies

Bind polices specifies binding of incoming advertisements to the local representation. The End-point-VLAN policies are required to be flexible enough to cover different requirements.

Announce Policies

Announce policies apply when advertising L2NBVPN information. The announce policies identify which reachability information need to be advertised and the scope of the advertisement.

5.2 End-point VLAN tag translation

The infrastructure MAY support translation of customersÆ VLAN tags. Such service simplifies connectivity of sites that want to keep their tag assignments or sites that belong to different administrative entities. In the latter case, the connectivity is sometimes referred to as L2 extranet.

5.3 Group (Multicast) Address Registration Services (C)

Layer 2 NBVPN providers MAY provide Group Address Registration service. The Group Address Registration service facilitates customers to define applicable flooding scopes for multicast addresses. Some available alternatives for Group Address Registration services are, IGMP snooping, Generic Multicast Address Registration Protocol (GMRP).

In the absence of Group Address Registration services either Provider has to manually configure the scope of the multicast addresses or flood the multicast addresses to all endpoints.

5.4 Support for Layer 2 control protocols such as GVRP and STP (C)

The Layer 2 NBVPN in theory emulates a LAN. Hence like physical LAN Layer 2 NBVPN SHOULD be transparent to Layer 2 Control protocols such as STP. However, optionally, Laver 2 NBVPN PE device MAY participate in GVRP to create forwarding scope dynamically, within a given customer domain. (This may be viewed same as IGP in VPRN)

GVRP is used as a dynamic VLAN registration protocol. Traditional Layer 2 devices create dynamic Layer 2 forwarding databases based on the GVRP registrations from CPE. Layer 2 VPN PE devices are required to maintain separate VFI instance for each customer's VLAN. In addition, VFI defines binding between the Local VFI and remote end points.

Hence, PE devices that support dynamic VLAN registration via GVRP MUST have capabilities to create Layer 2 VFI instances based on incoming GVRP registrations. Also PE devices MUST have capabilities to bind GVRP registrations to Layer 2 VFI.

Spanning Tree Protocol (STP) is a widely used Layer 2 protocol to prevent loops. Some customers may wish to implement the redundancy using STP rather than purchasing a redundant Layer 2 VPN service from the provider. The multi-point-to-multi-point Layer 2 VPN services in theory emulates legacy LAN topology. Unlike GVRP, the PE devices are not required to participate in STP processing. It is sufficient to transparently pass incoming BPDU to remote sites.

Consider the following fully meshed Layer 2 VPN deployment and the logical representation of the deployment. It is clear that ability of all end devices to receive BPDU is sufficient for proper operation of STP.

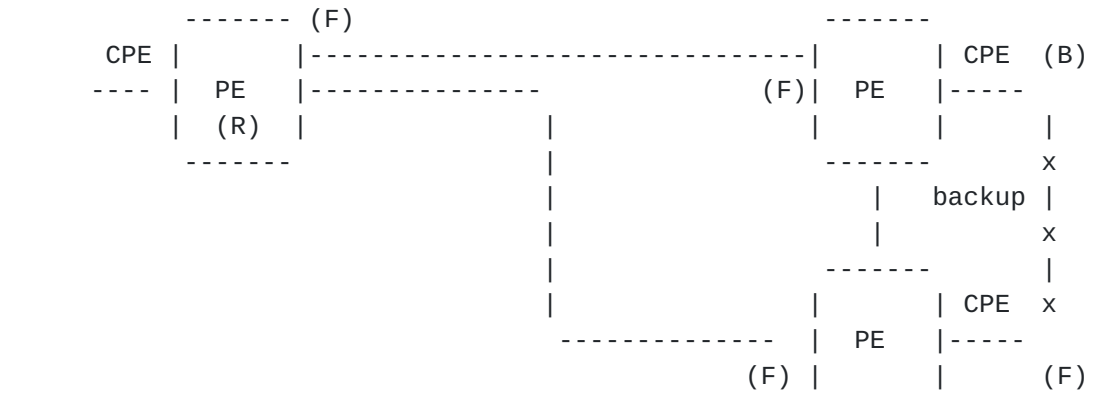

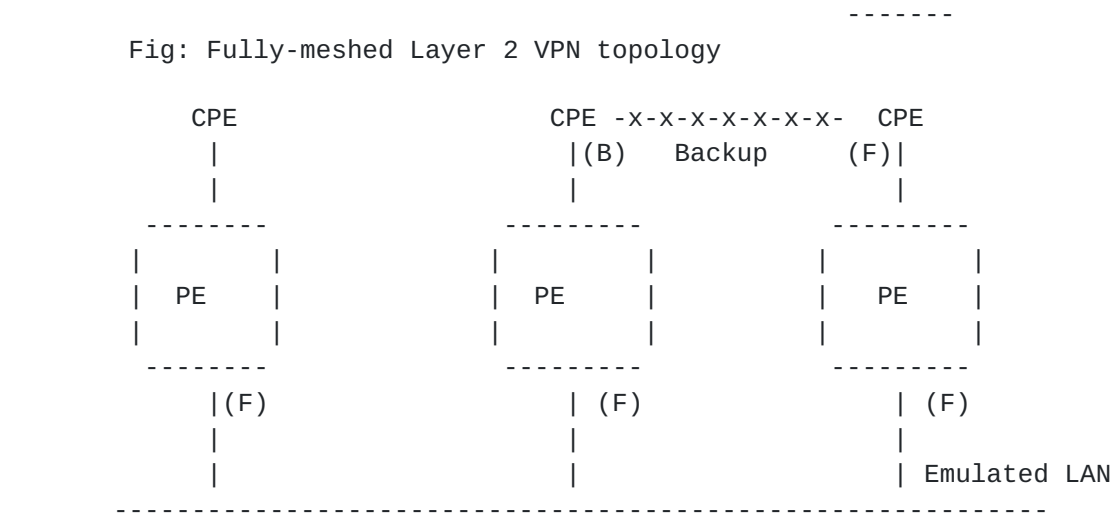

Fig: Fully-meshed Layer 2 VPN topology

Fig: Logical Representation of Fully Meshed Layer 2 VPN

(R) - Root Bridge

(F) - Forwarding Port
(B) - Block Port

5.5 Recovery and Restoration (C,B,F)

Redundancy of Layer 2 topologies are mostly implemented using Spanning Tree Protocols. Any major change in Layer 2 topology requires STP to re-converge. Re-convergence of STP is in the order of 10's of seconds. It is sufficient to have the same degree of re-convergence capabilities in Layer 2 NBVPN.

Alternatively Layer 2 NBVPN MAY provide redundant paths to assure high availability. The reaction to failures should result in an attempt to restore the service using alternative paths. It is desirable to make the restoration times as small as possible.

The restoration time SHOULD be less than the failure detection time of L3 service running over the same VLAN.

When implementing fast-reroute methods in the Back-end, the architecture MUST not violate requirements specified in this document. As an example a fail-over may not cause any packet mis-ordering.

5.6 Dynamic Service Signaling (C)

A provider MAY offer an in-band method for selecting services from the list specified in the SLA.

6. Forwarding Plane Requirements

6.1 Layer 2 Virtual Forwarding Instance (F)

Traditional Layer 2 devices maintain a separate Layer 2 forwarding database per VLAN instance. All forwarding decision for the VLAN is made using the Layer 2 forwarding database of that VLAN.Layer 2 NBVPN providers are required separate forwarding decision between customers. Within a given customer domain, forwarding decisions are again separated based on VLAN usage. Thus Layer 2 NBVPN requires two-level Layer 2 Virtual Forwarding Instance.

Layer 2 NBVPN, Provider Edge devices MUST have capabilities to classify incoming customer traffic into the appropriate customer domain and identification of the proper Layer 2 Virtual Forwarding instance based on Customer domain and VLAN identifier.

6.2 MAC address learning (F,C)

The Layer2 NBVPN is intended to be transparent to L2 customer networks.

Traditional Layer 2 devices learn MAC addresses in the context of a VLAN. Such devices learn MAC addresses against a physical/logical port.

In Layer 2 NBVPN MAC address learning takes place in the context of Layer 2 Virtual Forwarding Instance. That is in the context of Customer domain and a VLAN. The reachability of MAC addresses may be either via a directly attached physical port or a virtual port where remote MAC address may be reached. The virtual port may represent a physical port, a tunnel or some other logical representation. Hence, Layer 2 NBVPN PE devices are required to support virtual port concepts.

It is expected that the Layer 2 NBVPN is able to derive all topology and forwarding information from packets originating at end user sites. The service MAY implement a MAC addresses learning mechanism for this purpose.

The Layer 2 NBVPN is intended to be transparent to Layer 2 end user networks. It MUST NOT require any special packet processing by the end users before sending packets to the providerÆs network.

6.3 Unknown, Multicast and Broadcast forwarding (F)

Unknown, Broadcast and Multicast packets are required to be forwarded to all endpoints of the given customers given VLAN. The scope of Broadcast, Multicast and Unknown is called Broadcast domain. For a given customer there are multiple Broadcast domains, one for each VLAN.

The PE device is required to have capabilities to forward traffic to multiple end points within a given Broadcast domain.

The PE device is required to separate traffic between different broadcast domains.

Each Layer 2 Virtual Forwarding Instance (VFI) is required to have flooding capabilities in the scope of its broadcast domain to facilitate proper forwarding of Broadcast, Multicast and Unknown traffic.
The Layer2 NBVPN MUST be aware of the existence and the designated roles of special MAC addresses such as Multicast and Broadcast addresses. The Layer 2 NBVPN MUST forward these packets according to their intended functional meaning and scope.

If the Layer 2 NBVPN relies on MAC learning for its operations, it MUST assure proper forwarding of packets with MAC addresses that have not been learned.

6.4 Multi site spanning broadcast domains (F)

Broadcast domain is defined as the flooding scope of the Layer 2 network. Flooding scope of a Layer 2 network is the scope of end points that are included in Multicast, Broadcast and Unknown traffic forwarding. Each broadcast domain is defined in the context of a give customer and the scope of a given VLAN. In other words each Layer 2 VFI MUST contain its own Broadcast domain or a flooding scope.

6.5 Scope of Unknown MAC Addresses (F,C)

In general, unknown MAC addresses are flooded to all end point of the Layer 2 NBVPN. In order to conserve bandwidth and allow security, some customers MAY require Layer 2 NBVPN PE devices provide methods to define scope of unicast MAC addresses.

6.6 Mac address transparency in the core (F,B)

Traditionally, Layer 2 forwarding requires learning of MAC addresses. A given device can only learn a finite set of MAC addresses. If Layer 2 NBVPN is implemented with the requirement to learn MAC addresses in the core; then the number of Layer 2 NBVPN that could be supported in the core will depend on the total number of MAC addresses a device can learn. On the other hand any topology change in the core may require to relearn MAC addresses. In order to keep data flow uninterrupted, during the interval of relearning flooding of data is required. Flooding of customer data in the network core may affect the available bandwidth. The Layer 2 address space of each customer is mutually exclusive from one-another. Hence the devices in the core are required to maintain multiple Layer 2 VFI.

Large scale Layer 2 NBVPN deployments MAY require to implement tunneling methods that interconnect remote sites. Tunneling protocol defines the forwarding in the core. Topology convergence and recovery is part of the tunneling protocol. The customer MAC addresses and VLAN are require to be transparent to the devices in the core. Hence the devices in the core are not require to implement Layer 2 VFI for all the customers in the network. On the other hand the devices in the core may not be Layer 2 devices at all in its configuration. Only the PE devices required having Layer 2 forwarding capabilities in its configuration.

The Layer 2 NBVPN is intended to be transparent to all customers per-VLAN basis. The Layer 2 NBVPN MUST NOT rely on global uniqueness of MAC addresses.

6.7 Minimum MTU (F)

The service MUST support customer frames 1500 bytes long. The service MAY offer support for longer frames.

The service MUST NOT fragment packets. Packets exceeding committed MTU size MUST be discarded.

6.8 Packet re-ordering or duplication (F)

The service MUST preserve the order of packet sent from one end point to another end point within given VLAN with given committed topology.

The service MUST NOT duplicate packets.

6.9 Support for MAC Services (G,F)

Layer 2 NBVPN are required to provide MAC service as specified in IEEE 802.1D specification [4] Section 6. Some of the MAC services defined in [4] are directly applicable to Layer 2 NBVPN. Some other MAC services require changes in order to be meaningful in Layer 2 NBVPN applications. In this section each of the MAC service requirements specified in [4] are revisited. All Layer 2 NBVPN devices are required to support MAC services presented in this section. Compliance with this section facilitates proper operation of 802.1 LAN and seamless integration of Layer 2 NBVPN with bridged Local Area Networks.

A MAC service in the context of Layer 2 NBVPN is defined as; Transfer of user data between source and a destination end stations via the service access points using the information specified in the Layer 2 NBVPN VFI.

6.9.1 Preservation of MAC services

MAC services offered by LAN's interconnected by Layer 2 NBVPN devices must be similar to MAC services provided in a single LAN. Hence,

1. A Layer 2 NBVPN must not be directly accessed by end stations except for explicit management purposes.

2. All MAC addresses must be unique within a given customer domain and a VLAN, i.e. within VFI.

3. The MAC addresses of end stations must not be restricted by the topology and configuration of the Layer 2 NBVPN.

6.9.2 Quality of service maintenance

The quality of services provided by Layer 2 NBVPN must not be significantly inferior to that of LAN or IEEE 802.1 bridges. Following areas, at minimum, must be considered when evaluating the quality of service maintenance in Layer 2 NBVPN. Section 4.14 above present quality of service requirements that may be specific to Layer 2 NBVPN.

. Service availability

. Frame loss
. Frame misordering
. Frame duplication
. The transit delay experienced by frames
. Frame lifetime
. The undetected frame error rate
. MTU size support
. User priority
. Throughput
. Scope of Layer 2 forwarding

Service availability

Service availability is defined as a fraction of some total time during which the Layer 2 NBVPN services are available.

Automatic reconfiguration and other methods may increase Service availability. During service failures or automatic reconfiguration, Layer 2 NBVPN devices may deny access and discard frames to preserve forwarding aspects and other requirements of Layer 2 NBVPN.

Frame loss

Layer 2 NBVPN devices do not guarantee frame delivery. Frames may be discarded due to several reasons. PE device must provide statistics on frame loss that occur within the PE. Layer NBVPN PE devices are often connected via some tunneling method. Unlike 802.1 bridges the Layer 2 NBVPN PE devices may be separated by several hops. Hence Frame loss can occur some where in the

transit. Collection of such loss of frames is OPTIONAL and some implementation may not provide them.

Frame misordering

Layer 2 NBVPN must not permit misordering of frames of a given customer domain, for a given VLAN with a given user priority.

Frame duplication

The Layer 2 NBVPN must not permit frame duplication. If there exists multiple paths between PE devices forwarding pollicies of the local PE must ensure that only a single copy of the frame is transmitted to the remote PE device. Definition of frame forwarding polices are beyond the scope of this document. Section 7.7 of [4] provides detail explanation of frame forwarding in IEEE

bridges. These forwarding policies may be extended to Layer 2 NBVPN.

Transit delay

It is difficult to measure the total transit time taken from end station to end station. However, transit time of Layer NBVPN PE devices may be measured in terms of total time taken for reception, classification and transmission of the frame. Transit delay of Layer 2 NBVPN MUST be in compliance with 802.1D specification [4]. The transit delay introduced by Layer 2 NBVPN must not be arbitrarily large.

Frame lifetime

In order to ensure proper operation of upper layer protocols, Layer 2 NBVPN devices must specify an upper bound to the transit delay specified above.

Traditional bridges consider transmit time as the time taken to transmit the packet in to the wire. However, due to multi-hop separation between PE devices, the transmission time of a frame at PE device must take in to account the delay introduced by the transit network (tunnels).

As per 801.1D[4] specification lower bound for maximum frame life time is 1.0 seconds and upper bound for maximum frame life time is 4.0 seconds.

Layer 2 NBVPN MAY specify different set of frame life time parameters.

MTU Size.

Layer 2 NBVPN PE devices and the network must be capable of supporting the largest MTU size that customer is required to transmit. In another words, customer is capable of transmitting

the smallest of all MTU sizes supported by the Layer 2 NBVPN devices and the network.

Priority

Layer 2 NBVPN PE devices maps priority of incoming user traffic in to different internal traffic classes. In the egress to the back end (tunnels) PE device may map these traffic classes in to different planes of the NBVPN. Each of these planes MAY be defined

with different traffic engineering parameters. Mapping of user priorities to egress traffic planes is entirely a local policy and beyond the scope of this publication.

Throughput

Based on the service level agreement and bandwidth provisioning, the total throughput provided by Layer 2 NBVPN network may be significantly less than that of the local LAN. Hence PE device may discard frames that exceed the provisioned bandwidth.

Scope of Layer 2 forwarding

In order to optimize forwarding of MAC addresses, Layer 2 NBVPN PE devices may provide methods to specify the scope of a MAC address. The scope of the MAC address is defined as potential end-point where such MAC address can appear or where potential stations that wish to receive frames are located.

7. General/Architecture Requirements

7.1 Layer 2 Domain representation and VLAN allocation (G)

Layer 2 Network based VPN infrastructure MUST distinguish different customer domains. Each of these customer domains MUST appear as a L2 broadcast domain network behaving like a LAN (Local Area Network).

Configurations within the provider MUST not constrain customers ability to configure VLANs or any other Layer 2 parameters.

As explained in above section 3.1, Layer 2 NBVPN appears as a logical plane in the service infrastructure. The Layer 2 NBVPN plane MAY span across multiple service provider infrastructures. Layer 2 NBVPN Domain Identifier (L2NBVPNDoamin) uniquely identify a given Layer 2 NBVPN. The L2NBVPNDoamin identification is hence required to be unique within all service providers.

It is suggested that L2NBVPNDoamin {AS:L2NBVPNdomain-ID} MAY be derived as a combination of AS (Autonomous System) Number and a unique number that represent the customer's domain (L2NBVPNdomain-ID) within the AS.

Within a given AS L2NBVPNdomain-ID MUST be unique for all the end

points in a given L2NBVPN plane.

7.2 Customer end point inter connection (G,B)

Layer 2 networks require a symmetric connectivity between devices. In other words each device MUST have connectivity to all other devices.

Hub and spoke or point-to-multi-point is a popular deployment model in legacy Frame Relay networks. When providing routed services, inter site traffic is routed via the hub site.

When providing Layer 2 Connectivity, Hub is required to implement complex forwarding policies to achieve symmetric connectivity. On the other hand hub and spoke model may be used for deployment that does not require symmetric connectivity. Example of such deployment is - remote sites that requires connectivity to the data center and does not require connectivity between remote sites.

Many-to-many (either logical or physical) is required to be deployed for Layer 2 services that require symmetric connectivity. Enterprise networks that extend the private LAN using Layer2 NBVPN require many-to-many connectivity.

7.3 Class of Service Model (G)

The VLAN service SHOULD define a graded selection of classes of traffic. This includes, but is not limited to

- o range of priorities
- o best effort vs. guaranteed effort
- o range of minimum delay characteristics

7.4 L3, and higher, service access point (G)

The VLAN SHOULD a allow for a Provider based Service Access Point for orderly injection of L3 or higher services to the customerÆs VLAN.

As a value added service, L2 NBVPN may provide access to other services such as, IP gateways, Storage networks, Content delivery etc..

8. Management Plane Requirements

Layer 2 NBVPN infrastructure MUST have capabilities to manage and monitor different components.

Layer 2 NBVPN infrastructure MAY provide methods between CPE-PE devices for self provisioning and management.

8.1 Graceful reconfiguration (M)

In cases where the provider knows a priori about impending fault, the network SHOULD be reconfigured without a loss, duplication, or re-ordering of customer packets. This situation typically arises with planned network upgrades, or scheduled maintenance activities.

9. Security Plane Requirements

All layer 2 NBVPN devices MUST require to have ability to isolate traffic for each Layer 2 VFI. Ingress classification MUST be well defined such that there exists one and only one VFI for each VLAN of each customer domain. In other words, provider MUST provide traffic separation between different customers and between different VLANs of the same customer. The traffic separation MUST prevent leaking of the traffic in or out of VLANs in a normally functioning Layer 2 NBVPN.

Optionally; PE devices may offer data encryption and authentication services. Protection against denial of service attacks. Protection against bandwidth and connection hijacking. Ability to provide some of these optional security requirements may depend on tunneling protocol used in the core.

In addition, each key requirement plane may specify itÆs own security requirements. Such requirements are discussed under each of those sections. As an example, Access plane MAY have very specific set of requirements that are not related to security requirements of Management plane.

9.1 Immunity from malformed customer traffic (S)

The providerÆs infrastructure MUST NOT be compromised by malformed, or maliciously altered, customer traffic. This includes, but is not limited to, duplicate or invalid MAC addresses, short packets, long packets, etc.

10. Back-End Tunneling Requirements (B)

The main difference of Layer 2 NBVPN with traditional bridging is use of back-end tunneling.

The tunneling protocol used MUST be capable of providing a logical connectivity between PE device in the NBVPN domain.

Tunneling protocol MUST be capable of providing required connectivity. If the chosen Layer 2 NBVPN service requires many-

to-many connectivity then PE device SHOULD require to support the many-to-many connectivity in the back-end. If the tunneling protocol is not capable of providing the required connectivity then PE device MUST have capability to support appropriate forwarding behavior.

The back-end tunneling protocol SHALL not affect the scalability of Layer 2 NBVPN service.

The back-end tunneling protocol MUST not require customers to use special equipment or alteration to the normal network topology/protocols or other..

There are several choices for back-end tunneling protocols. The choice of tunneling protocol depends on various factors. Popular choices are MPLS, IP-tunneling. There are proposals to use variations of 802.1Q to achieve required tunneling behavior, discussion of such tunneling protocols are beyond the scope of this document.

11. Access Requirements:

Access requirements broadly specifies the requirements for CPE to PE connection.

CPE-PE access connection MAY require spanning other providers or public networks.

CPE-PE access MAY require to support different technologies such as Ethernet, ATM (DSL), Frame Relay etc.

PE device that provide Layer 2 NBVPN services MAY require to support multiple logical access connections over a single physical access connection. The logical connection MAY be either in the Form IEEE802.1Q VLAN Tagged format or MPLS Label Switched Path or ATM connection or some other technology. When providing logical access, PE devices MUST honor QoS and other properties of the logical connection. As an example when logical connections are provided using IEEE 802.1Q VLAN Tagging, PE device MUST provide

capabilities honor QoS parameters of the logical connection.

PE device MUST have capabilities to translate properties of logical connections to corresponding properties of the back-end tunneling. Translation of these properties are a local policy and beyond the scope of this document.

PE devices SHOULD not restrict the physical connectivity of the access to a single technology such as optical Ethernet (10G). PE device SHOULD have capabilities to provide different physical access technologies.

Access speeds

The Layer 2 NBVPN service shall allow for sites specific access speeds, i.e. it shall be possible to interconnect sites operating on and accessing to the backbone on any of the different standard access speeds. Neither the access speed to the backbone or the bandwidth on the LANs on the sites needs to be identical.

11.1 Security Requirements for Access Plane

The PE device SHOULD support all the security requirements mandated by the technology used to specify access connections.

The PE device MAY provide additional security features based on the SLA (Service Level Agreement).

The PE device MAY require to provide additional security services based on the access is a direct connection or over a public network.

The security methods used in the access plane MAY not depend on the security mechanisms in the back-end.

12. References

1 Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

2 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

3 Senevirathne, T., and et.al., Framework For Virtual Metropolitan Internetworks (VMI), Work in Progress, February 2001

4 ANSI/IEEE Std 802.1D, Media Access Control (MAC) Bridges, Internatyional Electrotechnical Commission, 1998.

13. Acknowledgments

Ongoing discussion in PPVPN has influenced work presented in this document.

14. Author's Addresses

Tissa Senevirathne
Force10 Networks
1475 McCarthy Blvd

Milpitas, CA
Phone: 408-965-5103
Email: tissa@force10networks.com

Waldemar Augustyn
Nortel Networks
600 Technology Park
Billerica, MA 01821
Phone: 978 288 4993
Email: waldemar@nortelnetworks.com

Loa Andersson
Utfors Bredband AB
Box 525
SE-169 29 Solna
Sweden
Phone: +46 8 5270 5038
Email: loa.Andersson@utfors.se

Tove Madse
Utfors Bredband AB
Box 525
SE-169 29 Solna
Sweden
Phone: +46 8 5270 5038
Email: tove.madsen@utfors.se

Pascal Menezes
TeraBeam Networks
2300 Seventh Ave Seattle, WA 98121

Email: Pascal.Menezes@Terabeam.com

Appendix A: Acronyms and Abbreviations

NBVPN û Network Based Virtual Private Networks

PE û Provider Edge Device

CE û Customer Edge Device

VFI û Virtual Forwarding Instance

GVRP û Generic VLAN registration Protocol

SLA û Service Level Agreement

STP û Spanning Tree Protocol

VLAN û Virtual Local Area Network. VLAN in this document refers to VLAN identifiers assigned by customers

Full Copyright Statement