

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 21, 2011

F. Brockners
Cisco
Y. Lee
Comcast
October 18, 2010

Multicast Considerations for Gateway-Initiated Dual-Stack Lite
draft-brockners-softwire-mcast-gi-ds-lite-00

Abstract

This document discusses multicast deployment aspects for networks which leverage Gateway-Initiated Dual-Stack lite.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 21, 2011.

Copyright Notice

Copyright (c) 2010 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 and Overview 3
- 2. Abbreviations 4
- 3. Multicast Deployment Considerations 4
 - 3.1. Architectural Attributes 4
 - 3.2. Overlapping private IPv4 addresses 5
 - 3.3. Considerations for the Gateway and AFTR 6
- 4. Acknowledgements 6
- 5. IANA Considerations 6
- 6. Security Considerations 7
- 7. References 7
 - 7.1. Normative References 7
 - 7.2. Informative References 7
- Authors' Addresses 8

Multicast traffic. Currently, most IPv4-Multicast deployments require the access devices to receive multicast traffic but not to source multicast traffic. This memo considers the scenario where the access device subscribes an IPv4-Multicast group and recommends how the multicast routing could be done. The following cases are out of scope in this memo:

- o IPv4-Multicast sourced by Access Devices.
- o Network Address Translation (NAT) for IPv4-Multicast traffic.

2. Abbreviations

The following abbreviations are used within this document:

AFTR: Address Family Transition Router (also known as "Large Scale NAT (LSN)" or "Dual-Stack lite Tunnel Concentrator", or "Carrier Grade NAT (CGN)"). An AFTR combines IP-in-IPv6 tunnel termination and IPv4-IPv4 NAT.

DS-lite: Dual-stack lite

GI-DS-lite: Gateway-initiated DS-lite

NAT: Network Address Translation

3. Multicast Deployment Considerations

This section details the IPv4-Multicast deployment considerations for GI-DS-lite. Several networks which follow the architecture shown in Figure 1 above deploy IPv4-Multicast. If GI-DS-lite is introduced, the Gateway continues to perform the role of the first hop IPv4-Multicast router and the overall multicast distribution architecture is left unchanged. Deployment dependent, the introduction of GI-DS-lite could go hand in hand with the Gateway no longer having native IPv4-Multicast connectivity. If the Gateway does not have native IPv4-Multicast connectivity it should create a tunnel (e.g. IP-in-IPv6 or IP-over-GRE6) to an IPv4-Multicast router (e.g. the closest). The Gateway peers with that IPv4-Multicast router via the tunnel to join the IPv4-Multicast routing domain.

3.1. Architectural Attributes

Deployment details for IP-Multicast are defined for several architectures which leverage tunnel-based access, such as [TR101] for DSL-Broadband, 3GPP TS 23.246 for mobile Multimedia Broadcast Service

(MBMS) [TS23246], or [I-D.ietf-multimob-pmipv6-base-solution] for multicast in Proxy Mobile-IP deployments). Multicast in mobile or broadband deployments with tunnel based access architectures share a set of common architectural attributes:

- o Subscribers are able to receive IP multicast, but are not assumed to send IP multicast (inline with the scope of this document).
- o The Gateway is an IP-Multicast router, which is attached to the IP-Multicast distribution network of the service provider.
- o Architectures often include devices which perform IGMP/MLD snooping and proxy reporting between the access device and the Gateway. Proxy Mobile IPv6 deployments [I-D.ietf-multimob-pmipv6-base-solution] are an example: Mobile devices (i.e. the Access Devices) are connected via a Mobile Access Gateway (MAG) implementing an IGMP proxy function to the Local Mobility Anchor (LMA) which performs the role of the Gateway.
- o In several broadband multicast deployments IP-Multicast traffic is not forwarded over the access tunnels used for unicast traffic, but uses an alternate vehicle, which allows for traffic replication between access devices and Gateway. DSL-broadband networks with Ethernet aggregation are an example: While unicast traffic is forwarded between the access devices and the Broadband Network Gateway (BNG) over dedicated point to point VLANs, a separate VLAN is used to forward multicast traffic. This allows taking advantage of the multicast replication capabilities of Ethernet within the aggregation network.

3.2. Overlapping private IPv4 addresses

GI-DS-lite supports deployments with (potentially overlapping) IPv4 addresses assigned to the access devices. This could present challenges from a theoretical point of view for the following scenarios:

1. The network deploys Source Specific Multicast (SSM) and IP-multicast is sourced from an access device: Per the note above, this scenario is out of the scope of this document.
2. The network deploys IGMPv3 and leverages explicit tracking (see appendix 2 of [RFC3376], or appendix 2 of [RFC3810]) only based on the source IP address of the IGMP messages: Explicit tracking is in use by several networks today, though one often does not rely (only) on the source IP-address to identify different hosts. Several multicast networks deploy devices performing IGMP

snooping with proxy reporting between the multicast host and the first hop IP-Multicast router. In those deployments, the source IP address of the IGMP join messages does no longer represent the multicast host. The Broadband Forum, for example, requires the source IP address of IGMP packets sent by the proxy reporting function to be 0.0.0.0 [TR101]. If explicit tracking is still desired in those environments, identifiers other than the source IP address need to be considered. Depending on the deployment and architecture, those could for example be the interface (as recommended for proxy Mobile-IP multicast deployments [I-D.ietf-multimob-pmipv6-base-solution]) or the MAC-address (see [TR101]), an access tunnel identifier etc.).

3.3. Considerations for the Gateway and AFTR

The Gateway's role with regards to IPv4-Multicast traffic forwarding and routing does not change if GI-DS-lite is deployed within the network and multicast traffic bypasses the AFTR. For deployments which require explicit tracking and the use of overlapping IPv4 address ranges at a Gateway, this Gateway needs to support explicit tracking based on identifiers other than the source IP-address of IGMP messages. The Gateway functions as IPv4-Multicast first hop router for the access devices. The Gateway is a multicast replication point for multicast flows towards receivers on or attached to the access devices. IPv4-Multicast traffic will be forwarded according to the group/source-specific forwarding states. If there are multiple receivers within the scope of the Gateway, its still a single flow which the Gateway receives. IPv4-Multicast forwarding at the Gateway is also not impacted in case IGMP-proxies exist between the access devices and the Gateway. This can be the case in broadband architectures (see [TR101]) as well as in mobile architectures (e.g., with PMIP, the MAG acts as MLD proxy, see [I-D.ietf-multimob-pmipv6-base-solution]).

4. Acknowledgements

The authors would like to acknowledge their discussions on this topic with Wojciech Dec and Sri Gundavelli.

5. IANA Considerations

This document includes no request to IANA.

6. Security Considerations

This draft does not introduce additional messages or novel protocol operations. Consequently, no new threats are introduced by this document in addition to those identified as security concerns for IP-Multicast deployments.

7. References

7.1. Normative References

- [I-D.ietf-multimob-pmipv6-base-solution]
Schmidt, T., Waehlich, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in PMIPv6 Domains", draft-ietf-multimob-pmipv6-base-solution-05 (work in progress), July 2010.
- [I-D.ietf-softwire-dual-stack-lite]
Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", draft-ietf-softwire-dual-stack-lite-06 (work in progress), August 2010.
- [I-D.ietf-softwire-gateway-init-ds-lite]
Brockners, F., Gundavelli, S., Speicher, S., and D. Ward, "Gateway Initiated Dual-Stack Lite Deployment", draft-ietf-softwire-gateway-init-ds-lite-00 (work in progress), May 2010.
- [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.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, May 2006.

7.2. Informative References

- [TR101] Broadband Forum, "TR-101: Migration to Ethernet-Based DSL Aggregation", April 2006.
- [TS23246] 3GPP, "3GPP TS 23.246: Multimedia Broadcast/Multicast

Service (MBMS), Architecture and functional description,
Release 9", June 2010.

Authors' Addresses

Frank Brockners
Cisco
Hansaallee 249, 3rd Floor
DUESSELDORF, NORDRHEIN-WESTFALEN 40549
Germany

Email: fbrockne@cisco.com

Yiu L. Lee
Comcast
One Comcast Center
Philadelphia, PA 19103
U.S.A.

Email: yiulee@cable.comcast.com
URI: <http://www.comcast.com>