Internet Engineering Task Force (IETF)
Request for Comments: 8293
Category: Informational
ISSN: 2070-1721

A. Ghanwani
Dell
L. Dunbar
M. McBride
Huawei
V. Bannai
Google
R. Krishnan
Dell

January 2018

A Framework for Multicast in Network Virtualization over Layer 3

Abstract

This document provides a framework for supporting multicast traffic in a network that uses Network Virtualization over Layer 3 (NVO3). Both infrastructure multicast and application-specific multicast are discussed. It describes the various mechanisms that can be used for delivering such traffic as well as the data plane and control plane considerations for each of the mechanisms.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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 https://www.rfc-editor.org/info/rfc8293.

Table of Contents

1. Introduction

Network Virtualization over Layer 3 (NVO3) [RFC7365] is a technology that is used to address issues that arise in building large, multi-tenant data centers (DCs) that make extensive use of server virtualization [RFC7364].

This document provides a framework for supporting multicast traffic in a network that uses NVO3. Both infrastructure multicast and application-specific multicast are considered. It describes various mechanisms, and the considerations of each of them, that can be used for delivering such traffic in networks that use NVO3.

The reader is assumed to be familiar with the terminology and concepts as defined in the NVO3 Framework [RFC7365] and NVO3 Architecture [RFC8014] documents.

1.1. Infrastructure Multicast

Infrastructure multicast refers to networking services that require multicast or broadcast delivery, such as Address Resolution Protocol (ARP), Neighbor Discovery (ND), Dynamic Host Configuration Protocol (DHCP), multicast Domain Name Server (mDNS), etc., some of which are described in Sections 5 and 6 of RFC 3819 [RFC3819]. It is possible to provide solutions for these services that do not involve multicast in the underlay network. For example, in the case of ARP/ND, a Network Virtualization Authority (NVA) can be used for distributing the IP address to Media Access Control (MAC) address mappings to all of the Network Virtualization Edges (NVEs). An NVE can then trap ARP Request and/or ND Neighbor Solicitation messages from the Tenant Systems (TSs) that are attached to it and respond to them, thereby eliminating the need for the broadcast/multicast of such messages. In the case of DHCP, the NVE can be configured to forward these messages using the DHCP relay function [RFC2131].

Of course, it is possible to support all of these infrastructure multicast protocols natively if the underlay provides multicast transport. However, even in the presence of multicast transport, it may be beneficial to use the optimizations mentioned above to reduce the amount of such traffic in the network.

1.2. Application-Specific Multicast

Application-specific multicast traffic refers to multicast traffic that originates and is consumed by user applications. Several such applications are described elsewhere [DC-MC]. Application-specific multicast may be either Source-Specific Multicast (SSM) or Any-Source Multicast (ASM) [RFC3569] and has the following characteristics:

1. Receiver hosts are expected to subscribe to multicast content using protocols such as IGMP [RFC3376] (IPv4) or Multicast Listener Discovery (MLD) [RFC2710] (IPv6). Multicast sources and listeners participate in these protocols using addresses that are in the TS address domain.

2. The set of multicast listeners for each multicast group may not be known in advance. Therefore, it may not be possible or practical for an NVA to get the list of participants for each multicast group ahead of time.

2. Terminology and Abbreviations

In this document, the terms host, Tenant System (TS), and Virtual Machine (VM) are used interchangeably to represent an end station that originates or consumes data packets.

ASM: Any-Source Multicast

IGMP: Internet Group Management Protocol

LISP: Locator/ID Separation Protocol

MSN: Multicast Service Node

RLOC: Routing Locator

NVA: Network Virtualization Authority

NVE: Network Virtualization Edge

NVGRE: Network Virtualization using GRE

PIM: Protocol-Independent Multicast

SSM: Source-Specific Multicast

TS: Tenant System

VM: Virtual Machine

VN: Virtual Network

VTEP: VXLAN Tunnel End Point

VXLAN: Virtual eXtensible LAN

3. Multicast Mechanisms in Networks That Use NVO3

In NVO3 environments, traffic between NVEs is transported using an encapsulation such as VXLAN [RFC7348] [VXLAN-GPE], Network Virtualization using Generic Routing Encapsulation (NVGRE) [RFC7637], Geneve [Geneve], Generic UDP Encapsulation [GUE], etc.

What makes networks using NVO3 different from other networks is that some NVEs, especially NVEs implemented in servers, might not support regular multicast protocols such as PIM. Instead, the only capability they may support would be that of encapsulating data packets from VMs with an outer unicast header. Therefore, it is important for networks using NVO3 to have mechanisms to support multicast as a network capability for NVEs, to map multicast traffic from VMs (users/applications) to an equivalent multicast capability inside the NVE, or to figure out the outer destination address if NVE does not support native multicast (e.g., PIM) or IGMP.

With NVO3, there are many possible ways that multicast may be handled in such networks. We discuss some of the attributes of the following four methods:

1. No multicast support

2. Replication at the source NVE

3. Replication at a multicast service node

4. IP multicast in the underlay

These methods are briefly mentioned in the NVO3 Framework [RFC7365] and NVO3 Architecture [RFC8014] documents. This document provides more details about the basic mechanisms underlying each of these methods and discusses the issues and trade-offs of each.

We note that other methods are also possible, such as [EDGE-REP], but we focus on the above four because they are the most common.

It is worth noting that when selecting a multicast mechanism, it is useful to consider the impact of these on any multicast congestion control mechanisms that applications may be using to obtain the desired system dynamics. In addition, the same rules for Explicit

Congestion Notification (ECN) would apply to multicast traffic being encapsulated, as for unicast traffic [RFC6040].

3.1. No Multicast Support

In this scenario, there is no support whatsoever for multicast traffic when using the overlay. This method can only work if the following conditions are met:

1. All of the application traffic in the network is unicast traffic, and the only multicast/broadcast traffic is from ARP/ND protocols.

2. An NVA is used by all of the NVEs to determine the mapping of a given TS's MAC and IP address to the NVE that it is attached to. In other words, there is no data-plane learning. Address resolution requests via ARP/ND that are issued by the TSs must be resolved by the NVE that they are attached to.

With this approach, it is not possible to support application-specific multicast. However, certain multicast/broadcast applications can be supported without multicast; for example, DHCP, which can be supported by use of DHCP relay function [RFC2131].

The main drawback of this approach, even for unicast traffic, is that it is not possible to initiate communication with a TS for which a mapping to an NVE does not already exist at the NVA. This is a problem in the case where the NVE is implemented in a physical switch and the TS is a physical end station that has not registered with the NVA.

3.2. Replication at the Source NVE

With this method, the overlay attempts to provide a multicast service without requiring any specific support from the underlay, other than that of a unicast service. A multicast or broadcast transmission is achieved by replicating the packet at the source NVE and making copies, one for each destination NVE that the multicast packet must be sent to.

For this mechanism to work, the source NVE must know, a priori, the IP addresses of all destination NVEs that need to receive the packet. For the purpose of ARP/ND, this would involve knowing the IP addresses of all the NVEs that have TSs in the VN of the TS that generated the request.

For the support of application-specific multicast traffic, a method similar to that of receiver-sites registration for a particular multicast group, described in [LISP-Signal-Free], can be used. The registrations from different receiver sites can be merged at the NVA, which can construct a multicast replication list inclusive of all NVEs to which receivers for a particular multicast group are attached. The replication list for each specific multicast group is maintained by the NVA. Note that using receiver-sites registration does not necessarily mean the source NVE must do replication. If the NVA indicates that multicast packets are encapsulated to multicast service nodes, then there would be no replication at the NVE.

The receiver-sites registration is achieved by egress NVEs performing IGMP/MLD snooping to maintain the state for which attached TSs have subscribed to a given IP multicast group. When the members of a multicast group are outside the NVO3 domain, it is necessary for NVO3 gateways to keep track of the remote members of each multicast group. The NVEs and NVO3 gateways then communicate with the multicast groups that are of interest to the NVA. If the membership is not communicated to the NVA, and if it is necessary to prevent TSs attached to an NVE that have not subscribed to a multicast group from receiving the multicast traffic, the NVE would need to maintain multicast group membership information.

In the absence of IGMP/MLD snooping, the traffic would be delivered to all TSs that are part of the VN.

In multihoming environments, i.e., in those where a TS is attached to more than one NVE, the NVA would be expected to provide information to all of the NVEs under its control about all of the NVEs to which such a TS is attached. The ingress NVE can then choose any one of those NVEs as the egress NVE for the data frames destined towards the multi-homed TS.

This method requires multiple copies of the same packet to all NVEs that participate in the VN. If, for example, a tenant subnet is spread across 50 NVEs, the packet would have to be replicated 50 times at the source NVE. Obviously, this approach creates more traffic to the network that can cause congestion when the network load is high. This also creates an issue with the forwarding performance of the NVE.

Note that this method is similar to what was used in Virtual Private LAN Service (VPLS) [RFC4762] prior to support of Multiprotocol Label Switching (MPLS) multicast [RFC7117]. While there are some similarities between MPLS Virtual Private Network (VPN) and NVO3, there are some key differences:

- The attachment from Customer Edge (CE) to Provider Edge (PE) in VPNs is somewhat static, whereas in a DC that allows VMs to migrate anywhere, the TS attachment to NVE is much more dynamic.

- The number of PEs to which a single VPN customer is attached in an MPLS VPN environment is normally far less than the number of NVEs to which a VN's VMs are attached in a DC.

When a VPN customer has multiple multicast groups, "Multicast VPN" [RFC6513] combines all those multicast groups within each VPN client to one single multicast group in the MPLS (or VPN) core. The result is that messages from any of the multicast groups belonging to one VPN customer will reach all the PE nodes of the client. In other words, any messages belonging to any multicast groups under customer X will reach all PEs of the customer X. When the customer X is attached to only a handful of PEs, the use of this approach does not result in an excessive waste of bandwidth in the provider's network.

In a DC environment, a typical hypervisor-based virtual switch may only support on the order of 10's of VMs (as of this writing). A subnet with N VMs may be, in the worst case, spread across N virtual switches (vSwitches). Using an "MPLS VPN multicast" approach in such a scenario would require the creation of a multicast group in the core in order for the VN to reach all N NVEs. If only a small percentage of this client's VMs participate in application-specific multicast, a great number of NVEs will receive multicast traffic that is not forwarded to any of their attached VMs, resulting in a considerable waste of bandwidth.

Therefore, the multicast VPN solution may not scale in a DC environment with the dynamic attachment of VNs to NVEs and a greater number of NVEs for each VN.

3.3. Replication at a Multicast Service Node

With this method, all multicast packets would be sent using a unicast tunnel encapsulation from the ingress NVE to a Multicast Service Node (MSN). The MSN, in turn, would create multiple copies of the packet and would deliver a copy, using a unicast tunnel encapsulation, to each of the NVEs that are part of the multicast group for which the packet is intended.

This mechanism is similar to that used by the Asynchronous Transfer Mode (ATM) Forum's LAN Emulation (LANE) specification [LANE]. The MSN is similar to the Rendezvous Point (RP) in Protocol Independent Multicast - Sparse Mode (PIM-SM), but different in that the user data traffic is carried by the NVO3 tunnels.

The following are possible ways for the MSN to get the membership information for each multicast group:

- The MSN can obtain this membership information from the IGMP/MLD report messages sent by TSs in response to IGMP/MLD query messages from the MSN. The IGMP/MLD query messages are sent from the MSN to the NVEs, which then forward the query messages to TSs attached to them. An IGMP/MLD query message sent out by the MSN to an NVE is encapsulated with the MSN address in the outer IP source address field and the address of the NVE in the outer IP destination address field. An encapsulated IGMP/MLD query message also has a virtual network (VN) identifier (corresponding to the VN that the TSs belong to) in the outer header and a multicast address in the inner IP destination address field. Upon receiving the encapsulated IGMP/MLD query message, the NVE establishes a mapping for "MSN address" to "multicast address", decapsulates the received encapsulated IGMP/MLD message, and multicasts the decapsulated query message to the TSs that belong to the VN attached to that NVE. An IGMP/MLD report message sent by a TS includes the multicast address and the address of the TS. With the proper "MSN address" to "multicast address" mapping, the NVEs can encapsulate all multicast data frames containing the "multicast address" with the address of the MSN in the outer IP destination address field.

- The MSN can obtain the membership information from the NVEs that have the capability to establish multicast groups by snooping native IGMP/MLD messages (note that the communication must be specific to the multicast addresses) or by having the NVA obtain the information from the NVEs and in turn have MSN communicate with the NVA. This approach requires additional protocol between MSN and NVEs.

Unlike the method described in Section 3.2, there is no performance impact at the ingress NVE, nor are there any issues with multiple copies of the same packet from the source NVE to the MSN. However, there remain issues with multiple copies of the same packet on links that are common to the paths from the MSN to each of the egress NVEs. Additional issues that are introduced with this method include the availability of the MSN, methods to scale the services offered by the MSN, and the suboptimality of the delivery paths.

Finally, the IP address of the source NVE must be preserved in packet copies created at the multicast service node if data-plane learning is in use. This could create problems if IP source address Reverse Path Forwarding (RPF) checks are in use.

3.4. IP Multicast in the Underlay

In this method, the underlay supports IP multicast and the ingress NVE encapsulates the packet with the appropriate IP multicast address in the tunnel encapsulation header for delivery to the desired set of NVEs. The protocol in the underlay could be any variant of PIM, or a protocol-dependent multicast, such as [ISIS-Multicast].

If an NVE connects to its attached TSs via a Layer 2 network, there are multiple ways for NVEs to support the application-specific multicast:

- The NVE only supports the basic IGMP/MLD snooping function, while the "TS routers" handle the application-specific multicast. This scheme doesn't utilize the underlay IP multicast protocols. Instead routers, which are themselves TSs attached to the NVE, would handle multicast protocols for the application-specific multicast. We refer to such routers as TS routers.

- The NVE can act as a pseudo multicast router for the directly attached TSs and support the mapping of IGMP/MLD messages to the messages needed by the underlay IP multicast protocols.

With this method, there are none of the issues with the methods described in Sections 3.2 and 3.3 with respect to scaling and congestion. Instead, there are other issues described below.

With PIM-SM, the number of flows required would be (n*g), where n is the number of source NVEs that source packets for the group, and g is the number of groups. Bidirectional PIM (BIDIR-PIM) would offer better scalability with the number of flows required being g. Unfortunately, many vendors still do not fully support BIDIR or have limitations on its implementation. [RFC6831] describes the use of SSM as an alternative to BIDIR, provided that the NVEs have a way to learn of each other's IP addresses so that they can join all of the SSM Shortest Path Trees (SPTs) to create/maintain an underlay SSM IP multicast tunnel solution.

In the absence of any additional mechanism (e.g., using an NVA for address resolution), for optimal delivery, there would have to be a separate group for each VN for infrastructure multicast plus a separate group for each application-specific multicast address within a tenant.

An additional consideration is that only the lower 23 bits of the IP address (regardless of whether IPv4 or IPv6 is in use) are mapped to the outer MAC address, and if there is equipment that prunes multicasts at Layer 2, there will be some aliasing.

Finally, a mechanism to efficiently provision such addresses for each group would be required.

There are additional optimizations that are possible, but they come with their own restrictions. For example, a set of tenants may be restricted to some subset of NVEs, and they could all share the same outer IP multicast group address. This, however, introduces a problem of suboptimal delivery (even if a particular tenant within the group of tenants doesn't have a presence on one of the NVEs that another one does, the multicast packets would still be delivered to that NVE). It also introduces an additional network management burden to optimize which tenants should be part of the same tenant group (based on the NVEs they share), which somewhat dilutes the value proposition of NVO3 (to completely decouple the overlay and physical network design allowing complete freedom of placement of VMs anywhere within the DC).

Multicast schemes such as Bit Indexed Explicit Replication (BIER) [RFC8279] may be able to provide optimizations by allowing the underlay network to provide optimum multicast delivery without requiring routers in the core of the network to maintain per-multicast group state.

3.5. Other Schemes

There are still other mechanisms that may be used that attempt to combine some of the advantages of the above methods by offering multiple replication points, each with a limited degree of replication [EDGE-REP]. Such schemes offer a trade-off between the amount of replication at an intermediate node (e.g., router) versus performing all of the replication at the source NVE or all of the replication at a multicast service node.

4. Simultaneous Use of More Than One Mechanism

While the mechanisms discussed in the previous section have been discussed individually, it is possible for implementations to rely on more than one of these. For example, the method of Section 3.1 could be used for minimizing ARP/ND, while at the same time, multicast applications may be supported by one, or a combination, of the other methods. For small multicast groups, the methods of source NVE replication or the use of a multicast service node may be attractive, while for larger multicast groups, the use of multicast in the underlay may be preferable.

5. Other Issues

5.1. Multicast-Agnostic NVEs

Some hypervisor-based NVEs do not process or recognize IGMP/MLD frames, i.e., those NVEs simply encapsulate the IGMP/MLD messages in the same way as they do for regular data frames.

By default, a TS router periodically sends IGMP/MLD query messages to all the hosts in the subnet to trigger the hosts that are interested in the multicast stream to send back IGMP/MLD reports. In order for the MSN to get the updated multicast group information, the MSN can also send the IGMP/MLD query message comprising a client-specific multicast address encapsulated in an overlay header to all of the NVEs to which the TSs in the VN are attached.

However, the MSN may not always be aware of the client-specific multicast addresses. In order to perform multicast filtering, the MSN has to snoop the IGMP/MLD messages between TSs and their corresponding routers to maintain the multicast membership. In order for the MSN to snoop the IGMP/MLD messages between TSs and their router, the NVA needs to configure the NVE to send copies of the IGMP/MLD messages to the MSN in addition to the default behavior of sending them to the TSs' routers; e.g., the NVA has to inform the NVEs to encapsulate data frames with the Destination Address (DA) being 224.0.0.2 (DA of IGMP report) to the TSs' router and MSN.

This process is similar to "Source Replication" described in Section 3.2, except the NVEs only replicate the message to the TSs' router and MSN.

5.2. Multicast Membership Management for DC with VMs

For DCs with virtualized servers, VMs can be added, deleted, or moved very easily. When VMs are added, deleted, or moved, the NVEs to which the VMs are attached are changed.

When a VM is deleted from an NVE or a new VM is added to an NVE, the VM management system should notify the MSN to send the IGMP/MLD query messages to the relevant NVEs (as described in Section 3.3) so that the multicast membership can be updated promptly.

Otherwise, if there are changes of VMs attachment to NVEs (within the duration of the configured default time interval that the TSs routers use for IGMP/MLD queries), multicast data may not reach the VM(s) that moved.

6. Security Considerations

This document does not introduce any new security considerations beyond what is described in the NVO3 Architecture document [RFC8014].

7. IANA Considerations

This document does not require any IANA actions.

8. Summary

This document has identified various mechanisms for supporting application-specific multicast in networks that use NVO3. It highlights the basics of each mechanism and some of the issues with them. As solutions are developed, the protocols would need to consider the use of these mechanisms, and coexistence may be a consideration. It also highlights some of the requirements for supporting multicast applications in an NVO3 network.

9. References

9.1. Normative References

[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, <https://www.rfc-editor.org/info/rfc3376>.

[RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <https://www.rfc-editor.org/info/rfc6513>.

[RFC7364] Narten, T., Ed., Gray, E., Ed., Black, D., Fang, L., Kreeger, L., and M. Napierala, "Problem Statement: Overlays for Network Virtualization", RFC 7364, DOI 10.17487/RFC7364, October 2014, <https://www.rfc-editor.org/info/rfc7364>.

[RFC7365] Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for Data Center (DC) Network Virtualization", RFC 7365, DOI 10.17487/RFC7365, October 2014, <https://www.rfc-editor.org/info/rfc7365>.

[RFC8014] Black, D., Hudson, J., Kreeger, L., Lasserre, M., and T. Narten, "An Architecture for Data-Center Network Virtualization over Layer 3 (NVO3)", RFC 8014, DOI 10.17487/RFC8014, December 2016, <https://www.rfc-editor.org/info/rfc8014>.

9.2. Informative References

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/info/rfc2131>.

[RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, DOI 10.17487/RFC2710, October 1999, <https://www.rfc-editor.org/info/rfc2710>.

[RFC3569] Bhattacharyya, S., Ed., "An Overview of Source-Specific Multicast (SSM)", RFC 3569, DOI 10.17487/RFC3569, July 2003, <https://www.rfc-editor.org/info/rfc3569>.

[RFC3819] Karn, P., Ed., Bormann, C., Fairhurst, G., Grossman, D., Ludwig, R., Mahdavi, J., Montenegro, G., Touch, J., and L. Wood, "Advice for Internet Subnetwork Designers", BCP 89, RFC 3819, DOI 10.17487/RFC3819, July 2004, <https://www.rfc-editor.org/info/rfc3819>.

[RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, DOI 10.17487/RFC4762, January 2007, <https://www.rfc-editor.org/info/rfc4762>.

[RFC6040] Briscoe, B., "Tunnelling of Explicit Congestion Notification", RFC 6040, DOI 10.17487/RFC6040, November 2010, <https://www.rfc-editor.org/info/rfc6040>.

[RFC6831] Farinacci, D., Meyer, D., Zwiebel, J., and S. Venaas, "The Locator/ID Separation Protocol (LISP) for Multicast Environments", RFC 6831, DOI 10.17487/RFC6831, January 2013, <https://www.rfc-editor.org/info/rfc6831>.

[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, <https://www.rfc-editor.org/info/rfc7117>.

[RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <https://www.rfc-editor.org/info/rfc7348>.

[RFC7637] Garg, P., Ed. and Y. Wang, Ed., "NVGRE: Network Virtualization Using Generic Routing Encapsulation", RFC 7637, DOI 10.17487/RFC7637, September 2015, <https://www.rfc-editor.org/info/rfc7637>.

[RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <https://www.rfc-editor.org/info/rfc8279>.

[DC-MC] McBride, M. and H. Liu, "Multicast in the Data Center Overview", Work in Progress, draft-mcbride-armd-mcast-overview-02, July 2012.

[EDGE-REP] Marques, P., Fang, L., Winkworth, D., Cai, Y., and P. Lapukhov, "Edge multicast replication for BGP IP VPNs.", Work in Progress, draft-marques-l3vpn-mcast-edge-01, June 2012.

[Geneve] Gross, J., Ganga, I., and T. Sridhar, "Geneve: Generic Network Virtualization Encapsulation", Work in Progress, draft-ietf-nvo3-geneve-05, September 2017.

[GUE] Herbert, T., Yong, L., and O. Zia, "Generic UDP Encapsulation", Work in Progress, draft-ietf-intarea-gue-05, December 2017.

[ISIS-Multicast]
Yong, L., Weiguo, H., Eastlake, D., Qu, A., Hudson, J., and U. Chunduri, "IS-IS Protocol Extension For Building Distribution Trees", Work in Progress, draft-yong-isis-ext-4-distribution-tree-03, October 2014.

[LANE] ATM Forum, "LAN Emulation Over ATM: Version 1.0", ATM Forum Technical Committee, af-lane-0021.000, January 1995.

[LISP-Signal-Free]
Moreno, V. and D. Farinacci, "Signal-Free LISP Multicast", Work in Progress, draft-ietf-lisp-signal-free-multicast-07, November 2017.

[VXLAN-GPE]
Maino, F., Kreeger, L., and U. Elzur, "Generic Protocol Extension for VXLAN", Work in Progress, draft-ietf-nvo3-vxlan-gpe-05, October 2017.

Acknowledgments

Many thanks are due to Dino Farinacci, Erik Nordmark, Lucy Yong, Nicolas Bouliane, Saumya Dikshit, Joe Touch, Olufemi Komolafe, and Matthew Bocci for their valuable comments and suggestions.

Authors' Addresses

Anoop Ghanwani
Dell

Email: anoop@alumni.duke.edu

Linda Dunbar
Huawei Technologies
5340 Legacy Drive, Suite 1750
Plano, TX 75024
United States of America

Phone: (469) 277 5840
Email: ldunbar@huawei.com

Mike McBride
Huawei Technologies

Email: mmcbride7@gmail.com

Vinay Bannai
Google

Email: vbannai@gmail.com

Ram Krishnan
Dell

Email: ramkri123@gmail.com