l3vpn Working Group
Internet-Draft
Intended status: Informational
Expires: May 17, 2007

T. Morin, Ed.
France Telecom R&D

November 13, 2006

Requirements for Multicast in L3 Provider-Provisioned VPNs
draft-ietf-l3vpn-ppvpn-mcast-reqts-10

Status of this Memo

Copyright Notice

Abstract

This document presents a set of functional requirements for network solutions that allow the deployment of IP multicast within L3 Provider Provisioned Virtual Private Networks (PPVPNs). It specifies requirements both from the end user and service provider standpoints. It is intended that potential solutions specifying the support of IP multicast within such VPNs will use these requirements as guidelines.

Working group

This document is a product of the IETF's Layer 3 Virtual Private Network (l3vpn) working group. Comments should be addressed to the WG's mailing list at <mailto:l3vpn@ietf.org>. The charter for l3vpn may be found at <http://www.ietf.org/html.charters/l3vpn-charter.html>

Table of Contents

1. Introduction

VPN services satisfying the requirements defined in [RFC4031] are now being offered by many service providers throughout the world. VPN services are popular because customers need not be aware of the VPN technologies deployed in the provider network. They scale well for the following reasons:

- because P routers (Provider Routers) need not be aware of VPN service details

- because the addition of a new VPN member requires only limited configuration effort

There is also a growing need for support of IP multicast-based services. Efforts to provide efficient IP multicast routing protocols and multicast group management have been done in standardization bodies which has led, in particular, to the definition of the PIM and IGMP protocols.

However, multicast traffic is not natively supported within existing L3 PPVPN solutions. Deploying multicast over an L3VPN today, with

only currently standardized solutions, requires designing customized solutions which will be inherently limited in terms of scalability, operational efficiency and bandwidth usage.

This document complements the generic L3VPN requirements [RFC4031] document, by specifying additional requirements specific to the deployment within PPVPNs of services based on IP multicast. It clarifies the needs of both VPN clients and providers and formulates the problems that should be addressed by technical solutions with the key objective being to remain solution agnostic. There is no intent to either specify solution-specific details in this document or application-specific requirements. Also this document does NOT aim at expressing multicast-related requirements that are not specific to L3 PPVPNs.

It is expected that solutions that specify procedures and protocol extensions for multicast in L3 PPVPNs SHOULD satisfy these requirements.

2. Conventions used in this document

2.1. Terminology

Although the reader is assumed to be familiar with the terminology defined in [RFC4031], [RFC4364], [RFC4601], [RFC4607] the following glossary of terms may be worthwhile.

Moreover we also propose here generic terms for concepts that naturally appear when multicast in VPNs is discussed.

ASM:
: Any Source Multicast. One of the two multicast service models, in which a terminal subscribes to a multicast group to receive data sent to the group by any source.

Multicast-enabled VPN, multicast VPN, or mVPN:
a VPN which supports IP multicast capabilities, i.e. for which some PE devices (if not all) are multicast-enabled and whose core architecture supports multicast VPN routing and forwarding.

PPVPN:
Provider-Provisioned Virtual Private Network.

PE/CE:
"Provider Edge", "Customer Edge" (as defined in [RFC4026]). As suggested in [RFC4026], we will use these notations to refer to the equipments/routers/devices themselves. Thus, "PE" will refer to the router on the provider's edge, which faces the "CE", the router on the customer's edge.

VRF or VR:
By this phrase, we refer to the entity defined in a PE dedicated to a specific VPN instance. "VRF" refers to "VPN Routing and Forwarding table" as defined in [RFC4364], and "VR" to "Virtual Router" as defined in [I-D.ietf-l3vpn-vpn-vr] terminology.

MDTunnel:
Multicast Distribution Tunnel, the means by which the customer's multicast traffic will be transported across the SP network. This is meant in a generic way: such tunnels can be either point-to-point or point-to-multipoint. Although this definition may seem to assume that distribution tunnels are unidirectional, the wording also encompasses bi-directional tunnels.

S:
Denotes a multicast source.

G:
Denotes a multicast group.

Multicast channel:
In the multicast SSM model[RFC4607], a "multicast channel"

designate traffic from a specific source S to a multicast group G. Also denominated as "(S,G)".

SP:
Service provider.

SSM:
Source Specific Multicast. One of the two multicast service models, where a terminal subscribes to a multicast group to receive data sent to the group by a specific source.

RP:
Rendez-vous point (PIM-SM [RFC4601]).

P2MP, MP2MP: Designate "Point to multipoint" and "Multipoint to multipoint" replication trees.

L3VPN, VPN:
Throughout this document, "L3VPN" or even just "VPN" will refer to "Provider-Provisioned Layer 3 Virtual Private Network" (PP L3VPNs), and will be preferred for readability.

Please refer to [RFC4026] for details about terminology specifically relevant to VPN aspects, and to [RFC2432] for multicast performance or QoS related terms.

2.2. Conventions

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].

3. Problem Statement

3.1. Motivations

More and more L3VPN customers use IP multicast services within their private infrastructures. Naturally, they want to extend these multicast services to remote sites that are connected via a VPN.

For instance, the customer could be a national TV channel with several geographical locations that wants to broadcast a TV program from a central point to several regional locations within its VPN.

A solution to support multicast traffic could consist of point-to-point tunnels across the provider network and requires the PEs (Provider Edge routers) to replicate traffic. This would obviously be sub-optimal as it would place the replication burden on the PE and hence would have very poor scaling characteristics. It would also probably waste bandwidth and control plane resources in the provider's network.

Thus, to provide multicast services for L3VPN networks in an efficient manner (that is, with a scalable impact on signaling and protocol state as well as bandwidth usage), in a large scale environment, new mechanisms are required to enhance existing L3VPN solutions for proper support of multicast-based services.

3.2. General Requirements

This document sets out requirements for L3 provider-provisioned VPN solutions designed to carry customers' multicast traffic. The main requirement is that a solution SHOULD first satisfy the requirements documented in [RFC4031]: as far as possible, a multicast service should have the same characteristics as the unicast equivalent, including the same simplicity (technology unaware), the same quality of service (if any), the same management (e.g. performance monitoring), etc.

Moreover, it also has to be clear that a multicast VPN solution MUST interoperate seamlessly with current unicast VPN solutions. It would also make sense that multicast VPN solutions define themselves as extensions to existing L3 provider-provisioned VPN solutions (such as for instance, [RFC4364] or [VRs]) and retain consistency with those, although this is not a core requirement.

The requirements in this document are equally applicable to IPv4 and IPv6, for both customer and provider related matters.

3.3. Scaling vs. Optimizing Resource Utilization

When transporting multicast VPN traffic over a service provider network, there intrinsically is tension between scalability and resource optimization, since the latter is likely to require the maintenance of control plane states related to replication trees in the core network [RFC3353].

Consequently, any deployment will require a trade-off to be made and this document will express some requirements related to this trade-off.

4. Use cases

The goal of this section is to highlight how different applications and network contexts may have a different impact on how a multicast VPN solution is designed, deployed and tuned. For this purpose we describe some typical use case scenarios and express expectations in terms of deployment orders of magnitude.

Most of the content of these sections originates from a survey done in summer 2005, among institutions and providers that expect to deploy such solutions. The full survey text, and raw results (13 responses) were published separately and we only present here the most relevant facts and expectations that the survey exposed.

For scalability figures, we considered that it was relevant to highlight the highest expectations, those that are expected to have the greatest impact on solution design ; for balance, we do also mention cases where such high expectations were expressed in only a few answers.

4.1. Scenarios

We don't provide here an exhaustive set of scenarios that a multicast VPN solution is expected to support - no solution should restrict the scope of multicast applications and deployments that can be done over a multicast VPN.

Hence, we only give here a short list of scenarios that are expected to have a large impact on the design of a multicast VPN solution.

4.1.1. Live content broadcast

Under this label, we group all applications that distribute content (audio, video, or other content) with the property that this content is expected to be consulted at once ("live") by the receiver. Typical applications are broadcast TV, production studios connectivity, distribution of market data feeds.

The characteristics of such applications are the following:

- one or few sources to many receivers
- sources are often in known locations, receivers are in less predictable locations (this latter point may depend on applications)
- in some cases, it is expected that the regularity of audience patterns may help improve how the bandwidth/state trade-off is

handled

- the number of streams can be as high as hundreds, or even thousands of streams
- bandwidth will depend on the application, but may vary between a few tens/hundreds of Kb/s (e.g audio or low quality video media) and tens of Mb/s (high quality video), with some demanding professional applications requiring as much as hundreds of Mb/s.
- QoS requirements include, in many cases, a low multicast group join delay
- QoS of these applications is likely to be impacted by packet loss (some applications may be robust to low packet loss), and to have low robustness against jitter
- delay sensitivity will depend on the application: some applications are not so delay sensitive (e.g. broadcast TV), whereas others may require very low delay (professional studio applications)
- some of these applications may involve rapid changes in customer multicast memberships as seen by the PE, but this will depend on audience patterns and on the amount of provider equipments deployed close to VPN customers

4.1.2. Symmetric applications

Some use cases exposed by the survey can be grouped under this label, and include many-to-many applications such as conferencing, server clusters monitoring.

They are characterized by the relatively high number of streams that they can produce, which has a direct impact on scalability expectations.

A sub-case of this scenario is the case of symmetric applications with small groups, when the number of receivers is low compared to the number of sites in the VPNs (e.g.: video conferencing and e-learning applications).

This latter case is expected to be an important input to solution design, since it may significantly impact how the bandwidth/state is managed.

Because of:

- small groups, and low predictability of the location of participants ("sparse groups")
- possibly significantly high bandwidth (a few Mb/s per participant)

...optimizing bandwidth may require introducing dedicated states in the core network (typically as much as the number of groups).

Lastly, some of these applications may involve realtime interactions, and will be highly sensitive to packet loss, jitter and delay.

4.1.3. Data distribution

Some applications which are expected to be deployed on multicast VPNs are non-realtime applications aimed at distributing data from few sources to many receivers.

Such applications may be considered to have lower expectations than their counterparts proposed in this document, since they would not necessarily involve more data streams and are more likely to adapt to the available bandwidth and to be robust to packet loss, jitter and delay.

One important property is that such applications may involve higher bandwidths (hundreds of Mb/s).

4.1.4. Generic multicast VPN offer

This ISP scenario is a deployment scenario where IP-Multicast connectivity is proposed for every VPN: if a customer requests a VPN, then this VPN will support IP-Multicast by default. In this case the number of multicast VPNs equals the number of VPNs. This implies a quite important scalability requirement (e.g. hundreds of PEs, hundreds of VPNs per PE, with a potential increase by one order of magnitude in the future).

The per mVPN traffic behavior is not predictable because how the service is used is completely up to the customer. This results in a traffic mix of the scenarios mentioned in section 4.1. QoS requirements are similar to typical unicast scenarios, with the need for different classes. Also in such a context, a reasonably large range of protocols should be made available to the customer for use at the PE-CE level.

Also, in such a scenario, customers may want to deploy multicast connectivity between two or more multicast VPNs as well as access to Internet Multicast.

4.2. Scalability orders of magnitude

This section proposes orders of magnitude for different scalability metrics relevant for multicast VPN issues. It should be noted that the scalability figures proposed here relate to scalability expectations of future deployments of multicast VPN solutions, as the authors chose to not restrict the scope to only currently known deployments.

4.2.1. Number of VPNs with multicast enabled

From the survey results, we see a broad range of expectations. There are extreme answers: from 5 VPNs (1 answer) to 10k VPNs (1 answer), but more typical answers are split between the low range -tens of VPNs- (7 answers) or in the higher range of hundreds or thousands of VPNs (2 + 4 answers).

A solution SHOULD support a number of multicast VPNs ranging from one

to several thousands.

A solution SHOULD NOT limit the proportion of multicast VPNs among all (unicast) VPNs.

4.2.2. Number of multicast VPNs per PE

The majority of survey answers express a number of multicast VPNs per PE of around tens (8 responses between 5 and 50); a significant number of them (4) expect deployments with hundreds or thousands (1 response) of multicast VPNs per PE.

A solution SHOULD support a number of multicast VPNs per PE of several hundreds, and may have to scale up to thousands of VPNs per PE.

4.2.3. Number of CEs per multicast VPN per PE

Survey responses span from 1 to 2000 CEs per multicast VPN per PE. Most typical responses are between tens (6 answers) and hundreds (4 responses).

A solution SHOULD support a number of CEs per multicast VPN per PE going up to several hundreds (and may target the support of thousands of CEs).

4.2.4. PEs per multicast VPN

People who answered the survey typically expect deployments with number of PEs per multicast VPN in the range of hundreds of PEs (6

responses) or tens of PEs (4 responses). Two responses were in the range of thousands (one mentioned a 10k figure).

A multicast VPN solution SHOULD support several hundreds of PEs per multicast VPN, and MAY usefully scale up to thousands.

4.2.4.1. ... with sources

The number of PEs, per VPN, that would be connected to sources, seems to be significantly lower than the number of PEs per VPN. This is obviously related to the fact that many respondents mentioned

deployments related to content broadcast applications (one to many).

Typical numbers are of tens of source-connected-PEs (6 responses), or hundreds (4 responses). One respondent expected a higher number of several thousands.

A solution SHOULD support hundreds of source-connected-PEs per VPN, and some deployment scenarios involving many-to-many applications, may require supporting a number of source-connected-PEs equal to the number of PEs (hundreds or thousands).

4.2.4.2. ... with receivers

The survey showed that the number of PEs with receivers is expected to be of the same order of magnitude as the number of PEs in a multicast VPN. This is consistent with the intrinsic nature of most multicast applications, which have few source only participants.

4.2.5. PEs with multicast VRFs

A solution SHOULD scale up to thousands of PEs having multicast service enabled.

4.2.6. Number of streams sourced

Survey responses led us to retain the following orders of magnitude for the number of streams that a solution SHOULD support:

per VPN: hundreds or thousands of streams

per PE: hundreds of streams

5. Requirements for supporting IP multicast within L3 PPVPNs

Again, the aim of this document is not to specify solutions but to give requirements for supporting IP multicast within L3 PPVPNs.

In order to list these requirements we have taken the standpoint of two different important entities: the end user (the customer using the VPN) and the service provider.

In the rest of the document, by "a solution" or "a multicast VPN solution", we mean a solution that allows multicast in an L3 provider-provisioned VPN, and which addresses the requirements listed in this document.

5.1. End user/customer standpoint

5.1.1. Service definition

As for unicast, the multicast service MUST be provider provisioned and SHALL NOT require customer devices (CEs) to support any extra features compared to those required for multicast in a non-VPN context. Enabling a VPN for multicast support SHOULD be possible with no (or very limited impact) on existing multicast protocols possibly already deployed on the CE devices.

5.1.2. CE-PE Multicast routing and group management protocols

Consequently to Section 5.1.1, multicast-related protocol exchanges between a CE and its directly connected PE SHOULD happen via existing multicast protocols.

Such protocols include: PIM-SM [RFC4601], bidirectional-PIM [I-D.ietf-pim-bidir], PIM-DM [RFC3973], and IGMPv3 [RFC3376] (this version implicitly supports hosts that only implements IGMPv1 [RFC1112] or IGMPv2 [RFC2236]).

Among those protocols, the support of PIM-SM (which includes the SSM model) and either IGMPv3 (for IPv4 solutions) and / or MLDv2 [RFC3810] (for IPv6 solutions) is REQUIRED. Bidir-PIM Support at the PE-CE interface is RECOMMENDED. And considering deployments, PIM-DM is considered as OPTIONAL.

When a multicast VPN solution is built on a VPN solution supporting IPv6 unicast, it MUST also support v6 variants of the above protocols, including MLDv2, and PIM-SM IPv6 specific procedures. For a multicast VPN solution built on a unicast VPN solution supporting only IPv4, it is RECOMMENDED that the design favors the definition of procedures and encodings that will provide an easy adaptation to

IPv6.

5.1.3. Quality of Service (QoS)

Firstly, general considerations regarding QoS in L3VPNs expressed in section 5.5 of [RFC4031] are also relevant to this section.

QoS is measured in terms of delay, jitter, packet loss, and availability. These metrics are already defined for the current unicast PPVPN services, and are included in Service Level Agreements (SLAs). In some cases, the agreed SLA may be different between unicast and multicast, and that will require differentiation mechanisms in order to monitor both SLAs.

The level of availability for the multicast service SHOULD be on par with what exists for unicast traffic. For instance comparable traffic protection mechanisms SHOULD be available for customer multicast traffic when it is carried over the service provider's network.

A multicast VPN solution SHALL allow a service provider to define at least the same level of quality of service as exists for unicast, and as exists for multicast in a non-VPN context. From this perspective, the deployment of multicast-based services within an L3VPN environment SHALL benefit from DiffServ [RFC2475] mechanisms that include multicast traffic identification, classification and marking capabilities, as well as multicast traffic policing, scheduling and conditioning capabilities. Such capabilities MUST therefore be supported by any participating device in the establishment and the maintenance of the multicast distribution tunnel within the VPN.

As multicast is often used to deliver high quality services such as TV broadcast, a multicast VPN solution MAY provide additional features to support high QoS such as bandwidth reservation and admission control.

Also, considering that multicast reception is receiver-triggered, group join delay (as defined in [RFC2432]) is also considered one important QoS parameter. It is thus RECOMMENDED that a multicast VPN solution be designed appropriately in this regard.

The group leave delay (as defined in [RFC2432]) may also be important on the CE-PE link for some usage scenarios: in cases where the typical bandwidth of multicast streams is close to the bandwidth of a PE-CE link, it will be important to have the ability to stop the emission of a stream on the PE-CE link as soon as it stops being requested by the CE, to allow for fast switching between two

different high throughput multicast streams. This implies that it

SHOULD be possible to tune the multicast routing or group management protocols (e.g. IGMP/MLD or PIM) used on the PE-CE adjacency to reduce the group leave delay to the minimum.

Lastly, a multicast VPN solution SHOULD as much as possible ensure that client multicast traffic packets are neither lost nor duplicated, even when changes occur in the way a client multicast data stream is carried over the provider network. Packet loss issues have also to be considered when a new source starts to send traffic to a group: any receiver interested in receiving such traffic SHOULD be serviced accordingly.

5.1.4. Operations and Management

The requirements and definitions for operations and management of L3VPNs that are defined in [RFC4176] equally apply to multicast, and are not extensively repeated in this document. This sub-section mentions the most important guidelines and details points of particular relevance in the context of multicast in L3VPNs.

A multicast VPN solution SHOULD allow a multicast VPN customer to manage the capabilities and characteristics of their multicast VPN services.

A multicast VPN solution MUST support SLA monitoring capabilities, which SHOULD rely upon techniques similar to those used for the unicast service for the same monitoring purposes. Multicast SLA-related metrics SHOULD be available through means similar to the ones already used for unicast-related monitoring, such as SNMP[RFC3411] or IPFIX[I-D.ietf-ipfix-protocol].

Multicast specific characteristics that may be monitored include: multicast statistics per stream, end-to-end delay, group join/leave delay (time to start/stop receiving a multicast group's traffic across the VPN, as defined in [RFC2432], Section 3).

The monitoring of multicast specific parameters and statistics MUST include multicast traffic statistics: total/incoming/outgoing/dropped traffic, by period of time ; and MAY include IP Performance Metrics related information (IPPM, [RFC2330]) that is relevant to the

multicast traffic usage: such information includes the one-way packet delay, the inter-packet delay variation, etc ([I-D.ietf-ippm-multimetrics]).

A generic discussion of SLAs is provided in [RFC3809].

Apart from statistics on multicast traffic, customers of a multicast VPN will need information concerning the status of their multicast

resource usage (multicast routing states and bandwidth). Indeed, as mentioned in Section 5.2.5, for scalability purposes, a service provider may limit the number (and/or throughput) of multicast streams that are received/sent to/from a client site. In such a case, a multicast VPN solution SHOULD allow customers to find out their current resource usage (multicast routing states and throughput), and to receive some kind of feedback if their usage exceeds the agreed bounds. Whether this issue will be better handled at the protocol level at the PE-CE interface or at the Service Management Level interface [RFC4176], is left for further discussion.

It is RECOMMENDED that any OAM mechanism designed to trigger alarms in relation to performance or resource usage metrics, integrate the ability to limit the rate at which such alarms are generated (e.g. some form of an hysteresis mecanism based on low/high thresholds defined for the metrics).

5.1.5. Security Requirements

Security is a key point for a customer who uses subscribes to a VPN service. For instance, the [RFC4364] model offers some guarantees concerning the security level of data transmission within the VPN.

A multicast VPN solution MUST provide an architecture with the same level of security for both unicast and multicast traffic.

Moreover, the activation of multicast features SHOULD be possible:

- per VRF / per VR
- per CE interface (when multiple CEs of a VPN are connected to a common VRF/VR)

- per multicast group and/or per channel

- with a distinction between multicast reception and emission

A multicast VPN solution may choose to make the optimality/ scalability trade-off stated in Section 3.3 by sometimes distributing multicast traffic of a client group to a larger set of PE routers that may include PEs which are not part of the VPN. From a security standpoint, this may be a problem for some VPN customers, thus a multicast VPN solution using such a scheme MAY offer ways to avoid this for specific customers (and/or specific customer multicast streams).

5.1.6. Extranet

In current PP L3VPN models, a customer site may be setup to be part of multiple VPNs and this should still be possible when a VPN is multicast-enabled. In practice it means that a VRF or VR can be part of more than one VPN.

A multicast VPN solution MUST support such deployments.

For instance, it must be possible to configure a VRF so that an enterprise site participating in a BGP/MPLS multicast-enabled VPN and connected to that VRF, can receive a multicast stream from, [or originate a multicast stream towards], another VPN that would be associated to that VRF.

This means that a multicast VPN solution MUST offer means for a VRF to be configured so that multicast connectivity can be setup for a chosen set of extranet VPNs. More precisely, it MUST be possible to configure a VRF so that:

- receivers behind attached CEs can receive multicast traffic sourced in the configured set of extranet VPNs

- sources behind attached CEs can reach multicast traffic receivers located in the configured set of extranet VPNs

- o multicast reception and emission can be independently enabled for each of the extranet VPNs

Moreover, a solution MUST allow service providers to control an extranet's multicast connectivity independently from the extranet's unicast connectivity. More specifically:

- o enabling unicast connectivity to another VPN MUST be possible without activating multicast connectivity with that VPN

- o enabling multicast connectivity with another VPN SHOULD NOT require more than the strict minimal unicast routing : sending multicast to a VPN SHOULD NOT require having unicast routes to that VPN, receiving multicast from a VPN SHOULD be possible with nothing more than unicast routes to the relevant multicast sources of that VPN

- o when unicast routes from another VPN are imported into a VR/VRF, for multicast RPF resolution, this SHOULD be possible without making those routes available for unicast routing

Proper support for this feature SHOULD NOT require replicating

multicast traffic on a PE-CE link, whether it is a physical or logical link.

5.1.7. Internet Multicast

Connectivity with Internet Multicast is a particular case of the previous section, where sites attached to a VR/VRF would need to receive/send multicast traffic from/to the Internet.

This should be considered OPTIONAL given the additional considerations, such as security, needed to fulfill the requirements for providing Internet Multicast.

5.1.8. Carrier's carrier

Many L3 PPVPN solutions, such as [RFC4364] and [VRs] define the "Carrier's Carrier" model, where a "carrier's carrier" service provider supports one or more customer ISP, or "sub-carriers". A multicast VPN solution SHOULD support the carrier's carrier model in

a scalable and efficient manner.

Ideally the range of tunneling protocols available for the sub-carrier ISP should be the same as those available for the carrier's carrier ISP. This implies that the protocols that may be used at the PE-CE level SHOULD NOT be restricted to protocols required as per Section 5.1.2 and SHOULD include some of the protocols listed in Section 5.2.4, such as for instance P2MP MPLS signaling protocols.

In the context of MPLS-based L3VPN deployments, such as BGP/MPLS VPNs [RFC4364], this means that MPLS label distribution SHOULD happen at the PE-CE level, giving the ability to the sub-carrier to use multipoint LSPs as a tunneling mechanism.

5.1.9. Multi-homing, load balancing and resiliency

A multicast VPN solution SHOULD be compatible with current solutions that aim at improving the service robustness for customers such as multi-homing, CE-PE link load balancing and fail-over. A multicast VPN solution SHOULD also be able to offer those same features for multicast traffic.

Any solution SHOULD support redundant topology of CE-PE links. It SHOULD minimize multicast traffic disruption and fail-over.

5.1.10. RP Engineering

When PIM-SM (or bidir-PIM) is used in ASM mode on the VPN customer side, the RP function (or RP-address in the case of bidir-PIM) has to

be associated to a node running PIM, and configured on this node.

5.1.10.1. RP Outsourcing

In the case of PIM-SM in ASM mode, engineering of the RP function requires the deployment of specific protocols and associated configurations. A service provider may offer to manage customers' multicast protocol operation on their behalf. This implies that it is necessary to consider cases where a customer's RPs are out-sourced (e.g., on PEs). Consequently, a VPN solution MAY support the hosting of the RP function in a VR or VRF.

5.1.10.2. RP Availability

Availability of the RP function (or address) is required for proper operation of PIM-SM (ASM mode) and bidir-PIM. Loss of connectivity to the RP from a receiver or source will impact the multicast service. For this reason different mechanisms exist, such as BSR [I-D.ietf-pim-sm-bsr] or anycast-RP (MSDP based [RFC3446] or PIM based [RFC4610]).

These protocols and procedures SHOULD work transparently through a multicast VPN, and MAY if relevant, be implemented in a VRF/VR.

Moreover, a multicast VPN solution MAY improve the robustness of the ASM multicast service regarding loss of connectivity to the RP, by providing specific features that help :

a) maintain ASM multicast service among all the sites within an MVPN that maintain connectivity among themselves, even when the site(s) hosting the RP lose their connectivity to the MVPN

b) maintain ASM multicast service within any site that loses connectivity to the service provider

5.1.10.3. RP Location

In the case of PIM-SM, when a source starts to emit traffic toward a group (in ASM mode), if sources and receivers are located in VPN sites that are different than that of the RP, then traffic may transiently flow twice through the SP network and the CE-PE link of the RP (from source to RP, and then from RP to receivers). This traffic peak, even short, may not be convenient depending on the traffic and link bandwidth.

Thus, a VPN solution MAY provide features that solve or help mitigate this potential issue.

5.1.11. Addressing

A multicast provider-provisioned L3VPN SHOULD NOT impose restrictions on multicast group addresses used by VPN customers.

In particular, like unicast traffic, an overlap of multicast group address sets used by different VPN customers MUST be supported.

The use of globally unique means of multicast-based service identification at the scale of the domain where such services are provided SHOULD be recommended. For IPv4 multicast, this implies the use of the multicast administratively scoped range, (239/8 as defined by [RFC2365]) for services which are to be used only inside the VPN, and of either SSM-range addresses (232/8 as defined by [RFC4607]) or globally assigned group addresses (e.g. GLOP [RFC3180], 233/8) for services for which traffic may be transmitted outside the VPN.

5.1.12. Minimum MTU

For customers, it is often a serious issue whether transmitted packets will be fragmented or not. In particular, some multicast applications might have different requirements than those that make use of unicast, and they may expect services that guarantee available packet length not to be fragmented.

Therefore, a multicast VPN solution SHOULD be designed with these considerations in mind. In practice:

- the encapsulation overhead of a multicast VPN solution SHOULD be minimized, so that customer devices can be free of fragmentation and reassembly activity as much as possible

- a multicast VPN solution SHOULD enable the service provider to commit to a minimum path MTU usable by multicast VPN customers

- a multicast VPN solution SHOULD be compatible with path MTU discovery mechanisms (see [RFC1191] and [RFC4459]), and particular care SHOULD be given to means to help troubleshoot MTU issues

Moreover, since Ethernet LAN segments are often located at first and last hops, a multicast VPN solution SHOULD be designed to allow for a minimum 1500 byte IP MTU for VPN customers multicast packet, when the provider backbone design allows it.

5.2. Service provider standpoint

Note: To avoid repetition and confusion with terms used in solution specifications, we introduced in Section 2.1 the term MDTunnel (for

Multicast Distribution Tunnel), which designates the data plane means used by the service provider to forward customer multicast traffic over the core network.

5.2.1. General requirement

The deployment of a multicast VPN solution SHOULD be possible with no (or very limited) impact on existing deployments of standardized multicast related protocols on P and PE routers.

5.2.2. Scalability

Some currently standardized and deployed L3VPN solutions have the major advantage of being scalable in the core regarding the number of customers and the number of customer routes. For instance, in the [RFC4364] and VRs [I-D.ietf-l3vpn-vpn-vr] models, a P router sees a number of MPLS tunnels that is only linked to the number of PEs and not to the number of VPNs, or customer sites.

As far as possible, this independence in the core, with respect to the number of customers and to customer activity, is recommended. Yet, it is recognized that in our context scalability and resource usage optimality are competing goals, so this requirement may be reduced to giving the possibility of bounding the quantity of states that the service provider needs to maintain in the core for MDTunnels, with a bound being independent of the multicast activity of VPN customers.

It is expected that multicast VPN solutions will use some kind of point-to-multipoint technology to efficiently carry multicast VPN traffic, and because such technologies require maintaining state information, this will use resources in the control plane of P and PE routers (memory and processing, and possibly address space).

Scalability is a key requirement for multicast VPN solutions. Solutions MUST be designed to scale well with an increase in the number of any of the following:

- the number of PEs
- the number of customer VPNs (total and per PE)
- the number of PEs and sites in any VPN
- the number of client multicast channels (groups or source-groups)

Please consult section 4.2 for typical orders of magnitude up to which a multicast VPN solution is expected to scale

Scalability of both performance and operation MUST be considered.

Key considerations SHOULD include:

- the processing resources required by the control plane (neighborhood or session maintenance messages, keep-alives, timers, etc.)
- the memory resources needed for the control plane
- the amount of protocol information transmitted to manage a multicast VPN (e.g. signaling throughput)
- the amount of control plane processing required on PE and P routers to add or remove a customer site (or a customer from a multicast session)
- the number of multicast IP addresses used (if IP multicast in ASM mode is proposed as a multicast distribution tunnel)
- other particular elements inherent to each solution that impact scalability (e.g., if a solution uses some distribution tree inside the core, topology of the tree and number of leaf nodes may be some of them)

It is expected that the applicability of each solution will be evaluated with regards to the aforementioned scalability criteria.

These considerations naturally lead us to believe that proposed solutions SHOULD offer the possibility of sharing such resources between different multicast streams (between different VPNs, between different multicast streams of the same or of different VPNs). This means for instance, if MDTunnels are trees, being able to share an MDTunnel between several customers.

Those scalability issues are expected to be more significant on P routers, but a multicast VPN solution SHOULD address both P and PE routers as far as scalability is concerned.

5.2.3. Resource optimization

5.2.3.1. General goals

One of the aims of the use of multicast instead of unicast is resource optimization in the network.

The two obvious suboptimal behaviors that a multicast VPN solution would want to avoid are needless duplication (when the same data

travels twice or more on a link, e.g. when doing ingress PE replication) and needless reception (e.g. a PE receiving traffic that it does not need because there are no downstream receivers).

5.2.3.2. Trade-off and tuning

As previously stated in this document, designing a scalable solution that makes an optimal use of resources is considered difficult. Thus what is expected from a multicast VPN solution is that it addresses the resource optimization issue while taking into account the fact that some trade-off has to be made.

Moreover, it seems that a "one size fits all" trade-off probably does not exist either. Thus a multicast VPN solution SHOULD offer service providers appropriate configuration settings that let them tune the trade-off according to their particular constraints (network topology, platforms, customer applications, level of service offered etc.).

As an illustration here are some example bounds of the trade-off space:

Bandwidth optimization: setting up optimized core MDTunnels whose topology (PIM or P2MP LSP trees, etc.) precisely follows a customer's multicast routing changes. This requires managing a large amount of state in the core, and also quick reactions of the core to customer multicast routing changes. This approach can be advantageous in terms of bandwidth, but it is poor in terms of state management.

State optimization: setting up MDTunnels that aggregate multiple customer multicast streams (all or some of them, across different VPNs or not). This will have better scalability properties, but at the expense of bandwidth since some MDTunnel leaves will very likely receive traffic they don't need, and because increased

constraints will make it harder to find optimal MDTunnels.

5.2.3.3. Traffic engineering

If the VPN service provides traffic engineering features for the connection used between PEs for unicast traffic in the VPN service, the solution SHOULD provide equivalent features for multicast traffic.

A solution SHOULD offer means to support key TE objectives as defined in [RFC3272], for the multicast service.

A solution MAY also usefully support means to address multicast-

specific traffic engineering issues: it is known that bandwidth resource optimization in the point-to-multipoint case is an NP-hard problem, and that techniques used for unicast TE may not be applicable to multicast traffic.

Also, it has been identified that managing the trade-off between resource usage and scalability may incur uselessly sending traffic to some PEs participating in a multicast VPN. For this reason, a multicast VPN solution MAY permit that the bandwidth/state tuning take into account the relative cost or availability of bandwidth toward each PE.

5.2.4. Tunneling Requirements

5.2.4.1. Tunneling technologies

Following the principle of separation between the control plane and the forwarding plane, a multicast VPN solution SHOULD be designed so that control and forwarding planes are not interdependent: the control plane SHALL NOT depend on which forwarding plane is used (and vice versa), and the choice of forwarding plane SHOULD NOT be limited by the design of the solution. Also, the solution SHOULD NOT be tied to a specific tunneling technology.

In a multicast VPN solution extending a unicast L3 PPVPN solution, consistency in the tunneling technology has to be favored: such a solution SHOULD allow the use of the same tunneling technology for multicast as for unicast. Deployment consistency, ease of operation

and potential migrations are the main motivations behind this requirement.

For MDTunnels, a solution SHOULD be able to use a range of tunneling technologies, including point-to-point and point-to-multipoint, such as GRE [RFC2784] (including GRE in multicast IP trees), MPLS [RFC3031] (including P2P or MP2P tunnels, and multipoint tunnels signaled with MPLS P2MP extensions to RSVP [I-D.ietf-mpls-rsvp-te-p2mp] or LDP [I-D.ietf-mpls-mp-ldp-reqs][I-D.ietf-mpls-ldp-p2mp]), L2TP (including L2TP for multicast [RFC4045]), IPsec [RFC4031], IP-in-IP [RFC2003], etc.

Naturally, it is RECOMMENDED that a solution is built so that it can leverage the point to multipoint variants of these techniques, that allow for packet replications to happen along a tree in the provider core network, and may help improve bandwidth efficiency in a multicast VPN context.

5.2.4.2. MTU and Fragmentation

A solution SHOULD support a method that provides the minimum MTU of the MDTunnel (e.g., to discover MTU, to communicate MTU via signaling, etc.) so that:

- fragmentation inside the MDTunnel does not happen, even when allowed by the underlying tunneling technology
- proper troubleshooting can be performed if packets that are too big for the MDTunnel happen to be encapsulated in the MDTunnel

5.2.5. Control mechanisms

The solution MUST provide some mechanisms to control the sources within a VPN. This control includes the number of sources that are entitled to send traffic on the VPN, and/or the total bit rate of all the sources.

At the reception level, the solution MUST also provide mechanisms to control the number of multicast groups or channels VPN users are

entitled to subscribe to and/or the total bit rate represented by the corresponding multicast traffic.

All these mechanisms MUST be configurable by the service provider in order to control the amount of multicast traffic and state within a VPN.

Moreover it MAY be desirable to be able to impose some bound on the quantity of state used by a VPN in the core network for its multicast traffic, whether on each P or PE router, or globally. The motivation is that it may be needed to avoid out-of-resources situations (e.g. out of memory to maintain PIM state if IP multicast is used in the core for multicast VPN traffic, or out of memory to maintain RSVP state if MPLS P2MP is used, etc.).

5.2.6. Support of Inter-AS, inter-provider deployments

A solution MUST support inter-AS multicast VPNs, and SHOULD support inter-provider multicast VPNs. Considerations about coexistence with unicast inter-AS VPN Options A, B and C (as described in section 10 of [RFC4364]) are strongly encouraged.

A multicast VPN solution SHOULD provide inter-AS mechanisms requiring the least possible coordination between providers, and keep the need for detailed knowledge of providers' networks to a minimum - all this being in comparison with corresponding unicast VPN options.

- Within each service provider the service provider SHOULD be able on its own to pick the most appropriate tunneling mechanism to carry (multicast) traffic among PEs (just like what is done today for unicast)

- If a solution does require a single tunnel to span P routers in multiple ASs, the solution SHOULD provide mechanisms to ensure that the inter-provider co-ordination to setup such a tunnel is minimized

Moreover such support SHOULD be possible without compromising other requirements expressed in this requirement document, and SHALL NOT incur penalties on scalability and bandwidth-related efficiency.

5.2.7. Quality of Service Differentiation

A multicast VPN solution SHOULD give a VPN service provider the ability to offer, guarantee and enforce differentiated levels of QoS for its different customers.

5.2.8. Infrastructure security

The solution SHOULD provide the same level of security for the service provider as what currently exists for unicast VPNs (for instance, as developed in the Security sections of [RFC4364] and [I-D.ietf-l3vpn-vpn-vr]). For instance, traffic segregation and intrinsic protection against DOS and DDOS attacks of the BGP/MPLS VPN solution must be supported by the multicast solution.

Moreover, since multicast traffic and routing are intrinsically dynamic (receiver-initiated), some mechanism SHOULD be proposed so that the frequency of changes in the way client traffic is carried over the core can be bounded and not tightly coupled to dynamic changes of multicast traffic in the customer network. For example, multicast route dampening functions would be one possible mechanism.

Network devices that participate in the deployment and the maintenance of a given L3VPN MAY represent a superset of the participating devices that are also involved in the establishment and the maintenance of the multicast distribution tunnels. As such the activation of IP multicast capabilities within a VPN SHOULD be device-specific, not only to make sure that only the relevant devices will be multicast-enabled, but also to make sure that multicast (routing) information will be disseminated to the multicast-enabled devices only, hence limiting the risk of multicast-inferred DOS attacks.

Traffic of a multicast channel for which there are no members in a

given multicast VPN MUST NOT be propagated within the multicast VPN, most particularly if the traffic comes from another VPN or from the Internet.

Security considerations are particularly important for inter-AS and inter-provider deployments. In such cases, it is RECOMMENDED that a multicast VPN solution support means to ensure the integrity and

authenticity of multicast-related exchanges across inter-AS or inter-provider borders. It is RECOMMENDED that corresponding procedures require the least possible coordination between providers; more precisely, when specific configurations or cryptographic keys have to be deployed, this shall be limited to ASBRs (Autonomous System Border Routers) or a subset of them, and optionally BGP Route Reflectors (or a subset of them).

Lastly, control mechanisms described in Section 5.2.5 are also to be considered from this infrastructure security point of view.

5.2.9. Robustness

Resiliency is also crucial to infrastructure security, thus a multicast VPN solution SHOULD either avoid single points of failures or propose some technical solution making it possible to implement a fail-over mechanism.

As an illustration, one can consider the case of a solution that would use PIM-SM as a means to setup MDTunnels. In such a case, the PIM RP might be a single point of failure. Such a solution SHOULD be compatible with a solution implementing RP resiliency, such as anycast-RP [RFC4610] or BSR [I-D.ietf-pim-sm-bsr].

5.2.10. Operation, Administration and Maintenance

The operation of a multicast VPN solution SHALL be as light as possible and providing automatic configuration and discovery SHOULD be a priority when designing a multicast VPN solution. Particularly the operational burden of setting up multicast on a PE or for a VR/VRF SHOULD be as low as possible.

Also, as far as possible, the design of a solution SHOULD carefully consider the number of protocols within the core network: if any additional protocols are introduced compared with the unicast VPN service, the balance between their advantage and operational burden SHOULD be examined thoroughly.

Moreover, monitoring of multicast specific parameters and statistics SHOULD be offered to the service provider, following the requirements expressed in [RFC4176].

Most notably the provider SHOULD have access to:

- Multicast traffic statistics (incoming/outgoing/dropped/total traffic conveyed, by period of time)

- Information about client multicast resource usage (multicast routing state and bandwidth usage)

- Alarms when limits are reached on such resources

- The IPPM (IP Performance Metrics [RFC2330]) -related information that is relevant to the multicast traffic usage: such information includes the one-way packet delay, the inter-packet delay variation, etc.

- Statistics on decisions related to how client traffic is carried on distribution tunnels (e.g. "traffic switched onto a multicast tree dedicated to such groups or channels")

- Statistics on parameters that could help the provider to evaluate its optimality/state trade-off

This information SHOULD be made available through standardized protocols such as SNMP [RFC3411] MIBs (Management Information Bases) or IPFIX [I-D.ietf-ipfix-protocol]. For instance, in the context of BGP/MPLS VPNs [RFC4364], multicast extensions to MIBs defined in [RFC4382] SHOULD be proposed, with proper integration with [RFC3811], [RFC3812], [RFC3813] and [RFC3814] when applicable.

Mechanisms similar to those described in Section 5.2.12 SHOULD also exist for proactive monitoring of the MDTunnels.

Proposed OAM mechanisms and procedures for multicast VPNs SHOULD be scalable with respect to the parameters mentioned in Section 5.2.2. In particular, it is RECOMMENDED that particular attention is given to the impact of monitoring mechanisms on performances and QoS.

Moreover, it is RECOMMENDED that any OAM mechanism designed to trigger alarms in relation to performance or resource usage metrics, integrate the ability to limit the rate at which such alarms are generated (e.g. some form of an hysteresis mecanism based on low/high thresholds defined for the metrics).

5.2.11. Compatibility and migration issues

It is a requirement that unicast and multicast services MUST be able to co-exist within the same VPN.

Likewise, a multicast VPN solution SHOULD be designed so that its activation in devices that participate in the deployment and the maintenance of a multicast VPN SHOULD be as smooth as possible, i.e. without affecting the overall quality of the services that are already supported by the underlying infrastructure.

A multicast VPN solution SHOULD prevent compatibility and migration issues, for instance by focusing on providing mechanisms facilitating forward compatibility. Most notably a solution supporting only a subset of the requirements expressed in this document SHOULD be designed to allow compatibility to be introduced in further revisions.

It SHOULD be an aim of any multicast VPN solution to offer as much backward compatibility as possible. Ideally a solution would have the ability to offer multicast VPN services across a network containing some legacy routers that do not support any multicast VPN specific features.

In any case a solution SHOULD state a migration policy from possibly existing deployments.

5.2.12. Troubleshooting

A multicast VPN solution that dynamically adapts the way some client multicast traffic is carried over the provider's network may incur the disadvantage of being hard to troubleshoot. In such a case, to help diagnose multicast network issues, a multicast VPN solution SHOULD provide monitoring information describing how client traffic is carried over the network (e.g. if a solution uses multicast-based MDTunnels, which provider multicast group is used for a given client multicast stream). A solution MAY also provide configuration options to avoid any dynamic changes, for multicast traffic of a particular VPN or a particular multicast stream.

Moreover, a solution MAY provide mechanisms that allows network operators to check that all VPN sites that advertised interest in a particular customer multicast stream are properly associated with the corresponding MDTunnel. Providing operators with means to check the proper setup and operation of MDTunnels MAY also be provided (e.g. when P2MP MPLS is used for MDTunnels, troubleshooting functionalities SHOULD integrate mechanisms compliant with [RFC4687], such as LSPPing [RFC4379][I-D.ietf-mpls-p2mp-lsp-ping]). Depending on the

implementation such verification could be initiated by a source-PE or e receiver-PE.

6. Security Considerations

This document does not by itself raise any particular security issue.

A set of security issues have been identified that MUST be addressed when considering the design and deployment of multicast-enabled L3 PP VPNs. Such issues have been described in Section 5.1.5 and Section 5.2.8.

7. IANA Considerations

This document has no actions for IANA.

8. Contributors

The main contributors to this document are listed below, in alphabetical order:

- Christian Jacquenet
  France Telecom
  3, avenue Francois Chateau
  CS 36901 35069 RENNES Cedex, France
  Email: christian.jacquenet@francetelecom.com

- 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

- Jean-Louis Le Roux
  France Telecom R & D
  2, avenue Pierre-Marzin
  22307 Lannion Cedex, France
  Email: jeanlouis.leroux@francetelecom.com

- Nicolai Leymann
  T-Systems International GmbH
  Engineering Networks, Products & Services
  Goslarer Ufer 3510589 Berlin, Germany
  Email: nicolai.leymann@t-systems.com

- Renaud Moignard
  France Telecom R & D
  2, avenue Pierre-Marzin
  22307 Lannion Cedex, France
  Email: renaud.moignard@francetelecom.com

- Thomas Morin
  France Telecom R & D
  2, avenue Pierre-Marzin
  22307 Lannion Cedex, France
  Email: thomas.morin@francetelecom.com

9. Acknowledgments

The authors would like to thank, by rough chronological order, Vincent Parfait, Zubair Ahmad, Elodie Hemon-Larreur, Sebastien Loye, Rahul Aggarwal, Hitoshi Fukuda, Luyuan Fang, Adrian Farrel, Daniel King, Yiqun Cai, Ronald Bonica, Len Nieman, Satoru Matsushima, Netzahualcoyotl Ornelas, Yakov Rekhter, Marshall Eubanks, Pekka Savola, Benjamin Niven-Jenkins, Thomas Nadeau, for their review, valuable input and feedback.

We also thank the people who kindly answered the survey, and Daniel King who took care of gathering and anonymizing its results.

10. References

10.1. Normative references

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4031] Carugi, M. and D. McDysan, "Service Requirements for Layer

3 Provider Provisioned Virtual Private Networks (PPVPNs)", RFC 4031, April 2005.

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[RFC4176] El Mghazli, Y., Nadeau, T., Boucadair, M., Chan, K., and A. Gonguet, "Framework for Layer 3 Virtual Private Networks (L3VPN) Operations and Management", RFC 4176, October 2005.

[RFC3973] Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", RFC 3973, January 2005.

10.2. Informative references

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[I-D.ietf-l3vpn-vpn-vr]
Ould-Brahim, H., "Network based IP VPN Architecture Using Virtual Routers", draft-ietf-l3vpn-vpn-vr-03 (work in progress), March 2006.

[RFC2432] Dubray, K., "Terminology for IP Multicast Benchmarking",

RFC 2432, October 1998.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.

[RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.

[I-D.ietf-mpls-rsvp-te-p2mp]
Aggarwal, R., "Extensions to RSVP-TE for Point-to-Multipoint TE LSPs", draft-ietf-mpls-rsvp-te-p2mp-06 (work in progress), August 2006.

[I-D.ietf-pim-sm-bsr]
Bhaskar, N., "Bootstrap Router (BSR) Mechanism for PIM", draft-ietf-pim-sm-bsr-09 (work in progress), June 2006.

[RFC4610] Farinacci, D. and Y. Cai, "Anycast-RP Using Protocol Independent Multicast (PIM)", RFC 4610, August 2006.

[RFC3446] Kim, D., Meyer, D., Kilmer, H., and D. Farinacci, "Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)", RFC 3446, January 2003.

[I-D.ietf-mpls-ldp-p2mp]
Minei, I., "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", draft-ietf-mpls-ldp-p2mp-02 (work in progress), October 2006.

[I-D.ietf-mpls-mp-ldp-reqs]
Roux, J., "Requirements for point-to-multipoint extensions to the Label Distribution Protocol", draft-ietf-mpls-mp-ldp-reqs-01 (work in progress), June 2006.

[RFC4687] Yasukawa, S., Farrel, A., King, D., and T. Nadeau, "Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks", RFC 4687, September 2006.

[I-D.ietf-pim-bidir]
Handley, M., "Bi-directional Protocol Independent Multicast (BIDIR-PIM)", draft-ietf-pim-bidir-08 (work in progress), October 2005.

[RFC2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.

[RFC3353] Ooms, D., Sales, B., Livens, W., Acharya, A., Griffoul, F., and F. Ansari, "Overview of IP Multicast in a Multi-Protocol Label Switching (MPLS) Environment", RFC 3353, August 2002.

[RFC3272] Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X. Xiao, "Overview and Principles of Internet Traffic Engineering", RFC 3272, May 2002.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.

[I-D.ietf-ipfix-protocol] Claise, B., "Specification of the IPFIX Protocol for the Exchange", draft-ietf-ipfix-protocol-24 (work in progress), November 2006.

[RFC4045] Bourdon, G., "Extensions to Support Efficient Carrying of Multicast Traffic in Layer-2 Tunneling Protocol (L2TP)", RFC 4045, April 2005.

[RFC3809] Nagarajan, A., "Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)", RFC 3809, June 2004.

[RFC3811] Nadeau, T. and J. Cucchiara, "Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management", RFC 3811, June 2004.

[RFC3812] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)", RFC 3812, June 2004.

[RFC3813] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)", RFC 3813, June 2004.

[RFC3814] Nadeau, T., Srinivasan, C., and A. Viswanathan, "Multiprotocol Label Switching (MPLS) Forwarding Equivalence Class To Next Hop Label Forwarding Entry (FEC-

To-NHLFE) Management Information Base (MIB)", RFC 3814, June 2004.

[RFC2365] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, July 1998.

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.

[I-D.ietf-ippm-multimetrics] Stephan, E., "IP Performance Metrics (IPPM) for spatial and multicast", draft-ietf-ippm-multimetrics-02 (work in progress), October 2006.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[RFC3180] Meyer, D. and P. Lothberg, "GLOP Addressing in 233/8", BCP 53, RFC 3180, September 2001.

[RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, December 2002.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, November 1990.

[RFC4382] Nadeau, T. and H. van der Linde, "MPLS/BGP Layer 3 Virtual Private Network (VPN) Management Information Base", RFC 4382, February 2006.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

[I-D.ietf-mpls-p2mp-lsp-ping] Farrel, A. and S. Yasukawa, "Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol", draft-ietf-mpls-p2mp-lsp-ping-02 (work in progress),

September 2006.

[RFC4459] Savola, P., "MTU and Fragmentation Issues with In-the-Network Tunneling", RFC 4459, April 2006.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

Appendix A. Changelog

This section lists changes made to this document (minor or editorial changes excepted) between major revisions.

It shall be removed before publication as an RFC.

A.1. Changes between -00 and -01

- integrated comments made on L3VPN WG mailing list after -00 submission
- completed Carrier's carrier section (5.1.9)
- updates in sections 5.1 and 5.2 about minimum MTU
- added a section about "Quality of Service Differentiation" as ISP requirement (section 5.2.5)
- added P2MP LDP extensions as possible MDTunnels techniques (section 5.2.3.1)
- started to build section 4 "Use Case"
- detailed section 5.1.3 "QoS", most notably about group join and leave delays
- additions to section 5.2.12 "Inter-AS, inter-provider"
- added MDTunnel verification requirement to section 5.2.11
- moved "Architectural Considerations" section

- moved contributors to top of document
- made draft content agnostic to unicast L3VPN solutions
- added two appendixes: "Changelog" and "Requirement summary"
- conversion to XML [RFC2629] with the help of some scripting and Bill Fenner's xml2rfc XMLMind plugin
- lot's of editorial changes

A.2. Changes between -01 and -02

- based on survey results:
  - restructure use case scenario section
  - fill in Scalability orders of magnitude section
  - better detail requirements for protocols at the PE-CE level
  - add considerations about PEs with scarce connectivity to section 5.2.3.3
  - step up requirement level for Extranet (Section 5.1.6)
- some editorial changes
- use capitalized wording for some requirements
- fill in requirements summary

A.3. Changes between -02 and -03

- made inter-AS a MUST (and moved the whole section up)

- add a requirement about security of multicast-related exchanges across providers/ASes, in Section 5.2.8

- some editorial changes and fixed typos

A.4. Changes between -03 and -04

- Integrated comments received during last call

- Lots of editorial comments

- Improved terminology section

- Number of VPNs with multicast enabled: A solution SHOULD NOT limit the proportion of multicast VPNs among all (unicast) VPNs.

- Customer-side service definition Enabling a VPN for multicast support SHOULD be possible with no (or very limited impact) on existing multicast protocols possibly already deployed on the CE devices

- Extranet: a solution MUST allow to control an extranet multicast connectivity independently from the extranet unicast connectivity

- Service provider standpoint, added a general statement: "The deployment of a multicast VPN solution SHOULD be possible with no (or very limited) impact on possibly existing deployments of multicast protocols on P and PE routers."

- Removed institutions and company names from the already long acknowledgments section.

A.5. Changes between -04 and -05

- Follow up of mailing list comments integration

- 2547bis is now 4364

- CE-PE Multicast routing and management protocols: clarified IPv6 related statements

- More precisions in 5.1.7. "Extranet"

- Revamped section 5.1.11 "RP Engineering" : separated text about different issues "RP Outsourcing", "RP Availability", "RP Location". Updated 5.1.10 accordingly.

- Updated 5.2.10 about proactive monitoring of tunnels

- Removed requirements summary

- Lots of editorial changes

A.6. Changes between -05 and -06

- editorial changes

- revamp OAM-related sections to integrate comments made on WG mailing-list

A.7. Changes between -06 and -08

- minor editorial changes

A.8. Changes between -08 and -09

- editorial changes integrating Last Call comments made by Ben Jenkins and Eric gray

- integrated some comment made by Stephen Farrell for the security directorate

- updated references to draft which became RFCs, and reordered references a bit

- other editorial changes

A.9. Changes between -09 and -10

- integration of comments made during IESG review : updated text on OAM and MTU issues, fixed/added references.

- editorial nits

Author's Address

Thomas Morin (editor)
France Telecom R&D

2, avenue Pierre Marzin
Lannion 22307
France

Email: thomas.morin@rd.francetelecom.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).