

PPVPN Working Group
Internet Draft
Document: [draft-augustyn-vpls-requirements-01.txt](#)
Category: Informational

Waldemar Augustyn

Giles Heron
PacketExchange Ltd

February 2002
Expires: August 2002

Vach Kompella
TiMetra Networks

Marc Lasserre
Riverstone Networks

Pascal Menezes
Terabeam

Hamid Ould-Brahim
Nortel Networks

Tissa Senevirathne
Force10 Networks

Requirements for Virtual Private LAN Services (VPLS)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

For potential updates to the above required-text see: <http://www.ietf.org/ietf/1id-guidelines.txt>

1 Abstract

This draft describes service requirements related to emulating a virtual private LAN segment over an IP or MPLS network infrastructure. The service is called VPLS. It is a class of Provider Provisioned Virtual Private Network [2].

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](#) [3].

3 Definitions

3.1 VPLS

A Virtual Private LAN Service or, when referring to a virtual LAN segment offered by that service, a Virtual Private LAN Segment.

A VPLS service allows the connection of multiple sites in a single bridged domain over a provider managed IP or MPLS network. All customer sites in the VPLS appear to be on the same LAN regardless of their location.

3.2 VPLS System

A collection of communication equipment, related protocols, and configuration elements that implements VPLS Services.

3.3 VPLS Domain

A Layer 2 VPN that is composed of a community of interest of L2 MAC addresses and VLANs. Each VPLS Domain MAY have multiple VLANs in it.

3.4 VLAN

A customer VLAN identification using some scheme such as IEEE 802.1Q tags, port configuration or any other means. A VPLS service can be extended to recognize customer VLANs as specified in 6.1 .

3.5 VLAN Flooding Scope (VLAN Broadcast Domain)

The scope of flooding for a given VLAN.

In a VPLS service, a VLAN flooding scope MUST be identical to the flooding scope of the VPLS it is part of. If a VPLS service is extended to recognize customer VLANs, the VLAN flooding scope SHOULD be limited to the broadcast domain of each recognized VLAN.

3.6 VFI

Virtual Forwarding Instance.

A virtual layer 2 forwarding entity that is closed to a VPLS Domain membership. VFI forwarding is based on MAC addresses, VLAN tags, policies, topologies, filters, QoS parameters, and other relevant information on a per VPLS basis.

4 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 architectures significantly increased the complexity of network topologies and required highly skilled personnel to manage and maintain the network.

One solution to the above has been for service providers to deploy VPLS services (often known in this context as "Transparent LAN" services, or TLS.) These have typically been offered over an ATM infrastructure. The TLS service is provisioned using a mesh of ATM PVCs between locations. While this reduces 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.

The aim of this effort is to develop a VPLS service that is simpler to manage and more scalable. By using the service provider's existing MPLS or IP backbone the overhead of running separate network backbones can be avoided, while by using an MPLS label stack or by tunneling traffic in IP the complexity (and scaling problems) of building a separate mesh of VCs for each TLS can be avoided.

VPLS can be contrasted with two other service models. One is the Virtual Private Routed Network (VPRN) [4], in which the service provider participates in the customer's routing protocol. The second model is a set of point-to-point circuits (L2VPN) [5].

In the VPRN model the PE devices are required to maintain the private IP routing tables of the customers. These routing tables may consist of hundreds, or thousands, or routes each. Requiring each PE device to hold all routes for each VPN present at that PE has a major impact on the cost of these devices and the scalability

of the service. Furthermore, PE devices are required to support different routing protocols such as BGP, OSPF and RIP in order to learn routing updates from the CE. These PE devices will have to maintain a separate routing instance per VPRN. This support increases the complexity of the service, and also impacts the cost of the PE devices.

In contrast, in a given LAN segment there is a reasonably small set of MAC devices with a limited number of MAC addresses to learn and manage. Layer 2 VPN services do not require any additional routing protocol support between the CE and the PE devices, and in fact they are transparent to the customer's choice of routing protocol. Layer 2 VPN services also benefit from being transparent to higher layer protocols, so the same technology can transport, for example, IPv4, IPv6, MPLS and also legacy protocols such as IPX and OSI.

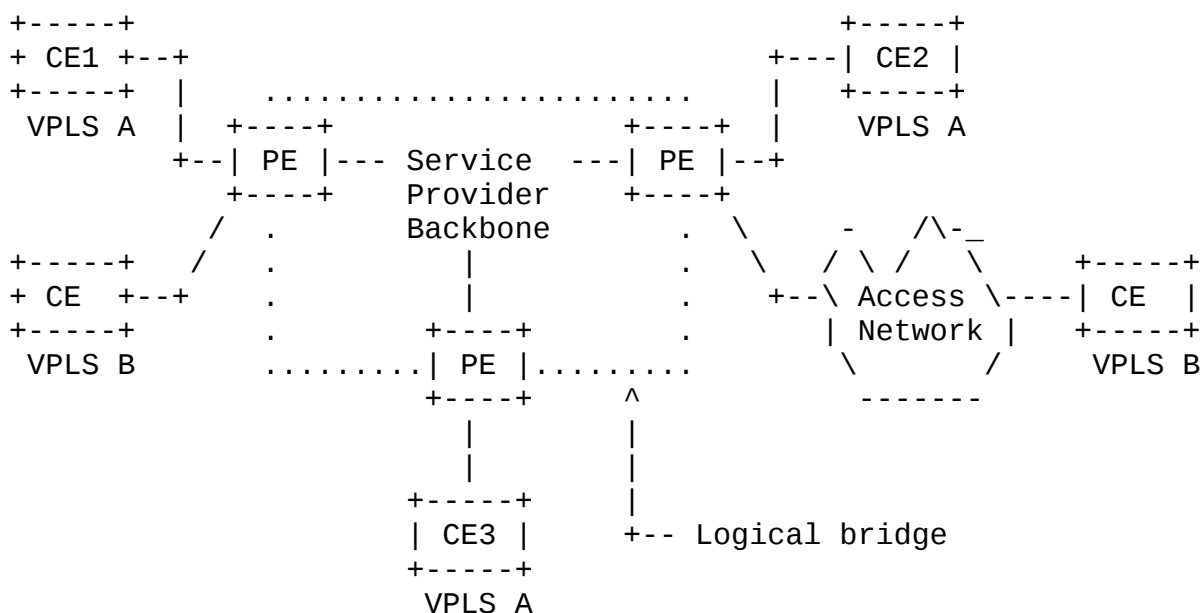
In the L2VPN model the CE devices are required to support a number of virtual circuits to other CE devices, and to form routing adjacencies over these virtual circuits. In addition, when building a large L2VPN, it becomes impractical to maintain a full mesh of virtual circuits, and the provisioned topology has to be designed per L2VPN. In contrast VPLS emulates a flat LAN segment and has learning and switching capabilities, and thus only requires the CE to support a single connection to the VPLS. Routing adjacencies can be formed with a designated router for the VPLS, or in a full mesh of all the CE devices on the VPLS (depending on the routing protocol used by the customer.)

The VPLS model, while offering significant benefits for both customers and service providers, retains all the quintessential characteristics of L2 networks including their well known limitations such as the maximum practical number of hosts on a single segment and the inability to assign different metrics for connectivity between different hosts on the same segment. A likely application of this model is to connect a few sites with only a single customer router at each site, or a small number of customer hosts, connected to the VPLS at each site.

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

5 VPLS Reference Model

The following diagram shows a VPLS reference model where PE devices that are VPLS-capable provide a logical interconnect such that CE devices belonging to a specific VPLS appear to be connected by a single logical Ethernet bridge. A VPLS can contain a single VLAN or multiple, tagged VLANs.



Separate L2 broadcast domains are maintained on a per VPLS basis by PE devices. Such domains are then mapped onto tunnels in the service provider network. These tunnels can either be specific to a VPLS (e.g. as with IP) or shared among several VPLSs (e.g. as with MPLS tunnel LSPs). In the above diagram, the top PE routers maintain separate forwarding instances for VPLS A and VPLS B.

The CE-to-PE links can either be direct physical links, e.g. 100BaseTX, or logical links, e.g. ATM PVC, T1/E1 TDM, or [RFC1490](#)-encapsulated link, over which bridged Ethernet traffic is carried.

The PE-to-PE links carry tunneled Ethernet frames using different tunneling technologies (e.g., GRE, IPsec, MPLS.).

Each PE device learns remote MAC addresses, and is responsible for proper forwarding of the customer traffic to the appropriate end nodes. It is responsible for guaranteeing each VPLS topology is loop free.

[6](#) VPLS General Requirements

[6.1](#) Layer 2 Domain representation and VLAN allocation

A VPLS system 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). These domains are referred to as VPLS domains.

A VPLS Domain MAY span multiple service providers. Each VPLS domain MUST carry a unique identification within a VPLS system. It is RECOMMENDED that VPLS identification be globally unique.

Each VPLS Domain MUST be capable of learning and forwarding based on MAC addresses thus emulating an Ethernet virtual switch to the customer CE devices attached to PEs.

A VPLS system MAY recognize customer VLAN identification. In that case, a VLAN MUST be recognized in the context of the VPLS it is part of. If customer VLANs are recognized, separate VLAN broadcast domains SHOULD be maintained.

A provider's implementation of a VPLS system SHOULD not constrain the customer's ability to configure VLAN topologies, tags, 802.1 p-bits, or any other Layer 2 parameters.

6.2 VPLS Topology

The VPLS system MAY be realized using one or more network tunnel topologies to interconnect PEs, ranging from simple point-to-point to distributed hierarchical arrangements. The typical topologies include:

- o point-to-point
- o point-to-multipoint, a.k.a. hub and spoke
- o any-to-any, a.k.a. full mesh
- o mixed, a.k.a. partial mesh
- o hierarchical

Regardless of the topology employed, the service to the customers MUST retain the typical LAN any-to-any connectivity. This requirement does not imply that all traffic characteristics (such as bandwidth, QoS, delay, etc.) be necessarily the same between any two end points.

6.3 Redundancy and Failure Recovery

The VPLS infrastructure SHOULD provide redundant paths to assure high availability. The reaction to failures SHOULD result in an attempt to restore the service using alternative paths.

The intention is to keep the restoration time small. It is RECOMMENDED that the restoration time be less than the time it takes the CE devices to detect a failure in the VPLS.

In cases where the provider knows a priori about impending changes in network topology, the network SHOULD have the capability to reconfigure without a loss, duplication, or re-ordering of customer packets. This situation typically arises with planned network upgrades, or scheduled maintenance activities.

6.4 Policy Constraints

A VPLS system MAY employ policy constraints governing various interconnection attributes for VPLS domains. Typical attributes include:

- o Selection of available network infrastructure
- o QoS services needed
- o Protection services needed
- o Availability of higher level service access points (see 9.11)

Policy attributes SHOULD be advertised via the VPLS system's control plane.

6.5 PE nodes

The PE nodes are the devices in the VPLS system that store information related to customer VPLS domains and employ methods to forward customer traffic based on that information. In this document, the PE nodes are meant in logical sense. In the actual implementations, the PE nodes may be comprised of several physical devices. Conversely, a single physical device may contain more than one PE node.

All forwarding decisions related to customer VPLS traffic MUST be made by PE nodes. This requirement prohibits any other network components from altering decisions made by PE nodes.

6.6 PE-PE Interconnection and Tunneling

A VPLS system MUST provide for connectivity between each pair of PE nodes. The connectivity is referred to as transport tunneling or simply tunneling.

There are several choices for implementing transport tunnels. Some popular choices include MPLS, IP in IP tunnels, variations of 802.1Q, etc. Regardless of the choice, the existence of the tunnels and their operations MUST be transparent to the customers.

6.7 PE-CE Interconnection and Profiles

A VPLS system MUST provide for connectivity between PE nodes and CE nodes. That connectivity is referred to as CE access connection, access connection, or simply access. Access connections MAY span networks of other providers or public networks.

There are several choices for implementing access. Some popular choices include Ethernet, ATM (DSL), Frame Relay, MPLS-based virtual circuits etc. Regardless of the choice, the access connection MUST use Ethernet frames as the Service Protocol Data Unit (SPDU).

A CE access connection MUST be bi-directional in nature.

PE devices MAY support multiple CE access connections on a single physical interface. In such cases, PE devices MUST NOT rely on customer controlled parameters, such as VLAN tags etc., for distinguishing between different access connections.

A CE access connection, whether direct or virtual, MUST maintain all committed characteristics of the customer traffic, such as QoS, priorities etc.

The characteristics of a CE access connection are only applicable to that connection.

7 Control Plane Requirements

7.1 Provider Edge Signaling

The control protocols SHOULD provide methods for signaling between PEs. The signaling SHOULD inform of membership, tunneling information, and other relevant parameters.

The infrastructure MAY employ manual configuration methods to provide this type of information.

The infrastructure SHOULD use policies to scope the membership and reachability advertisements for a particular VPLS.

7.2 VPLS Membership Discovery

The control protocols SHOULD provide methods to discover the PEs which connect CEs that form a VPLS.

7.3 Support for Layer 2 control protocols

A VPLS system MUST ensure that loops are prevented. This can be accomplished through a use of control protocols such as Spanning Tree or similar, or through other means such as full mesh topologies and appropriate forwarding rules.

The VPLS system's control protocols SHOULD allow transparent operation of Layer 2 control protocols employed by customers.

A VPLS system's control protocols MAY use indications from customer STP to improve the operation of a VPLS.

7.4 Scaling Requirements

Control plane traffic will increase correspondingly, as a VPLS scales in membership. The rate of growth of control plane traffic SHOULD be linear.

Control plane traffic will increase correspondingly, as the number of supported VPLS segments increases. The rate of growth of control plane traffic SHOULD be linear.

The use of control plane resources will increase correspondingly, as the number of hosts connected to a VPLS increases. The rate of growth of the demand for control process resources SHOULD be linear. The control plane MAY offer means for enforcing a limit on the number of customer hosts attached to a VPLS.

8 Data Plane Requirements

8.1 Transparency

VPLS service is intended to be transparent to Layer 2 customer networks. It MUST NOT require any special packet processing by the end users before sending packets to the provider's network.

8.2 Broadcast Domain

The Broadcast Domain is defined as the flooding scope of a Layer 2 network. In a VPLS system, a separate Broadcast Domain MUST be maintained for each VPLS.

In addition to VPLS Broadcast Domains, a VPLS system MAY recognize customer VLAN Broadcast Domains. In that case, the system SHOULD maintain a separate VLAN Broadcast Domain for each customer VLAN. A VLAN Broadcast Domain MUST be a subset of the owning VPLS Broadcast Domain.

8.3 Layer 2 Virtual Forwarding Instance

VPLS Provider Edge devices MUST maintain a separate Virtual Forwarding Instance (VFI) per VPLS. Each VFI MUST have capabilities to forward traffic based on customer's traffic parameters such as MAC addresses, VLAN tags (if supported), etc. as well as local policies.

Each VFI MUST have flooding capabilities for its Broadcast Domain to facilitate proper forwarding of Broadcast, Multicast and Unknown Unicast customer traffic.

VPLS Provider Edge devices MUST have capabilities to classify incoming customer traffic into the appropriate VFI.

8.4 MAC address learning

A VPLS service SHOULD derive all topology and forwarding information from packets originating at customer sites. Typically, MAC addresses learning mechanisms are used for this purpose.

In a VPLS system, MAC address learning MUST take place on a per Virtual Forwarding Instance (VFI) basis, i.e. in the context of a VPLS and, if supported, in the context of VLANs therein.

8.5 Unicast, Unknown Unicast, Multicast, and Broadcast forwarding

VPLS MUST be aware of the existence and the designated roles of special MAC addresses such as Multicast and Broadcast addresses. VPLS MUST forward these packets according to their intended functional meaning and scope.

Broadcast packets MUST be flooded to all destinations.

Multicast packets MUST be flooded to all destinations. However, a VPLS system MAY employ multicast snooping techniques, in which case multicast packets SHOULD be forwarded only to their intended destinations.

Unicast packets MUST be forwarded to their intended destinations.

Unknown Unicast packets MUST be flooded to all destinations in the flooding scope of the VPLS (or VLAN). If the VPLS service relies on MAC learning for its operations, it MUST assure proper forwarding of packets with MAC addresses that have not been learned. Once destination MAC addresses are learned, unicast packets SHOULD be forwarded only to their intended destinations.

A provider MAY employ a method to limit the scope of flooding of Unknown Unicast packets in cases where a customer desires to conserve its bandwidth or wants to implement certain security policies.

8.6 Multilink Access

The VPLS service SHOULD support multilink access for CE devices. The VPLS service MAY support multihome access for CE devices.

8.7 Minimum MTU

The service MUST support customer frames with payload 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.

The committed minimum MTU size MUST be the same for a given VPLS. If VLANs are supported, all VLANs within a given VPLS MUST inherit the

same MTU size. Different VPLS segments MAY have different committed MTU sizes.

8.8 QoS and packet re-ordering

A VPLS system SHOULD have capabilities to enforce QoS parameters.

The queuing and forwarding policies SHOULD preserve packet order for packets with the same QoS parameters.

The service SHOULD not duplicate packets.

8.9 Support for MAC Services

VPLS are REQUIRED to provide MAC service compliant with IEEE 802.1D specification [6] [Section 6](#). Compliance with this section facilitates proper operation of 802.1 LAN and seamless integration of VPLS with bridged Local Area Networks.

A MAC service in the context of VPLS is defined as the transfer of user data between source and destination end stations via the service access points using the information specified in the VFI.

1. A PE device that provides VPLS MUST NOT be directly accessed by end stations except for explicit management purposes.
2. All MAC addresses MUST be unique within a given broadcast domain.
3. The topology and configuration of the VPLS MUST NOT restrict the MAC addresses of end stations

9 Management and Operations Requirements

9.1 The VPLS system MUST have capabilities to manage and monitor its different components.

9.2 VPLS System Instantiation

It SHOULD be possible to create several disjoint instances of VPLS systems within the same underlying network infrastructures.

9.3 Monitoring

The infrastructure SHOULD monitor all characteristics of the service that are reflected in the customer SLA. This includes but is not limited to bandwidth usage, packet counts, packet drops, service outages, etc.

9.4 VPLS membership

The amount of configuration changes when adding or deleting customer ports to, or from, a given VPLS SHOULD be minimal. The configuration changes of a port to or from a given VPLS SHOULD involve only configuration on the device that this port is connected to.

9.5 End-point VLAN tag translation

If VLANs are recognized, the infrastructure MAY support translation of customers' 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.

9.6 MAC Address Limiting

The VPLS infrastructure MUST be able to limit the number of MAC addresses learned from the customers.

9.7 CE Provisioning

The VPLS MUST require only minimal or no configuration on the CE devices, depending on the CE device that connects into the infrastructure.

9.8 Customer traffic policing

The VPLS service SHOULD provide the ability to police and/or shape customer traffic entering and leaving the VPLS system.

9.9 Dynamic Service Signaling

A Provider MAY offer to Customers an in-band method for selecting services from the list specified in the SLA. A Provider MAY use the same mechanism for reporting statistical data related to the service.

9.10 Class of Service Model

The VPLS service MAY define a graded selection of classes of traffic. These include, but are not limited to

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

9.11 L3, and higher, service access point.

The VPLS service SHOULD allow for a Provider based Service Access Point for orderly injection of L3 or higher services to the customers' VPLS segments.

As a value added service, a Provider MAY offer access to other services such as, IP gateways, storage networks, content delivery etc..

[9.12](#) Testing

The VPLS service SHOULD provide the ability to test and verify operational and maintenance activities on a per VLAN basis.

[10](#) Security Requirements

[10.1](#) Traffic separation

VPLS system MUST provide traffic separation between different VPLS domains as well as between customer VLANs within each VPLS domain if VLANs are supported.

[10.2](#) Provider network protection.

The VPLS system MUST be immune to malformed or maliciously constructed customer traffic. This includes but is not limited to duplicate or invalid MAC addresses, short/long packets, spoofed management packets, spoofed VLAN tags, high volume traffic, etc.

Additionally, VPLS system MUST be immune to misconfigured or maliciously set up customer network topologies. These include customer side loops, backdoor links between sites, etc.

The VPLS infrastructure devices MUST NOT be accessible from the VPLS.

[10.3](#) Value added security services

Value added security services such as encryption and/or authentication of customer packets, certificate management, and similar are OPTIONAL.

Security measures employed by the VPLS system SHOULD NOT restrict implementation of customer based security add-ons.

[11](#) References

1. Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
2. Carugi, et al., "Service requirements for Provider Provisioned Virtual Private Networks ", Work in progress, December 2001.
3. Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
4. Gleeson, et al., "A Framework for IP Based Virtual Private Networks", [RFC 2764](#), February 2000.
5. Kompella et al., "MPLS-based Layer 2 VPNs", Work in progress, June 2001.
6. ANSI/IEEE Std 802.1D, Media Access Control (MAC) Bridges, International Electrotechnical Commission, 1998.6
7. IEEE Standard 802.1Q, "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks", 1998.
8. IEEE Standard 802.1u-2001, "IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks - Amendment 1: Technical and editorial corrections", 2001.
9. IEEE Standard 802.1v-2001, "IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks - Amendment 2: VLAN Classification by Protocol and Port", 2001.

[12](#) Acknowledgments

We would like to acknowledge extensive comments provided by Loa Anderson, Joel Halpern, and Eric Rosen. The authors, also, wish to extend appreciations to their respective employers and various other people who volunteered to review this work and provided feedback.

[13](#) Authors' Addresses

Waldemar Augustyn
Email: waldemar@nxp.com

Giles Heron
PacketExchange Ltd.
The Truman Brewery
91 Brick Lane
London E1 6QL

Augustyn, et al.

Expires August 2002

14

[draft-augustyn-vpls-requirements-01.txt](#) February 2002

United Kingdom
Email: giles@packetexchange.net

Vach Kompella
TiMetra Networks
274 Ferguson Dr.
Mountain View, CA 94043
Email: vkompella@timetra.com

Marc Lasserre
Riverstone Networks
5200 Great America Pkwy
Santa Clara, CA 95054
Phone: 408-878-6500
Email: marc@riverstonenet.com

Pascal Menezes
Terabeam
Phone: 206-686-2001
Email: pascal.menezes@terabeam.com

Hamid Ould-Brahim
Nortel Networks
P.O. Box 3511 Station C
Ottawa ON K1Y 4H7
Canada
Phone: 613-765-3418
Email: hbrahim@nortelnetworks.com

Tissa Senevirathne
Force10 Networks
1440 McCarthy Blvd
Milpitas, CA 95035
Phone: 408-965-5103
Email: tsenevir@hotmail.com

"Copyright (C) The Internet Society (2001). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."