Network Working Group	Y. Kamite	
Internet-Draft	NTT Communications	
Intended status: Informational	F. Jounay	
Expires: May 4, 2009	France Telecom	
	B. Niven-Jenkins	
	BT	
	D. Brungard	
	AT&T	
	L. Jin	
	Nokia Siemens Networks	
	October 31, 2008	

TOC

Framework and Requirements for Virtual Private Multicast Service (VPMS)
draft-kamite-l2vpn-vpms-frmwk-requirements-02.txt

Status of this Memo

Abstract

This document provides a framework and service level requirements for Virtual Private Multicast Service (VPMS). VPMS is defined as a Layer 2 VPN service that provides point-to-multipoint connectivity for a variety of Layer 2 link layers across an IP or MPLS-enabled PSN. This document outlines architectural service models of VPMS and states

generic and high level requirements. This is intended to aid in developing protocols and mechanisms to support VPMS.

Table of Contents

---

1. Introduction

TOC

1.1. Problem Statement

TOC

[RFC4664] (Andersson, L. and E. Rosen, “Framework for Layer 2 Virtual Private Networks (L2VPNs),” September 2006.) describes different types of Provider Provisioned Layer 2 VPNs (L2 PPVPNs, or L2VPNs); Some of them are widely deployed today, such as Virtual Private Wire Service (VPWS) and Virtual Private LAN Service (VPLS). A VPWS is a VPN service that supplies a Layer 2 (L2) point-to-point service. A VPLS is an L2 service that emulates Ethernet LAN service across a Wide Area Network (WAN).

For some use cases described hereafter, there are P2MP (point-to-multipoint) type services for Layer 2 traffic. However, there is no straightforward way to realize them based on the existing L2VPN specifications.

In a VPWS, a SP can set up point-to-point connectivity per a pair of CEs but it is not possible to replicate traffic for point-to-multipoint services in the SP's network side. A SP could build multiple PWs independently and have the CEs replicate traffic over them, but this is not only inconvenient for the customer, it’s a waste of bandwidth resources.

In a VPLS, SPs can natively offer multipoint connectivity across their backbone. Although it is seemingly applicable for point-to- multipoint service as well, there remains extra complexity for SPs to filter unnecessary traffic between irrelevant sites (i.e., from a receiver PE to another receiver PE) because VPLS provides multipoint-to-multipoint connectivity between CEs. Moreover, VPLS's MAC-based learning/ forwarding operation is unnecessary for some scenarios particularly if customers only need simple unidirectional point-to-multipoint service, or if they require non- Ethernet Layer 2 connectivity.

Consequently, there is a real need for a solution that natively provides point-to-multipoint service in L2VPN.

TOC

1.2. Scope of This Document

VPMS is defined as a Layer 2 service that provides point-to- multipoint connectivity for a variety of Layer2 link layers across an IP or MPLS-enabled PSN. VPMS is categorized as a class of provider- provisioned Layer 2 Virtual Private Networks (L2VPN).
This document introduces a new service framework, reference model and functional requirements for VPMS by extending the existing framework [RFC4664] (Andersson, L. and E. Rosen, “Framework for Layer 2 Virtual Private Networks (L2VPNs),” September 2006.) and requirements [RFC4665] (Augustyn, W. and Y. Serbest, “Service Requirements for Layer 2 Provider-Provisioned Virtual Private Networks,” September 2006.) for L2VPNs.
The technical specifications are outside the scope of this document. There is no intent to specify solution-specific details.
This document provides requirements from both the Service Provider's and the Customer's point of view.

---

2. Conventions used in this document

TOC

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.) .

---

3. Terminology

TOC

The content of this document makes use of the terminology defined in [RFC4026] (Andersson, L. and T. Madsen, “Provider Provisioned Virtual Private Network (VPN) Terminology,” March 2005.). For readability purposes, we list some of the terms here in addition to some specific terms used in this document.

3.1. Acronyms

TOC

P2P: Point-to-Point

P2MP: Point-to-Multipoint

PW: Pseudowire

VPMS:
Virtual Private Multicast Service

PE/CE: Provider/Customer Edge

P: Provider Router

AC: Attachment Circuit

PSN: Packet Switched Network

SP: Service Provider

---

4. Use Cases

TOC

4.1. Ethernet Use Case

TOC

For multicast traffic delivery, there is a requirement to deliver a unidirectional P2MP service in addition to the existing P2P service. The demand is growing to provide private (P2MP native Ethernet) services, for various applications such as IP- based delivery of TV broadcasting, content delivery networks, etc. Moreover, many digital audio/video devices (e.g., MPEG-TS, HD-SDI) that supports Ethernet interfaces are becoming available, which will make Ethernet P2MP service more common. Also there are some applications that naturally suited to static transport of VPMS. For example, MPEG-TS/IP/ Ethernet in DVB-H is typically static broadcast without any signaling in the upstream direction. VPMS could be a possible solution to provide these kinds of networking connectivity over PSNs.
Currently VPLS [RFC4761] (Kompella, K. and Y. Rekhter, “Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling,” January 2007.)[RFC4762] (Lasserre, M. and V. Kompella, “Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling,” January 2007.) is able to give P2MP-type replication for Ethernet traffic. Native VPLS already supports this capability via a full mesh of PWs, and an extension to optimize replication is also proposed [I-D.ietf-l2vpn-vpls-mcast] as an additional feature. However, VPLS by nature requires MAC-based learning and forwarding, which might not be needed in some cases by particular users. Generally, video distribution applications use unidirectional P2MP traffic, but may not always require any added complexity of MAC address management. In addition, VPLS is a service that essentially provides any-to-any connectivity

between all CEs in a L2VPN as it emulates a LAN service. However, if only P2MP connectivity is required, the traffic between different receivers is not always needed, and traffic from receiver to sender is not always needed, either. In these cases, VPMS is a service that provides much simpler operation.
Note that VPMS provides single coverage of receiver membership; that is, there is no distinct differentiation for multiple multicast groups. All traffic from a particular Attachment Circuit (AC) will be forwarded toward the same remote receivers, even if the destination MAC address is changed. Basically in VPMS, destination MAC addresses are not used for forwarding, which is significantly different from VPLS. If MAC-based forwarding is preferred (i.e., multicast/unicast differentiation of MAC address), VPLS should be chosen rather than VPMS.

4.2. ATM-based Use Case

TOC

A use case of ATM-based service in VPMS could be to offer the capability for service providers to support IP multicast wholesale services over ATM in case the wholesale customer relies on ATM infrastructure. The P2MP support alleviates the constraint in terms of replication for ATM to support IP multicast services.
Another use case of VPMS for ATM is for audio/video stream applications. Today many digital TV broadcasting networks adopt ATM-based distribution systems with point-to-multipoint PVPs/PVCs. The transport network supports replicating ATM cells in transit nodes to efficiently deliver programs to multiple terminals. For migrating such ATM-based networks onto IP/MPLS-based networks, VPMS is considered to be a candidate solution.

4.3. TDM-based Use Case

TOC

Today the existing VPWS already supports TDM emulation services (SAToP, CESoPSN or TDMoIP). It is a Layer 1 service, not Layer 2 service; however, a common architecture is being used since they are all packet-based emulations over a SP's network. VPMS is also considered to be a solution for such TDM applications that require point-to-multipoint topology.
In a PSN environment, the existing VPWS allows support for 2G/3G mobile backhauling (e.g. TDM traffic for GSM's Abis interface, ATM traffic for Release 99 UMTS's Iub interface). Currently, the Mobile backhauling architecture is always built as a star topology between the 2G/3G controller (e.g. BSC or RNC) and the 2G/3G Base Stations (BTS or NodeB). Therefore VPWSes (P2P services) are used between each Base

Station and their corresponding controller and nothing more is required.
As far as synchronization in a PSN environment is concerned, different mechanisms can be considered to provide frequency and phase clock required in the 2G/3G Mobile environment to guarantee mobile handover and strict QoS. One of them consists of using Adaptive Clock Distribution and Recovery. With this method a Master element distributes a reference clock at protocol level by regularly sending TDM PW packets (SAToP, CESoPSN or TDMoIP) to Slave elements. This process is based on the fact that the volume of transmitted data arrival is considered as an indication of the source frequency that could be used by the Slave element to recover the source clock frequency. Consequently, with the current methods, the PE connected to the Master must setup and maintain as many VPWS (P2P) as their are Slave elements, and the Master has to replicate the traffic. A better solution to deliver the clock frequency would be to use a VPMS which supports P2MP traffic. This may scale better than P2P services (VPWS) with regards to the forwarding plane at the Master since the traffic is no longer replicated to individual VPWSes (P2P) but only to the AC associated to the VPMS (P2MP). It may ease the provisioning process since only one source endpoint must be configured at the Ingress PE. This alleviated provisioning process would simplify the introduction of new Base Stations. The main gain would be to avoid replication on the Master and hence save bandwidth consumed by the synchronization traffic which typically requires the highest level of QoS. This kind of traffic will be competing with equivalent QOS traffic like VoIP, which is why it is significant to save the slightest bandwidth.

---

5. Reference Model

TOC

The VPMS reference model is shown in Figure 1.

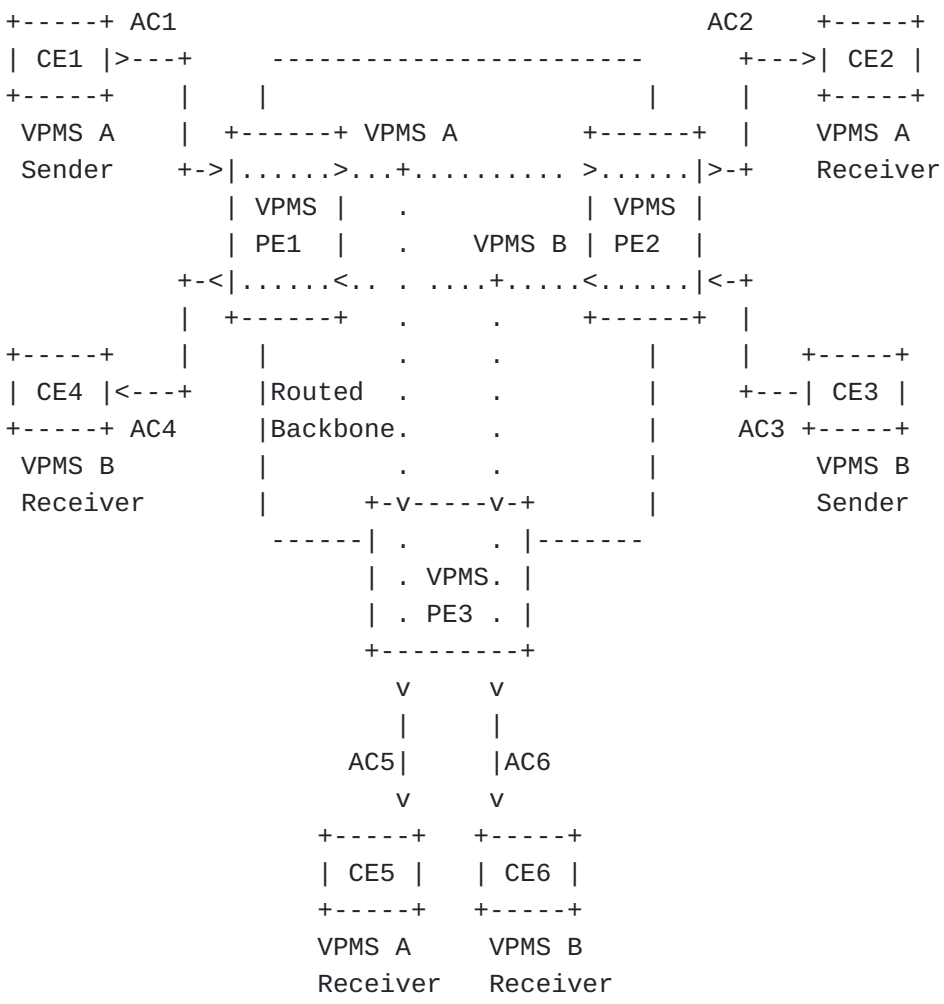

Figure 1: Reference Model for VPMS

A VPMS instance is defined as a service entity manageable in VPMS architecture. A single VPMS instance provides isolated service reachability domain to each CE, so it corresponds to a so-called "VPN" as a specific set of sites that allows communication. A single VPMS instance provides a unique unidirectional point-to-multipoint L2VPN service. In Figure 1, there are two VPMS instances shown, VPMS A and VPMS B. In principle, there is no traffic exchange allowed between these different instances, so they are treated as different VPNs.
In a VPMS, a single CE-PE connection is used for transmitting frames for delivery to multiple remote CEs, with point-to-multipoint duplication. The SP's network (PE as well as P) has a role to duplicate frames so that the traffic source does not need to send multiple frames to individual receivers.
Like VPWS, an Attachment Circuit (AC) is provided to accommodate CEs in a VPMS. In a VPMS, an AC attached to a VPMS MUST be configured as "sender" or "receiver" not both. That is, any AC is associated with the

role of either sending side (Tx) or receiving side (Rx) from the view of the CE. Thus every AC deals with unidirectional traffic flows. A sender AC does not have a capability of transmitting the traffic back to a CE at upstream side. Likewise a receiver AC does not have a capability of receive the traffic from a CE at downstream side. In Figure 1, AC1 and AC3 are configured as senders while AC2, AC4, AC5 and AC6 are configured as receivers. In VPMS A, CE1 could send traffic via AC1, but CE2 and CE5 could not send back traffic.

A CE which is locally connected to a sender AC is called a sender CE. Also a CE which is locally connected to a receiver AC is called a receiver CE. However, such CEs's roles will not be managed directly in VPMS because the configured AC's role (sender or receiver) will automatically determine them.

Basically there is a one-to-one mapping between an attachment circuit and a VPMS instance. For example, all traffic from CE1 to PE1 (thorough AC1) is mapped to VPMS A (to CE2 and CE5).

In a VPMS, PEs will be connected by PW technology which may include P2MP traffic optimization (i.e., P2MP PW. See section 7.2.). P2MP traffic optimization will provide the benefit of traffic replication for high bandwidth efficiency. The sender CE has only to transmit one stream towards the PE and it does not have to replicate traffic. Also routed backbone provides IP or MPLS-enabled PSN tunnels for transporting the PW traffic.

Regarding end-to-end traffic topology between the PEs, a single VPMS instance (i.e., one VPN) may correspond to a single unidirectional P2MP PW topology. In Figure 1, VPMS A (one instance) has a single P2MP PW topology (from PE1 to PE2 and PE3). However, there is also a case that a single VPMS consists of two or more P2MP PW topology grouped which is typically used for redundancy. The details are given in section 6.1.2.

VPMS can support various Layer 2 protocol services such as Ethernet, ATM, etc.

---

6. Customer Requirements

TOC

6.1. Service Topology

TOC

TOC

6.1.1. Point-to-Multipoint Support

A solution MUST support unidirectional point-to-multipoint connectivity from a sender to multiple receivers. A sender CE is assured to send traffic to one or more receiver CEs. Receiver CEs include not only the CEs which are located at remote sites, but also the local CEs which are connected to the same sender-side PE. If there is only one receiver in the instance, it is considered equivalent to unidirectional point-to-point traffic.

6.1.2. Multiple Source Support

TOC

A solution MUST support multiple sender topologies in one VPMS instance, where a common receiver group is reachable from two or more senders. This means that a solution needs to support having multiple P2MP topologies in the backbone whose roots are located apart in a common service. In other words, each P2MP topology MUST only have a single sender, however multiple P2MP topologies can be grouped together into a single VPMS instance. For example, in Figure 2, traffic from sender CE1 and CE2 both reach receivers CE3 and CE4 while CE1, CE2, CE3 and CE4 all are associated with a single service. This topology is useful for increasing service reliability by redundant sources. Note that every receiver has only to have one AC connected to each PE to receive traffic. (in Figure 2, AC3 and AC4 respectively). Thus a solution will also need to support protection and restoration mechanism combining these multiple P2MP topologies. (See section 6.4 too).

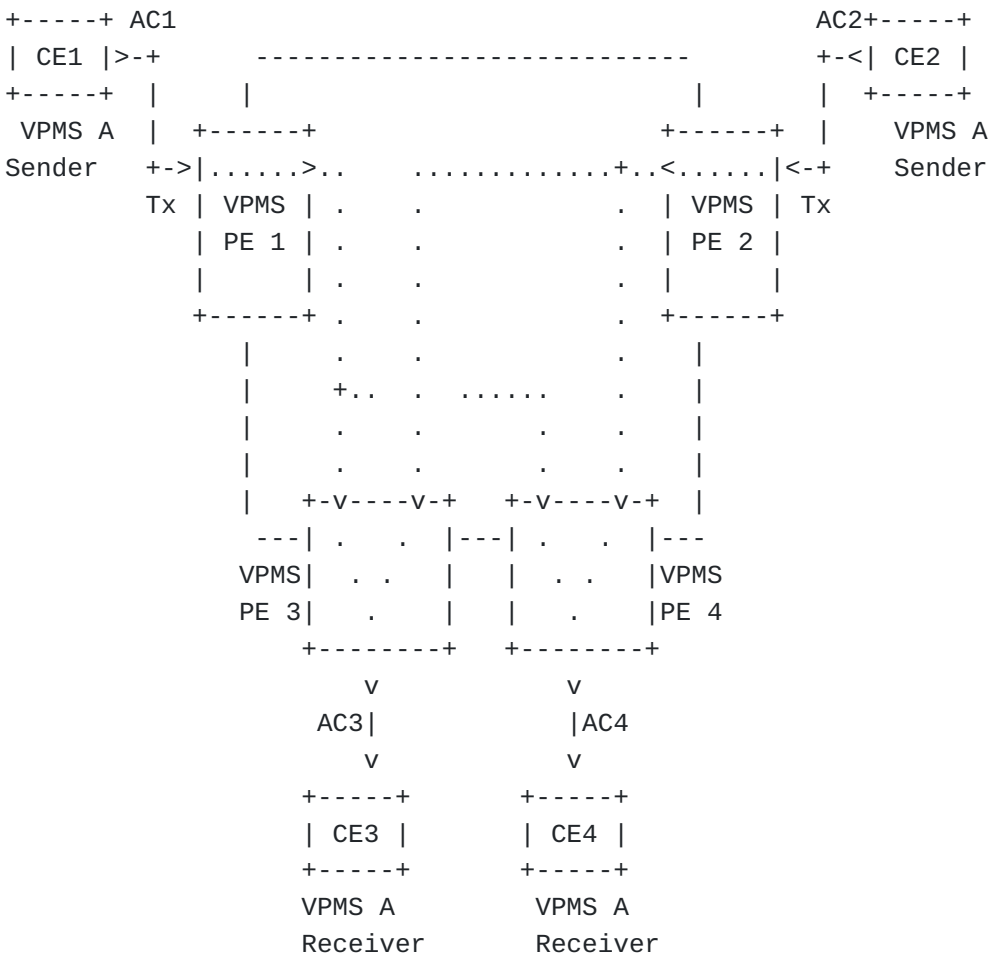

Figure 2: Multiple source support

6.1.3. Reverse Traffic Support

TOC

There are cases where a reverse traffic flow is necessary. A sender CE might sometimes want to receive traffic from a receiver CE. There are some usage scenarios for this, such as stream monitoring through a loopback mechanism, control channels which need feedback communication etc. The simplest way to accomplish this is to provide different VPMS instances for reverse traffic, i.e. a sender CE is a receiver of another VPMS instance.
Figure 3 illustrates this kind of reverse traffic scenario, where CE1 is configured as a sender in VPMS A and a receiver in VPMS B. VPMS B is used for reverse traffic. Note that a closed single network here is

composed of two VPMS instances. In operational terms, CE1 and CE4 belong to the same closed "VPN" by administrative policy (e.g., CE1, CE2, CE3 and CE4 are the devices in one enterprise's intranet network). Such bi-directional instances can be easily created if two distinct ACs are provisioned for sending and receiving exclusively (e.g., if VLAN id in dot1Q tagged frame is a service delimiter, different VLAN ids are uniquely allocated for Tx and Rx). This approach is acceptable if a receiver CE device can change Layer 2 interface appropriately in data transmitting and receiving.

Meanwhile it is also true that this might be considered a limitation in some deployment scenarios. If a CE is an IP router or Ethernet bridge, reverse traffic is normally expected to be received on the same interface as forward traffic on the receiver CE. (i.e., the same VLAN id is to be used for reverse traffic if the AC supports dot1Q tagged frames.)

Therefore, in a VPMS solution, both of the two type of ACs, sending (Tx) and receiving (Rx), SHOULD be allowed to be placed in the same physical/virtual circuit. In Figure 3, suppose AC5 of VPMS A is provisioned as {VLAN id = 100, direction= Rx}. It is expected that operators can provision AC6 of VPMS B in the same physical port as {VLAN id = 100, direction = Tx}. That is, the combination between VLAN id and the flow direction is now considered to be a service delimiter. Note, in most implementations of VPWS today, every AC is always considered bidirectional and a unique Layer 2 header/circuit (ATM VPI/VCI, an Ethernet port, a VLAN etc.) is considered the service delimiter. In contrast in VPMS, every AC is considered unidirectional and traffic direction is an additional element to identify a unique AC.

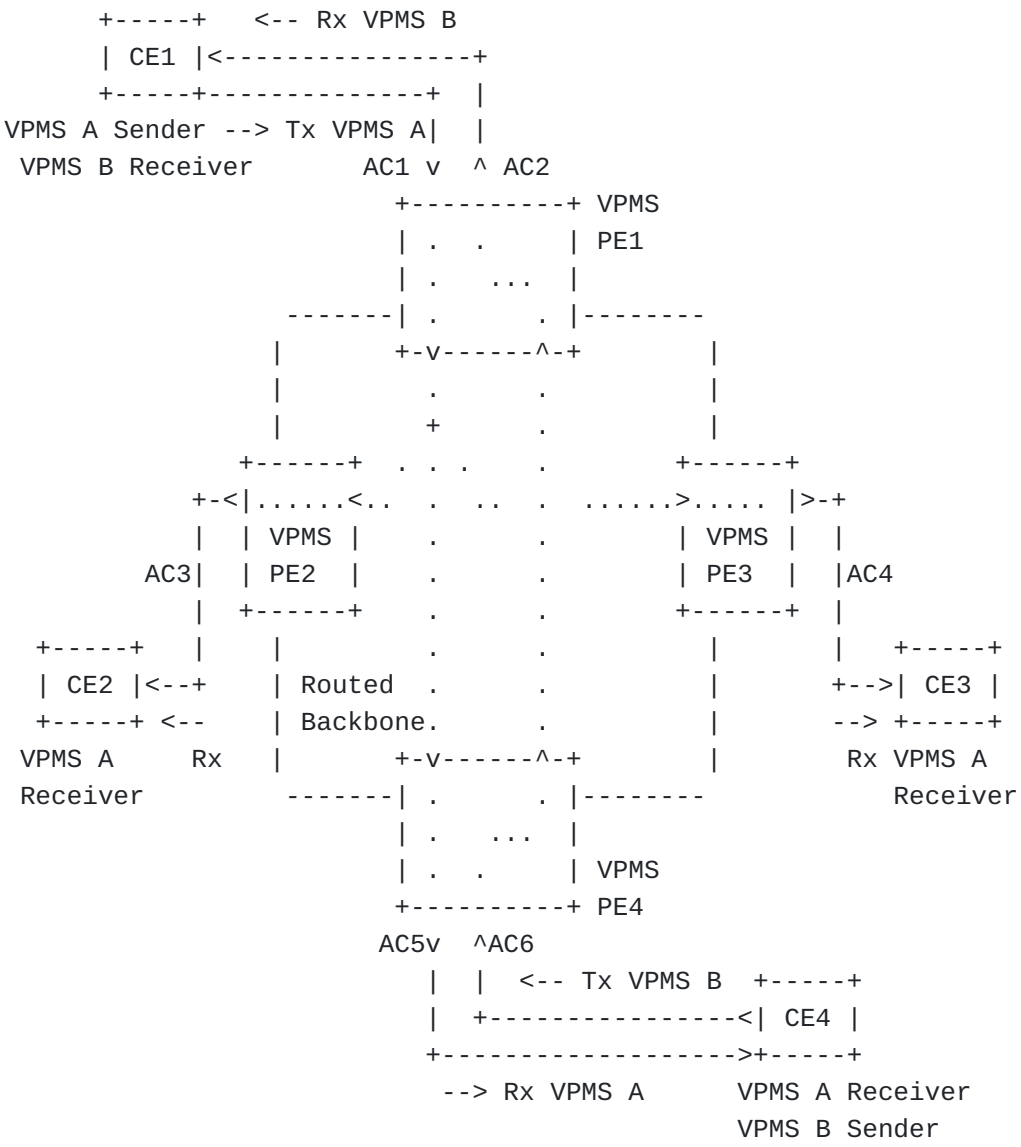

Figure 3: Reverse traffic support

6.2. Transparency

TOC

A solution is intended to provide Layer 2 traffic transparency. Transparency SHOULD be honoured per VPMS instance basis. In other words, Layer 2 traffic can be transparently transported from a sender CE to receiver CEs in a given instance. Note, however, if service delimiting fields (VLAN Id in Ethernet, VPI/VCI in ATM, DLCI in FR etc.) are assigned by SP, they are not transparent. It depends on SP's choice if they are assigned at each AC. Hence it could be that some of

receiver CEs are getting traffic with different delimiting fields than the other receiver CEs.
VPMS solution SHOULD NOT require any special packet processing by the end users (CEs).

6.3. Quality of Service (QoS)

TOC

A customer may require that the VPMS service provide the guaranteed QoS. In particular, for real time applications which are considered common in point-to-multipoint delivery, delay and loss sensitive traffic MUST be supported. The solution SHOULD provide native QoS techniques for service class differentiation, such as IEEE 802.1p CoS for Ethernet.
For bandwidth committed services (e.g., ATM CBR), a solution SHOULD guarantee end-to-end bandwidth. It MAY provide flow admission control mechanisms to achieve that.

6.4. Protection and Restoration

TOC

A solution MUST provide protection and restoration mechanism for end-to-end services.

6.4.1. Dual-homed Access Support

TOC

A solution MUST allow dual-homed redundant access from a CE to multiple PEs. Additionally, a solution SHOULD provide protection mechanism between the different PEs to which a CE is attached. This is because when an ingress PE node fails whole traffic delivery will fail unless a backup sender PE is provided, even in case of dual-homed access. Similarly, if an egress PE node fails, traffic toward that CE is never received unless a backup egress PE is provided. Figure 4 is an example for this access topology.

6.4.2. Single/Dual Traffic Support in Dual-homed Access

TOC

When dual-homed access to sender PEs is provided, a solution MAY allow a sender CE to transmit just a single copy of the traffic to either one of the two sender PEs, or to transmit a copy of the traffic to both the PEs simultaneously. The latter scenario consumes more resource of CE-PE

link than the single traffic scenario, but it is usually applicable when a source device has only a simple forwarding capability without any switchover functionality. In the dual traffic case, the backup ingress PE SHOULD be able to filter the incoming unnecessary traffic while active PE is working. Also in either case, single traffic or dual traffic, the protection mechanism of ingress PEs described in the previous subsection will be necessary to handle the traffic appropriately.

In the case of dual-homed access to receiver PEs, a solution MAY allow a receiver CE to receive a single copy of the traffic from either one of the two egress PEs, or receive a copy of the traffic from both PEs simultaneously. The dual traffic approach is applicable if CE has fast switchover capability as a receiver by selecting either one of incoming traffic, but note that additional traffic resources are always consumed at PE-CE link of backup side. Specifically in the single traffic case, it might be needed to support switchover mechanism between egress PEs in failure.

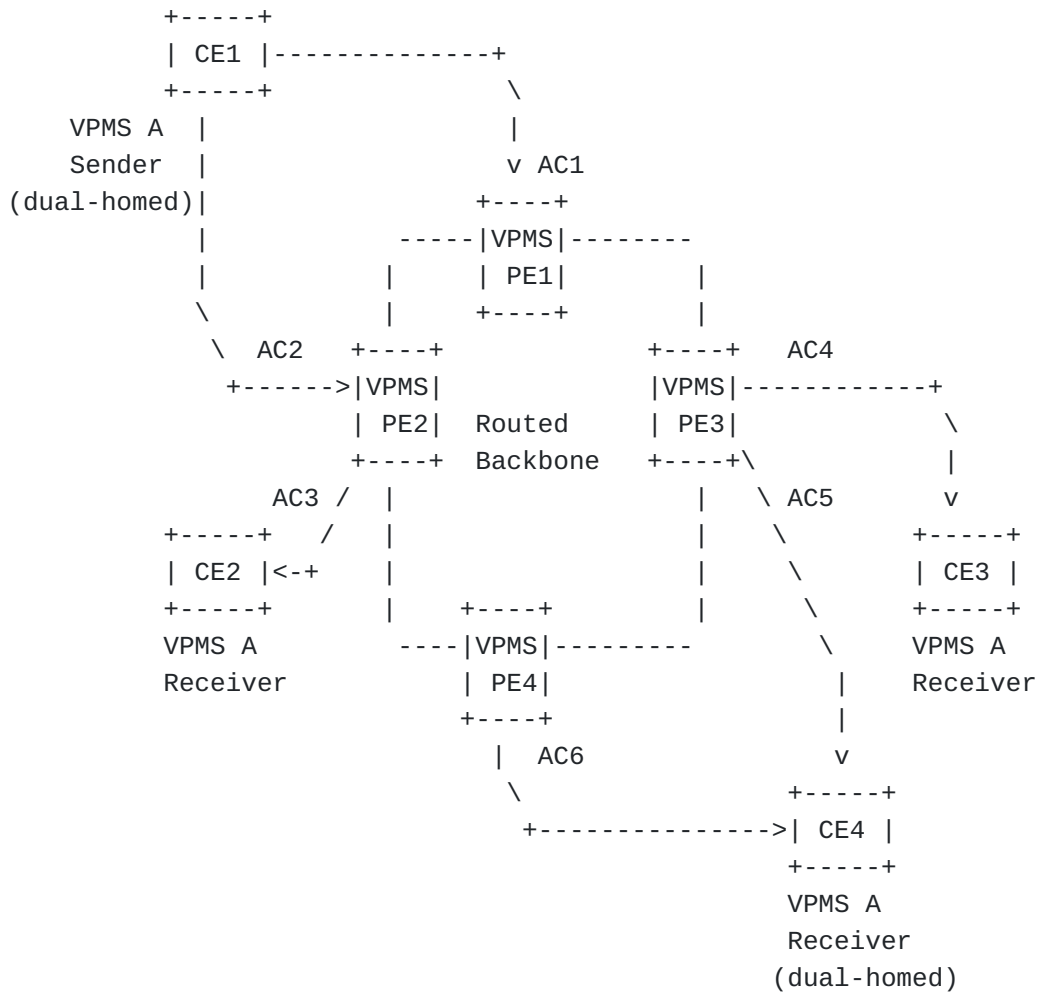

Figure 4: Dual homing support

6.5. Security

TOC

The basic security requirement raised in Section 6.5 of [RFC4665] (Augustyn, W. and Y. Serbest, “Service Requirements for Layer 2 Provider-Provisioned Virtual Private Networks,” September 2006.) also applies to VPMS.
In addition, a VPMS solution MAY have the mechanisms to activate the appropriate filtering capabilities (for example, MAC/VLAN filtering etc.), and it MAY be added with the filtering control mechanism between particular sender/receiver sites inside a VPMS instance. For example, in Figure 1, filtering can be added such that traffic from CE1 to CE4 and CE5 is allowed but traffic from CE1 to CE6 is filtered.

6.6. Reordering Prevention

TOC

A solution SHOULD prevent Layer 2 frame reordering when delivering customer traffic under normal conditions.

6.7. Failure reporting

TOC

A solution MAY provide information to the customer about failures. For example, if there is a loss of connectivity toward some of the receiver CEs, it is reported to the sender CE.

7. Service Provider Network Requirements

TOC

7.1. Scalability

TOC

A VPMS solution MUST be designed to scale well with an increase in the number of any of the following metrics:

- the number of PEs (per VPMS instance and total in a SP network)
- the number of VPMS instances (per PE and total)

- the number of sender CEs (per PE, VPMS instance and total)

- the number of receiver CEs (per PE, VPMS instance and total)

A VPMS solution SHALL document its scalability characteristics in quantitative terms. A solution SHOULD quantify the amount of state that a PE and a P device has to support.
The scalability characteristics SHOULD include:

- the processing resources required by the control plane in managing PWs (neighborhood or session maintenance messages, keepalives, timers, etc.)

- the processing resources required by the control plane in managing PSN tunnels

- the memory resources needed for the control plane

- other particular elements inherent to each solution that impact scalability

---

7.2. Pseudo Wire Signaling and PSN Tunneling

TOC

A VPMS solution SHOULD provide an efficient replication that can contribute to optimizing the bandwidth usage required in a SP's network. For supporting efficient replication, it is expected to take advantage of PW and PSN mechanisms that are capable of P2MP traffic. Regarding PW mechanism, [I-D.ietf-pwe3-p2mp-pw-requirements] (JOUNAY, F., Niger, P., Kamite, Y., DeLord, S., and L. Martini, “Requirements for Point-to-Multipoint Pseudowire,” September 2008.) introduces P2MP PW concept and its requirements, showing two basic approaches of providing replication. One is SS (Single Segment)-PW model that provides replication by PSN tunnel such as P2MP LSP (i.e., by outer label layer), and the other is MS (Multi Segment)-PW model that provides replication by multiple interconnected PWs (i.e., by inner label layer). In either case, end-to-end P2MP topology in VPMS is common from the view of PEs and ACs. Requirements as a provider service specified in this document will be commonly applied regardless of P2MP PW's signaling model.
This document does not raise any specific requirements for particular PSN tunneling schemes (point-to-point, point-to-multipoint and multipoint-to-multipoint) that is applied only to VPMS. The actual type of PSN tunnel used in VPMS will be dependent on individual deployment scenarios (e.g., which PSN protocol is available now in the core and how much network resources operators will want to optimize).

7.3. Discovering VPMS Related Information

TOC

A solution SHOULD support auto-discovery methods that dynamically allow VPMS information to be discovered by the PEs to minimize the amount of configuration the SP must perform.
All of the requirements on discovery described in Section 7.3 of [RFC4665] (Augustyn, W. and Y. Serbest, “Service Requirements for Layer 2 Provider-Provisioned Virtual Private Networks,” September 2006.) SHOULD be satisfied in VPMS as well.
Auto-discovery will help operators' initial configuration of adding a new VPN (i.e., VPMS instance), adding/deleting new sender/receiver, and so on.
The information related to remote sites will be as follows:

- Information to identify the VPMS instance
- PE router ID / IP address as location information
- Information to identify Attachment Circuits and their associated group information to compose a unique service (i.e., VPMS instance).
- AC role in each VPMS (Sender or Receiver)
- SP-related information (AS number, etc. for an inter-provider case)

Following is an example scenario: by default, every PE will have the association among the information described above. Suppose a new PE having an AC is provisioned in the existing VPMS instance and this AC is configured as receiver. This information will be automatically discovered by the other existing remote PEs (i.e., ingress and egress PEs in the same VPMS instance). Once the ingress PE discovers this new PE/AC, it can automatically add it as the new leaf of P2MP topology according to P2MP PW signaling mechanism. This operation does not require any new configuration at the existing PEs.

7.4. Activation and Deactivation

TOC

This section raises generic requirements for handling related information about remote sites after the initial provisioning to ease the total operation of VPMS.

A solution SHOULD provide a way to activate/deactivate the administrative status of each CE/AC. After initial provisioning, a SP might change connectivity configuration between particular CEs inside a single VPMS instance for operational reasons. This feature will be beneficial to help such a scenario.

For example, in Figure 5, CE1, CE3, CE4 and CE5 (and their ACs) are initially provisioned for VPMS A. CE2 is not provisioned for any VPMSes. In VPMS A, CE1 is a sender and CE3, CE4 and CE5 are receivers. Traffic will usually flow from CE1 to all receivers, CE3, CE4 and CE5. However, for maintenance operation, application's request (e.g., stream program has changed) or some other reasons, CE4 needs to be set as administratively deactivated. Then it becomes necessary to turn off traffic from PE4 to CE4. This operation must be appropriately distinguished from failure cases.

When deactivating a particular site, backbone PSN/PW resources (e.g., admission control of PSN tunnel) MAY be released for that particular direction in order to provide that bandwidth to other services. In Figure 5, CE3 is now administratively activated and receiving traffic. However, if CE3 comes to be administratively deactivated, and if RSVP-TE (including P2P and/or P2MP) is used for backbone PSN, then TE reserved resources from PE1 to PE3 may be released.

In addition, a solution SHOULD allow single-sided activation operation at a sender PE. In some scenarios, operators prefer centralized operation. This is often considered natural for one-way digital audio/video distribution applications: SPs often want to complete their service delivery by a single operation at one source PE, not by multiple operations at many receiver PEs. Figure 5 illustrates this scenario, where a SP only has to do single-sided operation at PE1 (source) to administratively activate/deactivate various connections from AC1 to AC3, AC4 and/or AC5. It is not needed to perform operations on PE3 and PE4 directly.

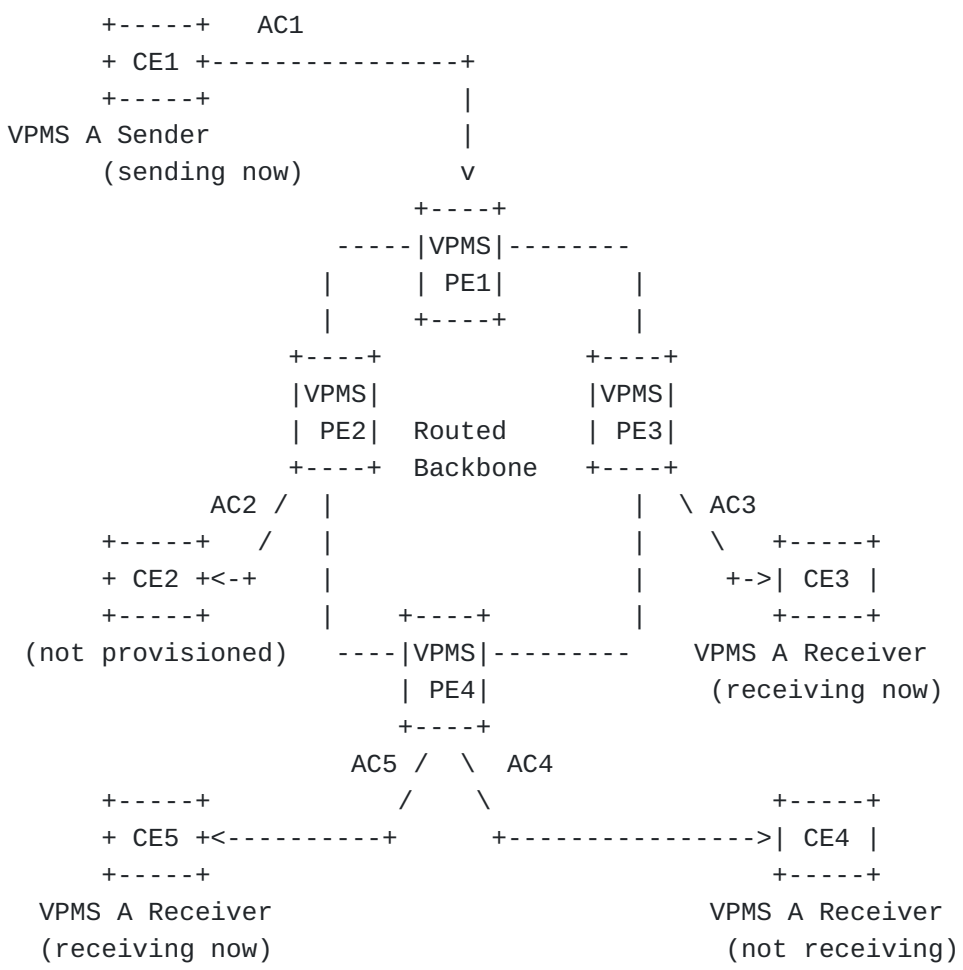

Figure 5: Site activation and deactivation

7.5. Inter-AS Support

TOC

A solution SHOULD support inter-AS scenarios, where there is more than one provider providing a common VPMS instance and VPN. More specifically, it is necessary to consider the case where some of the PEs that compose one VPMS belong to several different ASes.

7.6. Co-existence with Existing L2VPNs

TOC

A solution MUST co-exist with the existing L2VPNs (e.g., VPWS, VPLS) across the same SP's network. A solution MUST NOT impede the operation of auto-discovery and signalling mechanism that are already supported by the PEs for those existing L2VPNs.

7.7. Operation, Administration and Maintenance

TOC

7.7.1. Fault Management

TOC

7.7.1.1. Fault Detection

TOC

A solution MUST provide tools that detect reachability failure and traffic looping of P2MP transport in a VPMS instance. If multiple sources are supported (i.e., multiple P2MP topologies are grouped together into a single VPMS instance), such tools MUST be able to perform distinguishing each P2MP topology.

7.7.1.2. Fault Notification

TOC

A solution MUST provide fault notification and trouble tracking mechanisms. (e.g. SNMP-trap and syslog that notify fault to remote NMS.)
In VPMS one point of failure at upstream often affects a number of downstream PEs and ACs that might raise a notification message. Hence notification messages MAY be summarized or compressed for operators' ease of management.
In case of receiver-side failure (receiver PE or its AC), this fault status SHOULD be able to be monitored at sender PE. This will help an operator to monitor each receiver PEs/AC in a centralized manner; that is, a sender PE can collect receiver-side information. How this status is transferred depends on a solution.
In contrast, in case of sender-side failure (sender PE or its AC), this fault status SHOULD also be able to be monitored at receiver PEs. This will help an operator to troubleshoot at receiver PEs (i.e.,

distinguish local AC’s failure from remote upstream AC’s failure easily).
In any case of failure at SP’s network, fault information MAY be notified to the customer. Specifically, such fault MAY trigger generating customer OAM message toward CEs (e.g., AIS) and/or shutting down receiver ACs.

7.7.1.3. Fault Isolation

TOC

A solution MUST provide diagnostic/troubleshooting tools for P2MP transport in a VPMS instance.

7.7.2. Testing

TOC

A solution MUST provide a mechanism for testing each P2MP connectivity and verifying the associated information in a VPMS instance. The connectivity is between sender and all receiver ACs.
Operators will run testing before and after service activation. Testing mechanism SHOULD support end-to-end testing of the data path used by customer's data. End-to-end testing will have CE-to-CE path test and PE-to-PE path test. A solution MUST support PE-to-PE path test and MAY support CE-to-CE path test. In either case the data path provided for each VPMS is unidirectional, hence if loopback testing is supported, additional consideration about reverse-path might also be needed (see section 6.1.3).

7.7.3. Performance Management

TOC

A solution MUST offer mechanisms to monitor traffic performance parameters and statistics in each P2MP traffic.
A solution MUST provide access to:

- Traffic statistics (total traffic forwarded, incoming, outgoing, dropped, etc., by period of time)

A solution SHOULD provide access to:

- Performance information related to traffic usage, e.g., one-way delay, one-way jitter, one-way loss, delay variations (the difference of various one-way delay from a particular sender PE to multiple receiver PEs) etc.

All or part of this information SHOULD be made available through standardized SNMP MIB Modules (Management Information Base).
It is expected that such information can be used for SLA monitoring between sender and receiver, to give the SP a clear picture of current service providing to the customer.

7.8. Security

TOC

TBD (for further study for next revision)

---

8. Security Considerations

TOC

Security consideration will be covered by section 6.5. and section 7.8. (This is for further study for next revision.)

---

9. IANA Considerations

TOC

This document has no actions for IANA.

---

10. Acknowledgments

TOC

Many thanks to Ichiro Fukuda, Kazuhiro Fujihara, Ukyo Yamaguchi and Kensuke Shindome for their valuable review and feedback.

---

11. References

TOC

11.1. Normative References

TOC

[RFC2119]	Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC4026]	

	Andersson, L. and T. Madsen, “Provider Provisioned Virtual Private Network (VPN) Terminology,” RFC 4026, March 2005 (TXT).

11.2. Informative References

TOC

[I-D.ietf-l2vpn-vpls-mcast]	Aggarwal, R., Kamite, Y., Fang, L., and Y. Rekhter, “Multicast in VPLS,” draft-ietf-l2vpn-vpls-mcast-04 (work in progress), June 2008 (TXT).
[I-D.ietf-pwe3-p2mp-pw-requirements]	JOUNAY, F., Niger, P., Kamite, Y., DeLord, S., and L. Martini, “Requirements for Point-to-Multipoint Pseudowire,” draft-ietf-pwe3-p2mp-pw-requirements-00 (work in progress), September 2008 (TXT).
[RFC4664]	Andersson, L. and E. Rosen, “Framework for Layer 2 Virtual Private Networks (L2VPNs),” RFC 4664, September 2006 (TXT).
[RFC4665]	Augustyn, W. and Y. Serbest, “Service Requirements for Layer 2 Provider-Provisioned Virtual Private Networks,” RFC 4665, September 2006 (TXT).
[RFC4761]	Kompella, K. and Y. Rekhter, “Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling,” RFC 4761, January 2007 (TXT).
[RFC4762]	Lasserre, M. and V. Kompella, “Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling,” RFC 4762, January 2007 (TXT).

Authors' Addresses

TOC

	Yuji Kamite
	NTT Communications Corporation
	Tokyo Opera City Tower
	3-20-2 Nishi Shinjuku, Shinjuku-ku
	Tokyo 163-1421
	Japan
Email:	y.kamite@ntt.com
	Frederic Jounay
	France Telecom
	2, avenue Pierre-Marzin
	22307 Lannion Cedex

	France
Email:	frederic.jounay@orange-ftgroup.com
	Ben Niven-Jenkins
	BT
	208 Callisto House, Adastral Park
	Ipswich, IP5 3RE
	UK
Email:	benjamin.niven-jenkins@bt.com
	Deborah Brungard
	AT&T
	Rm. D1-3C22, 200 S. Laurel Ave.
	Middletown, NJ, 07748
	USA
Email:	dbrungard@att.com
	Lizhong Jin
	Nokia Siemens Networks
	Building 89, 1122 North QinZhou Road,
	Shanghai, 200211
	P.R.China
Email:	lizhong.jin@nsn.com

Full Copyright Statement

TOC

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the

procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.