

Internet Engineering Task Force (IETF)  
Request for Comments: 7899  
Updates: 6514  
Category: Standards Track  
ISSN: 2070-1721

T. Morin, Ed.  
S. Litkowski  
Orange  
K. Patel  
Cisco Systems  
Z. Zhang  
R. Kebler  
J. Haas  
Juniper Networks  
June 2016

## Multicast VPN State Damping

### Abstract

This document describes procedures to damp Multicast VPN (MVPN) routing state changes and control the effect of the churn due to the multicast dynamicity in customer sites. The procedures described in this document are applicable to BGP-based multicast VPN and help avoid uncontrolled control-plane load increase in the core routing infrastructure. The new procedures proposed were inspired by BGP unicast route damping principles that have been adapted to multicast.

### Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7899>.

## Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1.	Introduction . . . . .	3
2.	Terminology . . . . .	4
3.	Overview . . . . .	4
4.	Existing Mechanisms . . . . .	5
4.1.	Rate-Limiting Multicast Control Traffic . . . . .	5
4.2.	Existing PIM, IGMP, and MLD Timers . . . . .	6
4.3.	BGP Route Damping . . . . .	6
5.	Procedures for Multicast State Damping . . . . .	7
5.1.	PIM Procedures . . . . .	7
5.2.	Procedures for Multicast VPN State Damping . . . . .	10
6.	Procedures for P-Tunnel State Damping . . . . .	12
6.1.	Damping MVPN P-Tunnel Change Events . . . . .	12
6.2.	Procedures for Ethernet VPNs . . . . .	13
7.	Operational Considerations . . . . .	13
7.1.	Enabling Multicast Damping . . . . .	13
7.2.	Troubleshooting and Monitoring . . . . .	13
7.3.	Default and Maximum Values . . . . .	13
8.	Security Considerations . . . . .	15
9.	References . . . . .	16
9.1.	Normative References . . . . .	16
9.2.	Informative References . . . . .	17
	Acknowledgements . . . . .	17
	Authors' Addresses . . . . .	18

## 1. Introduction

In a multicast VPN [RFC6513] deployed with BGP-based procedures [RFC6514], when receivers in VPN sites join and leave a given multicast group or channel through multicast membership control protocols (Internet Group Management Protocol (IGMP) [RFC3376] and Multicast Listener Discovery (MLD) [RFC3810]), multicast routing protocols accordingly adjust multicast routing states and P-multicast tree states to forward or prune multicast traffic to these receivers. Similar challenges arise in the context of the multicast specification for Virtual Private LAN Service (VPLS) [RFC7117].

In VPN contexts, providing isolation between customers of a shared infrastructure is a core requirement resulting in stringent expectations with regard to risks of denial-of-service attacks.

By nature, multicast memberships change based on the behavior of multicast applications running on end hosts. Hence, the frequency of membership changes can legitimately be much higher than the typical churn of unicast routing states.

Therefore, mechanisms need to be put in place to ensure that the load put on the BGP control plane, and on the P-tunnel setup control plane, remains under control regardless of the frequency at which multicast membership changes are made by end hosts.

This document describes procedures inspired by existing BGP route damping [RFC2439] that are aimed at offering means to set an upper bound to the amount of processing for the MVPN control-plane protocols: more precisely, the BGP control plane in [RFC6514], the P-tunnel control-plane protocol in the contexts of [RFC6514], and the multicast specification for VPLS [RFC7117]. The goal of setting this upper bound is pursued simultaneously with the goal of preserving the service provided (delivering the multicast stream as requested by Customer Edge devices), although at the expense of a minimal increase of average bandwidth use in the provider network). The upper bound to the control-plane load due to the processing of a given multicast state is controlled indirectly via configurable parameters.

Section 16 of [RFC6514] specifically spells out the need for damping the activity of C-multicast and Leaf Auto-discovery routes and outlines how to do it by "delaying the advertisement of withdrawals of C-multicast routes". This specification provides appropriate detail on how to implement this approach and how to provide control to the operator; for this reason, it is an update to [RFC6514].

The base principle of this specification is described in Section 3. Existing mechanisms that could be relied upon are discussed in Section 4. Section 5 details the procedures introduced by this specification.

Section 6 provides specific details related to the damping of multicast VPNs P-tunnel state.

Finally, Section 7 discusses operational considerations related to the proposed mechanism.

## 2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

This document reuses terminology from [RFC7761] and [RFC6514].

In this specification, damping of a multicast state will be said to be "active" or "inactive". Note that in [RFC2439], the term used for a unicast route that is dampened is "suppressed", but we will avoid this term in this specification given that the proposed solution consists in holding active a damped multicast state.

## 3. Overview

The procedures described in this document allow the network operator to configure multicast VPN PEs (Provider Edge routers) so that they can delay the propagation of multicast state prune messages between PEs when faced with a rate of multicast state dynamicity exceeding a certain configurable threshold. Assuming that the number of multicast states that can be created by a receiver is bounded, delaying the propagation of multicast state pruning results in setting up an upper bound to the average frequency at which the router will send state updates to an upstream router.

From the point of view of a downstream router, such as a CE (Customer Edge router), this approach has no impact: the multicast routing state changes that it solicits to its PE will be honored without any additional delay. Indeed, the propagation of Joins is not impacted by the procedures specified here, and having the upstream router delay state prune propagation to its own upstream router does not affect what traffic is sent to the downstream router. In particular, the amount of bandwidth used on the PE-CE link downstream to a PE applying this damping technique is not increased.

This approach increases the average bandwidth utilization on a link upstream to a PE applying this technique, such as a PE-PE link: indeed, a given multicast flow will be forwarded for a longer time than if no damping was applied. That said, it is expected that this technique will meet the goals of protecting the multicast routing infrastructure control plane without a significant average increase of bandwidth; for instance, damping events happening at a frequency higher than one event per X seconds can be done without increasing by more than X seconds the time during which a multicast flow is present on a link.

That said, simulation of the exponential decay algorithm shows that the multicast state churn can be drastically reduced without significantly increasing the duration for which multicast traffic is forwarded. Hence, using this technique will efficiently protect the multicast routing infrastructure control plane against the issues described here without a significant average increase of bandwidth. The exception will be a scenario with strict constraints on multicast bandwidth, where extending the time a multicast flow is forwarded would result in congestion.

To be practical, such a mechanism requires configurability. In particular, means are required to control when damping is triggered and to allow delaying the pruning of a multicast state for a time increasing with the churn of this multicast state. This will let the operator control the trade-off made between minimizing the dynamicity and reducing bandwidth consumption.

#### 4. Existing Mechanisms

This section describes mechanisms that could be considered to address the issue but that end up appearing as not suitable or not efficient enough.

##### 4.1. Rate-Limiting Multicast Control Traffic

The Protocol Independent Multicast - Sparse Mode (PIM-SM) specification [RFC7761] examines multicast security threats and, among other things, the risk of denial-of-service attacks described in Section 1. A mechanism relying on rate-limiting PIM messages is proposed in Section 5.3.3 of [RFC4609] but has the identified drawbacks of impacting the service delivered and having side-effects on legitimate users.

#### 4.2. Existing PIM, IGMP, and MLD Timers

In the context of PIM multicast routing protocols [RFC7761], a mechanism exists that may offer a form of de facto damping of multicast states, under some conditions. Indeed, when active, the prune override mechanism consists in having a PIM upstream router introduce a delay ("prune override interval") before taking into account a PIM Prune message sent by a downstream neighbor.

This mechanism has not been designed specifically for the purpose of damping multicast state, but as a means to allow PIM to operate on multi-access networks. See Section 4.3.3 of [RFC7761]. However, when active, this mechanism will prevent a downstream router from producing multicast routing protocol messages that would cause, for a given multicast state, the upstream router to send to its own upstream router multicast routing protocol messages at a rate higher than  $1/[\text{JP\_Override\_Interval}]$ . This provides a form of de facto damping.

Similarly, the IGMP and MLD multicast membership control protocols can provide a similar behavior under the right conditions.

These mechanisms are not considered suitable to meet the goals spelled out in Section 1, the main reasons being that:

- o when enabled, these mechanisms require additional bandwidth on the local link on which the effect of a prune is delayed (in our case, the PE-CE link);
- o when enabled, these mechanisms require disabling explicit tracking (see Section 4.3.3 of [RFC7761]), even though enabling this feature may otherwise be desired;
- o on certain implementations, these mechanisms are incompatible with behaviors that cannot be turned off (e.g., implementation applying a fast-leave behavior on interfaces with only two neighbors);
- o they do not provide a suitable level of configurability; and
- o they do not provide a way to discriminate between multicast flows based on estimation of their dynamicity.

#### 4.3. BGP Route Damping

The procedures defined in [RFC2439] and [RFC7196] for BGP route flap damping are useful for operators who want to control the impact of unicast route churn on the routing infrastructure and offer a standardized set of parameters to control damping.

These procedures are not directly relevant in a multicast context for the following reasons:

- o they are not specified for multicast routing protocol in general, and
- o even in contexts where BGP routes are used to carry multicast routing states (e.g., [RFC6514]), these procedures do not allow the implementation of the principle described in this document; the main reason being that a damped route becomes suppressed while the target behavior would be to keep advertising when damping is triggered on a multicast route.

However, the set of parameters standardized to control the thresholds of the exponential decay mechanism can be relevantly reused. This is the approach proposed for the procedures described in this document (Section 5). Motivations for doing so are to help the network operator deploy this feature based on consistent configuration parameters and to obtain predictable results without the drawbacks of relying on rate-limiting multicast control protocol exchanges (as is exposed in Section 4.1) or on the use of existing PIM/IGMP timers (as is exposed in Section 4.2).

## 5. Procedures for Multicast State Damping

### 5.1. PIM Procedures

This section describes procedures for multicast state damping satisfying the goals spelled out in Section 1. This section describes procedures for (S,G) states in the PIM-SM protocol [RFC7761]; they apply unchanged for such states created based on multicast group management protocols (IGMP [RFC3376], MLD [RFC3810]) on downstream interfaces. The same procedures are applied to (\*,G) states in the context of PIM-SM Any-Source Multicast (ASM) groups (damping is not applied to (S,G,Rpt) Prune state).

The following notions of [RFC2439] are reused in these procedures:

**figure-of-merit:** A number reflecting the current estimation of recent past activity of an (S,G) multicast routing state, which increases based on routing events related to this state and decreases between these events following an exponential decay function (see below); the activation or inactivation of damping on the state is based on this number. This number is associated with the upstream state machine for (S,G) and is initialized to a value of zero on state creation.

exponential decay function: A mathematical function as defined in Section 2.3 of [RFC2439] (ignoring the first paragraph of the section, as it does not apply here).

decay-half-life: The duration used to control how fast the exponential decay of the *figure-of-merit* is; this parameter of the exponential decay function is the time duration during which the *figure-of-merit* will be reduced by half when in the absence of a routing event (configurable parameter).

cutoff-threshold: The value of the *figure-of-merit* over which damping is applied (configurable parameter).

reuse-threshold: The value of the *figure-of-merit* under which damping stops being applied (configurable parameter).

In addition to these values, a configurable *increment-factor* parameter is introduced that controls by how much the *figure-of-merit* is incremented on multicast state update events.

Section 7.3 proposes default and maximum values for the configurable parameters.

On reception of updated multicast membership or routing information on a downstream interface I for a given (S,G) state, which results in a change of the state of the PIM downstream state machine (see Section 4.5.3 of [RFC7761]), a router implementing these procedures MUST:

- o apply procedures of [RFC7761] unchanged, for everything relating to what multicast traffic ends up being sent on downstream interfaces, including interface I
- o update the *figure-of-merit* following the exponential decay algorithm
- o increase the *figure-of-merit* for the (S,G) by the *increment-factor*\*
- o update the damping state for the (S,G) state: damping becomes active on the state if the recomputed *figure-of-merit* is strictly above the configured *cutoff-threshold*:
  - \* if damping remains inactive on (S,G) state, update the upstream state machine as usual (as per Section 4.5.7 of [RFC7761]).



- \* if damping becomes active for the (S,G) state:
  - + if the received message has caused the upstream state machine to transition to Joined state, update the upstream state machine for (S,G) applying usual PIM procedures in Section 4.5.7 of [RFC7761] and including sending a PIM Join to the upstream neighbor
  - + if the received message has caused the upstream state machine to transition to NotJoined state, do not update the upstream state machine for (S,G)
  - + hold the upstream state machine in Joined state until the reuse threshold is reached: for the purpose of updating this state machine, events that may result in updating the state based on [RFC7761] SHOULD be ignored until the *\*reuse-threshold\** is reached. The effect is that in the meantime, while PIM Join messages may be sent as refreshes to the upstream neighbor, no PIM Prune message will be sent.
- \* if damping was already active, do not update the upstream state machine for (S,G); the upstream state machine was frozen after processing the previous message.

Once the *\*figure-of-merit\** for (S,G) damping state decays to a value strictly below the configured *\*reuse-threshold\**, the upstream state machine for (S,G) is recomputed based on states of downstream state machines, eventually leading to a PIM Join or Prune message to be sent to the upstream neighbor.

Given the specificity of multicast applications, it is REQUIRED for the implementation to let the operator configure the *\*decay-half-life\** in seconds, rather than in minutes.

This specification does not impose the use of a particular technique to update the *\*figure-of-merit\** following the exponential decay controlled by the configured *\*decay-half-life\**. For instance, the same techniques as the ones described in [RFC2439] can be applied. The only requirement is that the *\*figure-of-merit\** has to be updated prior to increasing it and that its decay below the *\*reuse-threshold\** has to be reacted upon in a timely manner: in particular, if the recomputation is done with a fixed time granularity, this granularity should be low enough to not significantly delay the inactivation of damping on a multicast state beyond what the operator wanted to configure (e.g., for a *\*decay-half-life\** of 10s, recomputing the *\*figure-of-merit\** each minute would result in a multicast state remaining damped for a much longer time than specified).

PIM implementations typically follow the suggestion from Section 4.1 of [RFC7761] that:

implementations will only maintain state when it is relevant to forwarding operations - for example, the 'NoInfo' state might be assumed from the lack of other state information, rather than being held explicitly.

To properly implement damping procedures, an implementation MUST keep an explicit (S,G) state as long as damping is active on an (S,G). Once an (S,G) state expires, and damping becomes inactive on this state, its associated \*figure-of-merit\* and damping state are removed as well.

Note that these procedures:

- o do not impact PIM procedures related to refreshes or expiration of multicast routing states: PIM Prune messages triggered by the expiration of the (S,G) keep-alive timer are not suppressed or delayed, and the reception of Join messages not causing transition of state on the downstream interface does not lead to incrementing the \*figure-of-merit\*;
- o do not impact the PIM Assert mechanism: in particular, PIM Prune messages triggered by a change of the PIM Assert winner on the upstream interface are not suppressed or delayed;
- o do not impact PIM Prune messages that are sent when the RPF neighbor is updated for a given multicast flow; and
- o do not impact PIM Prune messages that are sent in the context of switching between a Rendezvous Point Tree and a Shortest Path Tree.

Note also that no action is triggered based on the reception of PIM Prune messages (or corresponding IGMP/MLD messages) that relate to non-existing (S,G) state: in particular, no \*figure-of-merit\* or damping state is created in this case.

## 5.2. Procedures for Multicast VPN State Damping

The procedures described in Section 5.1 can be applied in the Virtual Routing and Forwarding (VRF) PIM-SM implementation (in the "C-PIM instance"), with the corresponding action to suppressing the emission of a Prune(S,G) message being to not withdraw the C-multicast Source Tree Join (C-S,C-G) BGP route. An implementation of [RFC6513] relying on the use of PIM to carry C-multicast routing information MUST support this technique to be compliant with this specification.

In the context of [RFC6514], where BGP is used to distribute C-multicast routing information, the following procedure is proposed as an alternative to the procedures in Section 5.1 and consists in applying damping in the BGP implementation based on existing BGP damping mechanisms applied to C-multicast Source Tree Join routes and Shared Tree Join routes (and as well to Leaf A-D routes - see Section 6) and modified to implement the behavior described in Section 3 along the following guidelines:

- o not withdrawing (instead of not advertising) damped routes;
- o providing means to configure the \*decay-half-life\* in seconds if that option is not already available; and
- o using parameters for the exponential decay that are specific to multicast based on default values and multicast-specific configuration.

While these procedures would typically be implemented on PE routers, in a context where BGP Route Reflectors (RRs) [RFC4456] are used it can be considered useful to also be able to apply damping on RRs as well to provide additional protection against activity created behind multiple PEs. Additionally, for MVPN Inter-AS deployments, it can be needed to protect one Autonomous System (AS) from the dynamicity of multicast VPN routing events from other ASes.

The choice to implement damping based on BGP routes or the procedures described in Section 5.1 is up to the implementor, but at least one of the two MUST be implemented. In the perspective of allowing damping to be done on RRs and Autonomous System Border Routers (ASBRs), implementing the BGP approach is recommended.

When not all routers in a deployment have the capability to drop traffic coming from the wrong PE (as spelled out in Section 9.1.1 of [RFC6513]), then the withdrawal of a C-multicast route resulting from a change in the Upstream Multicast Hop or Upstream Multicast PE SHOULD NOT be damped. An implementation of this specification MUST do at least one of the two following things:

- o not damp these withdrawals by default, and/or
- o provide a tuning knob to disable the damping of these withdrawals.

Additionally, in such a deployment context, it is RECOMMENDED not to enable any multicast VPN route damping on RRs and ASBRs since these types of equipment cannot distinguish the event having caused a C-multicast to be withdrawn.

Note well that it is out of scope of this section to consider the application of these damping techniques on MVPN BGP routes other than C-multicast routes.

## 6. Procedures for P-Tunnel State Damping

### 6.1. Damping MVPN P-Tunnel Change Events

When selective P-tunnels are used (see Section 7 of [RFC6513]), the effect of updating the upstream state machine for a given (C-S,C-G) state on a PE connected to multicast receivers is not only to generate activity to propagate C-multicast routing information to the source connected PE, but also to possibly trigger changes related to the P-tunnels carrying (C-S,C-G) traffic. Protecting the provider network from an excessive amount of change in the state of P-tunnels is required, and this section details how this can be done.

A PE implementing these procedures for MVPN MUST damp Leaf A-D routes in the same manner as it would for C-multicast routes (see Section 5.2).

A PE implementing these procedures for MVPN MUST damp the activity related to removing itself from a P-tunnel. Possible ways to do so depend on the type of P-tunnel, and local implementation details are left up to the implementor.

The following is proposed as an example of how the above can be achieved:

- o For P-tunnels implemented with the PIM protocol, this consists in applying multicast state damping techniques described in Section 5.1 to the P-PIM instance, at least for (S,G) states corresponding to P-tunnels.
- o For P-tunnels implemented with multipoint LDP (mLDP), this consists in applying damping techniques completely similar to the one described in Section 5 but generalized to apply to mLDP states.
- o For root-initiated P-tunnels (P-tunnels implemented with the Point-to-Multipoint (P2MP) RSVP-TE, or relying on ingress replication), no particular action needs to be implemented to damp P-tunnels membership, if the activity of Leaf A-D route themselves is damped.

- o Another possibility is to base the decision to join or not join the P-tunnel to which a given (C-S,C-G) is bound and to advertise or not advertise a Leaf A-D route related to (C-S,C-G) based on whether or not a C-multicast Source Tree Join route is being advertised for (C-S,C-G) rather than by relying on the state of the C-PIM Upstream state machine for (C-S,C-G).

## 6.2. Procedures for Ethernet VPNs

Specifications exist to support or optimize multicast and broadcast in the context of Ethernet VPNs [RFC7117] relying on the use of Selective P-Multicast Service Interface (S-PMSI) and P-tunnels. For the same reasons as for IP multicast VPNs, an implementation of [RFC7117] MUST follow the procedures described in Section 6.1 to be compliant with this specification.

## 7. Operational Considerations

### 7.1. Enabling Multicast Damping

In the context of multicast VPNs, these procedures would be enabled on PE routers. Additionally, in the case of C-multicast routing based on BGP extensions ([RFC6514]), these procedures can be enabled on ASBRs and RRs.

### 7.2. Troubleshooting and Monitoring

Implementing the damping mechanisms described in this document should be complemented by appropriate tools to observe and troubleshoot damping activity.

Complementing the existing interface providing information on multicast states with information on eventual damping of corresponding states (e.g., Multicast Routing Information Base (MRIB) states) is RECOMMENDED for C-multicast routing states and P-tunnel states.

### 7.3. Default and Maximum Values

Considering that, by design, multicast streams will be delivered unchanged to the end user independent of the value chosen for the configurable parameters, and that the only trade-off being made is an increase of bandwidth use, the default and maximum values do not have to be perfectly tuned.

This section proposes default and maximum values that are conservative, so as to not significantly impact network dimensioning but still prevent multicast state churn going beyond what can be

considered a reasonably low churn for a multicast state (see below for illustrations in order of magnitude of the effect of these values).

The following values are RECOMMENDED to be adopted as default values:

- o `*increment-factor*`: 1000
- o `*cutoff-threshold*`: 3000
- o `*decay-half-life*`: 10s
- o `*reuse-threshold*`: 1500

For unicast damping, it is common to set an upper bound to the time during which a route is suppressed. In the case of multicast state damping, which relies on not withdrawing a damped route, it may be desirable to avoid a situation where a multicast flow would keep flowing in a portion of the network for a very long time in the absence of receivers.

The proposed default maximum value for the `*figure-of-merit*` is  $20 \times \text{*increment-factor*}$ , i.e., 20000 with the proposed default `*increment-factor*` of 1000.

As illustrations, with these values:

- o a multicast state updated less frequently than once every 6 s will not be damped at all;
- o a multicast state changing once per second for 3 s, and then not changing, will not be damped;
- o a multicast state changing once per second for 4 s, and then not changing, will be damped after the fourth change for approximately 13 s;
- o a multicast state changing twice per second for 15 s, and then not changing, will be damped after the fourth change for approximately 50 s; and
- o a multicast state changing at a fast pace for a long time will reach the maximum of `*figure-of-merit*`; once the activity on this state stops, corresponding traffic may still flow in the network for approximately 37 s before dampening stops being active.

The following values are proposed as maximums:

- o \*decay-half-life\*: 60 s
- o \*cutoff-threshold\*: 50000

More aggressive protection against the risk of denial of service can be achieved by increasing the \*increment-factor\* or the \*decay-half-life\*, or by reducing the \*cutoff-threshold\* and/or \*reuse-threshold\*.

## 8. Security Considerations

The procedures defined in this document do not introduce additional security issues not already present in the contexts addressed and actually aim at addressing some of the identified risks without introducing as much denial-of-service risk as some of the mechanisms already defined.

The protection provided relates to the control plane of the multicast routing protocols, including the components implementing the routing protocols and the components responsible for updating the multicast forwarding plane.

The procedures described are meant to provide some level of protection for the router on which they are enabled by reducing the amount of routing state updates that it needs to send to its upstream neighbor or peers but do not provide any reduction of the control-plane load related to processing routing information from downstream neighbors. Protecting routers from an increase in control-plane load due to activity on downstream interfaces toward core routers (or in the context of BGP-based MVPN C-multicast routing, BGP peers) relies on the activation of damping on corresponding downstream neighbors (or BGP peers) and/or at the edge of the network. Protecting routers from an increase in control-plane load due to activity on customer-facing downstream interfaces or downstream interfaces to routers in another administrative domain is out of the scope of this document and should use already defined mechanisms (see [RFC4609]).

To be effective, the procedures described here must be complemented by configuration limiting the number of multicast states that can be created on a multicast router through protocol interactions with multicast receivers, neighbor routers in adjacent ASes, or in multicast VPN contexts with multicast CEs. Note well that the two mechanisms may interact: the state for which Prune has been requested may still remain taken into account for some time if damping has been triggered and hence result in an otherwise acceptable new state from being successfully created.

Additionally, it is worth noting that these procedures are not meant to protect against peaks of control-plane load but only address averaged load. For instance, assuming a set of multicast states are submitted to the same Join/Prune events, damping can prevent more than a certain number of Join/Prune messages to be sent upstream in the period of time that elapses between the reception of Join/Prune messages triggering the activation of damping on these states and when damping becomes inactive after decay.

## 9. References

### 9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2439] Villamizar, C., Chandra, R., and R. Govindan, "BGP Route Flap Damping", RFC 2439, DOI 10.17487/RFC2439, November 1998, <<http://www.rfc-editor.org/info/rfc2439>>.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<http://www.rfc-editor.org/info/rfc3376>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<http://www.rfc-editor.org/info/rfc3810>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<http://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<http://www.rfc-editor.org/info/rfc6514>>.
- [RFC7117] Aggarwal, R., Ed., Kamite, Y., Fang, L., Rekhter, Y., and C. Kodeboniya, "Multicast in Virtual Private LAN Service (VPLS)", RFC 7117, DOI 10.17487/RFC7117, February 2014, <<http://www.rfc-editor.org/info/rfc7117>>.



- [RFC7196] Pelsser, C., Bush, R., Patel, K., Mohapatra, P., and O. Maennel, "Making Route Flap Damping Usable", RFC 7196, DOI 10.17487/RFC7196, May 2014, <<http://www.rfc-editor.org/info/rfc7196>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<http://www.rfc-editor.org/info/rfc7761>>.

## 9.2. Informative References

- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<http://www.rfc-editor.org/info/rfc4456>>.
- [RFC4609] Savola, P., Lehtonen, R., and D. Meyer, "Protocol Independent Multicast - Sparse Mode (PIM-SM) Multicast Routing Security Issues and Enhancements", RFC 4609, DOI 10.17487/RFC4609, October 2006, <<http://www.rfc-editor.org/info/rfc4609>>.

## Acknowledgements

We would like to thank Bruno Decraene and Lenny Giuliano for discussions that helped shape this proposal. We would also like to thank Yakov Rekhter and Eric Rosen for their reviews and helpful comments. Thanks to Wim Henderickx for his comments and support of this proposal. Additionally, Martin Vigoureux, Gunter Van De Velde, and Alvaro Retana provided useful comments to finalize the document.

## Authors' Addresses

Thomas Morin (editor)  
Orange  
2, avenue Pierre Marzin  
Lannion 22307  
France

Email: [thomas.morin@orange.com](mailto:thomas.morin@orange.com)

Stephane Litkowski  
Orange  
France

Email: [stephane.litkowski@orange.com](mailto:stephane.litkowski@orange.com)

Keyur Patel  
Cisco Systems  
170 W. Tasman Drive  
San Jose, CA 95134  
United States of America

Email: [keyupate@cisco.com](mailto:keyupate@cisco.com)

Zhaohui Zhang  
Juniper Networks Inc.  
10 Technology Park Drive  
Westford, MA 01886  
United States of America

Email: [zzhang@juniper.net](mailto:zzhang@juniper.net)

Robert Kebler  
Juniper Networks Inc.  
10 Technology Park Drive  
Westford, MA 01886  
United States of America

Email: [rkebler@juniper.net](mailto:rkebler@juniper.net)

Jeffrey Haas  
Juniper Networks

Email: [jhaas@juniper.net](mailto:jhaas@juniper.net)