

MBONED Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: April 30, 2018

P. Tarapore, Ed.
R. Sayko
AT&T
G. Shepherd
Cisco
T. Eckert, Ed.
Futurewei Technologies
R. Krishnan
SupportVectors
October 27, 2017

Use of Multicast Across Inter-Domain Peering Points
draft-ietf-mboned-interdomain-peering-bcp-12

Abstract

This document examines the use of Source Specific Multicast (SSM) across inter-domain peering points for a specified set of deployment scenarios. The objective is to describe the setup process for multicast-based delivery across administrative domains for these scenarios and document supporting functionality to enable this process.

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 <https://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 30, 2018.

Copyright Notice

Copyright (c) 2017 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

(<https://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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	3
2.	Overview of Inter-domain Multicast Application Transport . .	5
3.	Inter-domain Peering Point Requirements for Multicast	6
3.1.	Native Multicast	6
3.2.	Peering Point Enabled with GRE Tunnel	7
3.3.	Peering Point Enabled with an AMT - Both Domains Multicast Enabled	9
3.4.	Peering Point Enabled with an AMT - AD-2 Not Multicast Enabled	10
3.5.	AD-2 Not Multicast Enabled - Multiple AMT Tunnels Through AD-2	11
4.	Functional Guidelines	13
4.1.	Network Interconnection Transport and Security Guidelines	13
4.2.	Routing Aspects and Related Guidelines	14
4.2.1.	Native Multicast Routing Aspects	15
4.2.2.	GRE Tunnel over Interconnecting Peering Point	15
4.2.3.	Routing Aspects with AMT Tunnels	16
4.3.	Back Office Functions - Provisioning and Logging Guidelines	18
4.3.1.	Provisioning Guidelines	18
4.3.2.	Application Accounting Guidelines	20
4.3.3.	Log Management Guidelines	20
4.4.	Operations - Service Performance and Monitoring Guidelines	21
4.5.	Client Reliability Models/Service Assurance Guidelines .	23
5.	Troubleshooting and Diagnostics	23

6.	Security Considerations	24
7.	IANA Considerations	25
8.	Conclusions	25
9.	Acknowledgments	25
10.	Change log [RFC Editor: Please remove]	26
11.	References	26
11.1.	Normative References	26
11.2.	Informative References	27
	Authors' Addresses	28

[1.](#) Introduction

Content and data from several types of applications (e.g., live video streaming, software downloads) are well suited for delivery via multicast means. The use of multicast for delivering such content/data offers significant savings of utilization of resources in any given administrative domain. End user demand for such content/data is growing. Often, this requires transporting the content/data across administrative domains via inter-domain peering points.

The objective of this Best Current Practices document is twofold:

- o Describe the technical process and establish guidelines for setting up multicast-based delivery of application content/data across inter-domain peering points via a set of use cases.
- o Catalog all required information exchange between the administrative domains to support multicast-based delivery. This enables operators to initiate necessary processes to support inter-domain peering with multicast.

The scope and assumptions for this document are stated as follows:

- o For the purpose of this document, the term "peering point" refers to an interface between two networks/administrative domains over which traffic is exchanged between them. A Network-Network Interface (NNI) is an example of a peering point.
- o Administrative Domain 1 (AD-1) is enabled with native multicast. A peering point exists between AD-1 and AD-2.
- o It is understood that several protocols are available for this purpose including PIM-SM and Protocol Independent Multicast - Source Specific Multicast (PIM-SSM) [[RFC7761](#)], Internet Group Management Protocol (IGMP) [[RFC3376](#)], and Multicast Listener Discovery (MLD) [[RFC3810](#)].

- o As described in [Section 2](#), the source IP address of the multicast stream in the originating AD (AD-1) is known. Under this condition, PIM-SSM use is beneficial as it allows the receiver's upstream router to directly send a JOIN message to the source without the need of invoking an intermediate Rendezvous Point (RP). Use of SSM also presents an improved threat mitigation profile against attack, as described in [\[RFC4609\]](#). Hence, in the case of inter-domain peering, it is recommended to use only SSM protocols; the setup of inter-domain peering for ASM (Any-Source Multicast) is not in scope for this document.
- o AD-1 and AD-2 are assumed to adopt compatible protocols. The use of different protocols is beyond the scope of this document.
- o An Automatic Multicast Tunnel (AMT) [\[RFC7450\]](#) is setup at the peering point if either the peering point or AD-2 is not multicast enabled. It is assumed that an AMT Relay will be available to a client for multicast delivery. The selection of an optimal AMT relay by a client is out of scope for this document. Note that AMT use is necessary only when native multicast is unavailable in the peering point (Use Case 3.3) or in the downstream administrative domain (Use Cases 3.4, and 3.5).
- o The collection of billing data is assumed to be done at the application level and is not considered to be a networking issue. The settlements process for end user billing and/or inter-provider billing is out of scope for this document.
- o Inter-domain network connectivity troubleshooting is only considered within the context of a cooperative process between the two domains.

Thus, the primary purpose of this document is to describe a scenario where two AD's interconnect via a peering point with each other. Security and operational aspects for exchanging traffic on a public Internet Exchange Point (IXP) with a large shared broadcast domain between many operators, is not in scope for this document.

It may be possible to have a configuration whereby a transit domain (AD-3) interconnects AD-1 and AD-2. Such a configuration adds complexity and may require manual provisioning if, for example, AD-3 is not multicast enabled. This configuration is out of scope for this document; it is for further study.

This document also attempts to identify ways by which the peering process can be improved. Development of new methods for improvement is beyond the scope of this document.

2. Overview of Inter-domain Multicast Application Transport

A multicast-based application delivery scenario is as follows:

- o Two independent administrative domains are interconnected via a peering point.
- o The peering point is either multicast enabled (end-to-end native multicast across the two domains) or it is connected by one of two possible tunnel types:
 - o A Generic Routing Encapsulation (GRE) Tunnel [[RFC2784](#)] allowing multicast tunneling across the peering point, or
 - o An Automatic Multicast Tunnel (AMT) [[RFC7450](#)].
- o A service provider controls one or more application sources in AD-1 which will send multicast IP packets for one or more (S,G)s. It is assumed that the service being provided is suitable for delivery via multicast (e.g. live video streaming of popular events, software downloads to many devices, etc.), and that the packet streams will be part of a suitable multicast transport protocol.
- o An End User (EU) controls a device connected to AD-2, which runs an application client compatible with the service provider's application source.
- o The application client joins appropriate (S,G)s in order to receive the data necessary to provide the service to the EU. The mechanisms by which the application client learns the appropriate (S,G)s are an implementation detail of the application, and are out of scope for this document.

The assumption here is that AD-1 has ultimate responsibility for delivering the multicast based service on behalf of the content source(s). All relevant interactions between the two domains described in this document are based on this assumption.

Note that domain 2 may be an independent network domain (e.g., Tier 1 network operator domain). Alternately, domain 2 could also be an Enterprise network domain operated by a single customer. The peering point architecture and requirements may have some unique aspects associated with the Enterprise case.

The Use Cases describing various architectural configurations for the multicast distribution along with associated requirements is described in [section 3](#). Unique aspects related to the Enterprise


network possibility will be described in this section. [Section 4](#) contains a comprehensive list of pertinent information that needs to be exchanged between the two domains in order to support functions to enable the application transport.

3. Inter-domain Peering Point Requirements for Multicast

The transport of applications using multicast requires that the inter-domain peering point is enabled to support such a process. There are five Use Cases for consideration in this document.

3.1. Native Multicast

This Use Case involves end-to-end Native Multicast between the two administrative domains and the peering point is also native multicast enabled - Figure 1.



AD = Administrative Domain (Independent Autonomous System)
AS = Application (e.g., Content) Multicast Source
BR = Border Router
I1 = AD-1 and AD-2 Multicast Interconnection (e.g., MBGP)
I2 = AD-2 and EU Multicast Connection

Figure 1: - Content Distribution via End to End Native Multicast

Advantages of this configuration are:

- o Most efficient use of bandwidth in both domains.
- o Fewer devices in the path traversed by the multicast stream when compared to an AMT enabled peering point.

From the perspective of AD-1, the one disadvantage associated with native multicast into AD-2 instead of individual unicast to every EU in AD-2 is that it does not have the ability to count the number of

End Users as well as the transmitted bytes delivered to them. This information is relevant from the perspective of customer billing and operational logs. It is assumed that such data will be collected by the application layer. The application layer mechanisms for generating this information need to be robust enough such that all pertinent requirements for the source provider and the AD operator are satisfactorily met. The specifics of these methods are beyond the scope of this document.

Architectural guidelines for this configuration are as follows:

- a. Dual homing for peering points between domains is recommended as a way to ensure reliability with full BGP table visibility.
- b. If the peering point between AD-1 and AD-2 is a controlled network environment, then bandwidth can be allocated accordingly by the two domains to permit the transit of non- rate adaptive multicast traffic. If this is not the case, then it is recommended that the multicast traffic should support rate-adaption.
- c. The sending and receiving of multicast traffic between two domains is typically determined by local policies associated with each domain. For example, if AD-1 is a service provider and AD-2 is an enterprise, then AD-1 may support local policies for traffic delivery to, but not traffic reception from, AD-2. Another example is the use of a policy by which AD-1 delivers specified content to AD-2 only if such delivery has been accepted by contract.
- d. Relevant information on multicast streams delivered to End Users in AD-2 is assumed to be collected by available capabilities in the application layer. The precise nature and formats of the collected information will be determined by directives from the source owner and the domain operators.
- e. The interconnection of AD-1 and AD-2 should, at a minimum, follow guidelines for traffic filtering between autonomous systems [[BCP38](#)]. Filtering guidelines specific to the multicast control-plane and data-plane are described in [section 6](#).

[3.2](#). Peering Point Enabled with GRE Tunnel

The peering point is not native multicast enabled in this Use Case. There is a Generic Routing Encapsulation Tunnel provisioned over the peering point. In this case, the interconnection I1 between AD-1 and AD-2 in Figure 1 is multicast enabled via a Generic Routing Encapsulation Tunnel (GRE) [[RFC2784](#)] and encapsulating the multicast

protocols across the interface. The routing configuration is basically unchanged: Instead of BGP (SAFI2) across the native IP multicast link between AD-1 and AD-2, BGP (SAFI2) is now run across the GRE tunnel.

Advantages of this configuration:

- o Highly efficient use of bandwidth in both domains, although not as efficient as the fully native multicast Use Case.
- o Fewer devices in the path traversed by the multicast stream when compared to an AMT enabled peering point.
- o Ability to support only partial IP multicast deployments in AD- 1 and/or AD-2 (the two Border Routers in Figure 1 do not need to be the two "unicast" domain border routers; instead they can be anywhere in AD-1 and AD-2).
- o GRE is an existing technology and is relatively simple to implement.

Disadvantages of this configuration:

- o Per Use Case 3.1, current router technology cannot count the number of end users or the number bytes transmitted.
- o GRE tunnel requires manual configuration.
- o The GRE must be established prior to stream starting.
- o The GRE tunnel is often left pinned up.

Architectural guidelines for this configuration include the following:

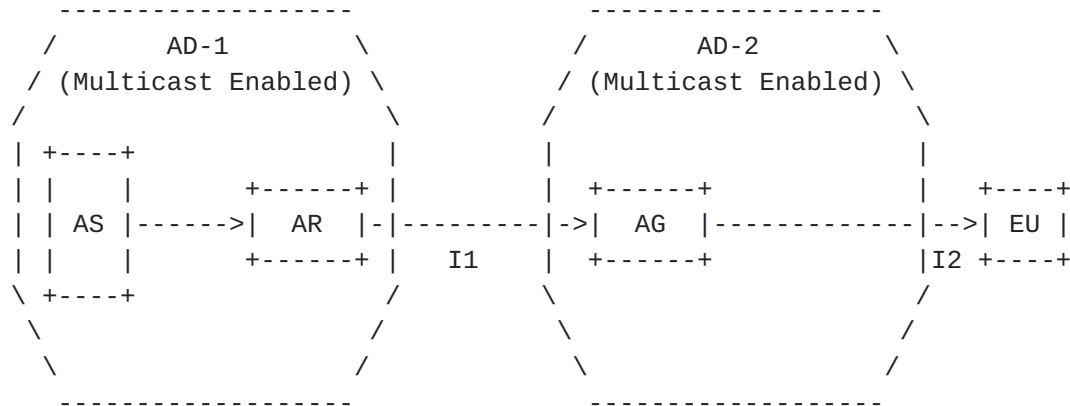
Guidelines (a) through (d) are the same as those described in Use Case 3.1. Two additional guidelines are as follows:

- e. GRE tunnels are typically configured manually between peering points to support multicast delivery between domains.
- f. It is recommended that the GRE tunnel (tunnel server) configuration in the source network is such that it only advertises the routes to the application sources and not to the entire network. This practice will prevent unauthorized delivery of applications through the tunnel (e.g., if application - e.g., content - is not part of an agreed inter-domain partnership).

3.3. Peering Point Enabled with an AMT - Both Domains Multicast Enabled

Both administrative domains in this Use Case are assumed to be native multicast enabled here; however, the peering point is not.

The peering point is enabled with an Automatic Multicast Tunnel. The basic configuration is depicted in Figure 2.



AR = AMT Relay

AG = AMT Gateway

I1 = AMT Interconnection between AD-1 and AD-2

I2 = AD-2 and EU Multicast Connection

Figure 2: - AMT Interconnection between AD-1 and AD-2

Advantages of this configuration:

- o Highly efficient use of bandwidth in AD-1.
- o AMT is an existing technology and is relatively simple to implement. Attractive properties of AMT include the following:
 - o Dynamic interconnection between Gateway-Relay pair across the peering point.
 - o Ability to serve clients and servers with differing policies.

Disadvantages of this configuration:

- o Per Use Case 3.1 (AD-2 is native multicast), current router technology cannot count the number of end users or the number of bytes transmitted to all end users.

- o Additional devices (AMT Gateway and Relay pairs) may be introduced into the path if these services are not incorporated in the existing routing nodes.
- o Currently undefined mechanisms for the AG to automatically select the optimal AR.

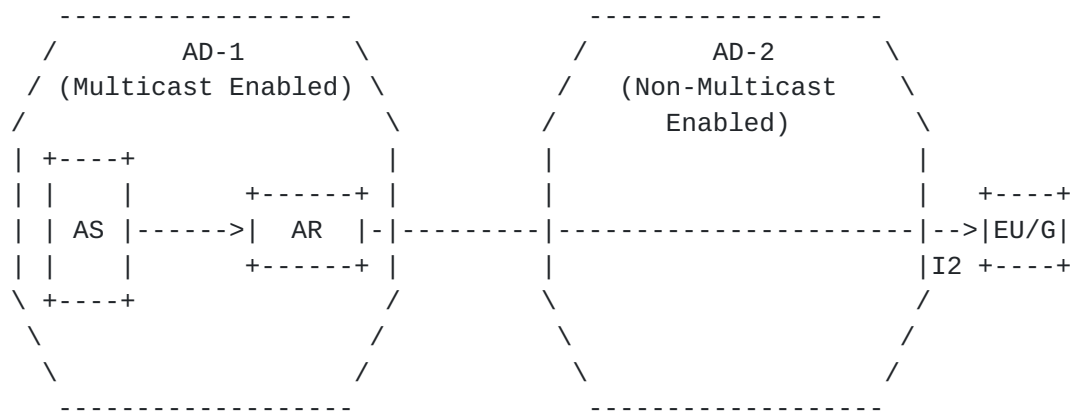
Architectural guidelines for this configuration are as follows:

Guidelines (a) through (d) are the same as those described in Use Case 3.1. In addition,

- e. It is recommended that AMT Relay and Gateway pairs be configured at the peering points to support multicast delivery between domains. AMT tunnels will then configure dynamically across the peering points once the Gateway in AD-2 receives the (S, G) information from the EU.

3.4. Peering Point Enabled with an AMT - AD-2 Not Multicast Enabled

In this AMT Use Case, the second administrative domain AD-2 is not multicast enabled. Hence, the interconnection between AD-2 and the End User is also not multicast enabled. This Use Case is depicted in Figure 3.



AS = Application Multicast Source

AR = AMT Relay

EU/G = Gateway client embedded in EU device

I2 = AMT Tunnel Connecting EU/G to AR in AD-1 through Non-Multicast Enabled AD-2.

Figure 3: - AMT Tunnel Connecting AD-1 AMT Relay and EU Gateway

This Use Case is equivalent to having unicast distribution of the application through AD-2. The total number of AMT tunnels would be

equal to the total number of End Users requesting the application. The peering point thus needs to accommodate the total number of AMT tunnels between the two domains. Each AMT tunnel can provide the data usage associated with each End User.

Advantages of this configuration:

- o Highly efficient use of bandwidth in AD-1.
- o AMT is an existing technology and is relatively simple to implement. Attractive properties of AMT include the following:
 - o Dynamic interconnection between Gateway-Relay pair across the peering point.
 - o Ability to serve clients and servers with differing policies.
- o Each AMT tunnel serves as a count for each End User and is also able to track data usage (bytes) delivered to the EU.

Disadvantages of this configuration:

- o Additional devices (AMT Gateway and Relay pairs) are introduced into the transport path.
- o Assuming multiple peering points between the domains, the EU Gateway needs to be able to find the "correct" AMT Relay in AD-1.

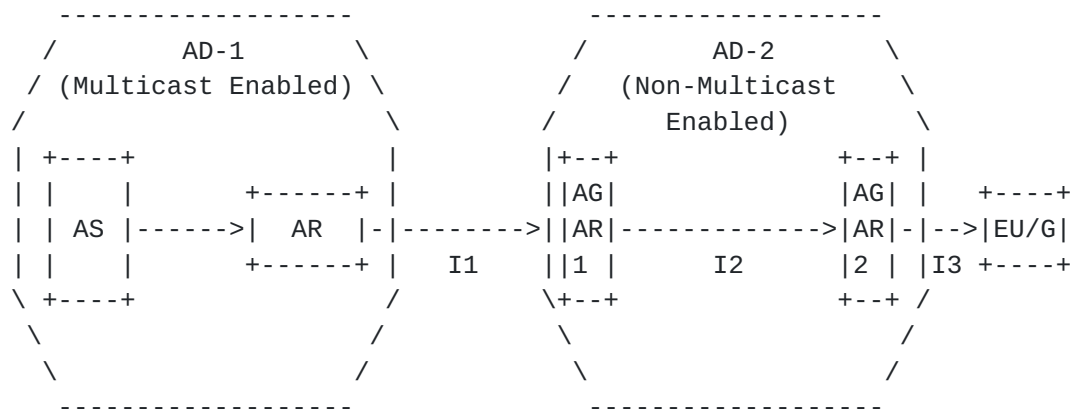
Architectural guidelines for this configuration are as follows:

Guidelines (a) through (c) are the same as those described in Use Case 3.1.

- d. It is recommended that proper procedures are implemented such that the AMT Gateway at the End User device is able to find the correct AMT Relay in AD-1 across the peering points. The application client in the EU device is expected to supply the (S, G) information to the Gateway for this purpose.
- e. The AMT tunnel capabilities are expected to be sufficient for the purpose of collecting relevant information on the multicast streams delivered to End Users in AD-2.

3.5. AD-2 Not Multicast Enabled - Multiple AMT Tunnels Through AD-2

This is a variation of Use Case 3.4 as follows:



AS = Application Source

AR = AMT Relay in AD-1

AGAR1 = AMT Gateway/Relay node in AD-2 across Peering Point

I1 = AMT Tunnel Connecting AR in AD-1 to GW in AGAR1 in AD-2

AGAR2 = AMT Gateway/Relay node at AD-2 Network Edge

I2 = AMT Tunnel Connecting Relay in AGAR1 to GW in AGAR2

EU/G = Gateway client embedded in EU device

I3 = AMT Tunnel Connecting EU/G to AR in AGAR2

Figure 4: - AMT Tunnel Connecting AMT Relay and Relays

Use Case 3.4 results in several long AMT tunnels crossing the entire network of AD-2 linking the EU device and the AMT Relay in AD-1 through the peering point. Depending on the number of End Users, there is a likelihood of an unacceptably large number of AMT tunnels - and unicast streams - through the peering point. This situation can be alleviated as follows:

- o Provisioning of strategically located AMT nodes at the edges of AD-2. An AMT node comprises co-location of an AMT Gateway and an AMT Relay. One such node is at the AD-2 side of the peering point (node AGAR1 in Figure 4).
- o Single AMT tunnel established across peering point linking AMT Relay in AD-1 to the AMT Gateway in the AMT node AGAR1 in AD-2.
- o AMT tunnels linking AMT node AGAR1 at peering point in AD-2 to other AMT nodes located at the edges of AD-2: e.g., AMT tunnel I2 linking AMT Relay in AGAR1 to AMT Gateway in AMT node AGAR2 in Figure 4.
- o AMT tunnels linking EU device (via Gateway client embedded in device) and AMT Relay in appropriate AMT node at edge of AD-2: e.g., I3 linking EU Gateway in device to AMT Relay in AMT node AGAR2.

The advantage for such a chained set of AMT tunnels is that the total number of unicast streams across AD-2 is significantly reduced, thus freeing up bandwidth. Additionally, there will be a single unicast stream across the peering point instead of possibly, an unacceptably large number of such streams per Use Case 3.4. However, this implies that several AMT tunnels will need to be dynamically configured by the various AMT Gateways based solely on the (S,G) information received from the application client at the EU device. A suitable mechanism for such dynamic configurations is therefore critical.

Architectural guidelines for this configuration are as follows:

Guidelines (a) through (c) are the same as those described in Use Case 3.1.

- d. It is recommended that proper procedures are implemented such that the various AMT Gateways (at the End User devices and the AMT nodes in AD-2) are able to find the correct AMT Relay in other AMT nodes as appropriate. The application client in the EU device is expected to supply the (S, G) information to the Gateway for this purpose.
- e. The AMT tunnel capabilities are expected to be sufficient for the purpose of collecting relevant information on the multicast streams delivered to End Users in AD-2.

4. Functional Guidelines

Supporting functions and related interfaces over the peering point that enable the multicast transport of the application are listed in this section. Critical information parameters that need to be exchanged in support of these functions are enumerated, along with guidelines as appropriate. Specific interface functions for consideration are as follows.

4.1. Network Interconnection Transport and Security Guidelines

The term "Network Interconnection Transport" refers to the interconnection points between the two Administrative Domains. The following is a representative set of attributes that will need to be agreed to between the two administrative domains to support multicast delivery.

- o Number of Peering Points.
- o Peering Point Addresses and Locations.

- o Connection Type - Dedicated for Multicast delivery or shared with other services.
- o Connection Mode - Direct connectivity between the two AD's or via another ISP.
- o Peering Point Protocol Support - Multicast protocols that will be used for multicast delivery will need to be supported at these points. Examples of protocols include eBGP [[RFC4760](#)] and MBGP [[RFC4760](#)].
- o Bandwidth Allocation - If shared with other services, then there needs to be a determination of the share of bandwidth reserved for multicast delivery. When determining the appropriate bandwidth allocation, parties should consider use of a multicast protocol suitable for live video streaming that is consistent with Congestion Control Principles [[BCP41](#)].
- o QoS Requirements - Delay/latency specifications that need to be specified in an SLA.
- o AD Roles and Responsibilities - the role played by each AD for provisioning and maintaining the set of peering points to support multicast delivery.

[4.2.](#) Routing Aspects and Related Guidelines

The main objective for multicast delivery routing is to ensure that the End User receives the multicast stream from the "most optimal" source [[INF ATIS 10](#)] which typically:

- o Maximizes the multicast portion of the transport and minimizes any unicast portion of the delivery, and
- o Minimizes the overall combined network(s) route distance.

This routing objective applies to both Native and AMT; the actual methodology of the solution will be different for each. Regardless, the routing solution is expected:

- o To be scalable,
- o To avoid/minimize new protocol development or modifications, and
- o To be robust enough to achieve high reliability and automatically adjust to changes/problems in the multicast infrastructure.

For both Native and AMT environments, having a source as close as possible to the EU network is most desirable; therefore, in some cases, an AD may prefer to have multiple sources near different peering points. However, that is entirely an implementation issue.

4.2.1. Native Multicast Routing Aspects

Native multicast simply requires that the Administrative Domains coordinate and advertise the correct source address(es) at their network interconnection peering points(i.e., border routers). An example of multicast delivery via a Native Multicast process across two Administrative Domains is as follows assuming that the interconnecting peering points are also multicast enabled:

- o Appropriate information is obtained by the EU client who is a subscriber to AD-2 (see Use Case 3.1). This information is in the form of metadata and it contains instructions directing the EU client to launch an appropriate application if necessary, as well as additional information for the application about the source location and the group (or stream) id in the form of the "S,G" data. The "S" portion provides the name or IP address of the source of the multicast stream. The metadata may also contain alternate delivery information such as specifying the unicast address of the stream.
- o The client uses the join message with S,G to join the multicast stream [[RFC4604](#)]. To facilitate this process, the two AD's need to do the following:
 - o Advertise the source id(s) over the Peering Points.
 - o Exchange relevant Peering Point information such as Capacity and Utilization.
 - o Implement compatible multicast protocols to ensure proper multicast delivery across the peering points.

4.2.2. GRE Tunnel over Interconnecting Peering Point

If the interconnecting peering point is not multicast enabled and both AD's are multicast enabled, then a simple solution is to provision a GRE tunnel between the two AD's - see Use Case 3.2.2. The termination points of the tunnel will usually be a network engineering decision, but generally will be between the border routers or even between the AD 2 border router and the AD 1 source (or source access router). The GRE tunnel would allow end-to-end native multicast or AMT multicast to traverse the interface. Coordination and advertisement of the source IP is still required.

The two AD's need to follow the same process as described in 4.2.1 to facilitate multicast delivery across the Peering Points.

4.2.3. Routing Aspects with AMT Tunnels

Unlike Native Multicast (with or without GRE), an AMT Multicast environment is more complex. It presents a dual layered problem because there are two criteria that should be simultaneously met:

- o Find the closest AMT relay to the end-user that also has multicast connectivity to the content source, and
- o Minimize the AMT unicast tunnel distance.

There are essentially two components to the AMT specification

AMT Relays: These serve the purpose of tunneling UDP multicast traffic to the receivers (i.e., End-Points). The AMT Relay will receive the traffic natively from the multicast media source and will replicate the stream on behalf of the downstream AMT Gateways, encapsulating the multicast packets into unicast packets and sending them over the tunnel toward the AMT Gateway. In addition, the AMT Relay may perform various usage and activity statistics collection. This results in moving the replication point closer to the end user, and cuts down on traffic across the network. Thus, the linear costs of adding unicast subscribers can be avoided. However, unicast replication is still required for each requesting End-Point within the unicast-only network.

AMT Gateway (GW): The Gateway will reside on an End-Point - this may be a Personal Computer (PC) or a Set Top Box (STB). The AMT Gateway receives join and leave requests from the Application via an Application Programming Interface (API). In this manner, the Gateway allows the End-Point to conduct itself as a true Multicast End-Point. The AMT Gateway will encapsulate AMT messages into UDP packets and send them through a tunnel (across the unicast-only infrastructure) to the AMT Relay.

The simplest AMT Use Case ([section 3.3](#)) involves peering points that are not multicast enabled between two multicast enabled AD's. An AMT tunnel is deployed between an AMT Relay on the AD 1 side of the peering point and an AMT Gateway on the AD 2 side of the peering point. One advantage to this arrangement is that the tunnel is established on an as needed basis and need not be a provisioned element. The two AD's can coordinate and advertise special AMT Relay Anycast addresses with each other. Alternately, they may decide to simply provision Relay addresses, though this would not be an optimal solution in terms of scalability.

Use Cases 3.4 and 3.5 describe more complicated AMT situations as AD-2 is not multicast enabled. For these cases, the End User device needs to be able to setup an AMT tunnel in the most optimal manner. There are many methods by which relay selection can be done including the use of DNS based queries and static lookup tables [[RFC7450](#)]. The choice of the method is implementation dependent and is up to the network operators. Comparison of various methods is out of scope for this document; it is for further study.

An illustrative example of a relay selection based on DNS queries and Anycast IP addresses process for Use Cases 3.4 and 3.5 is described here. Using an Anycast IP address for AMT Relays allows for all AMT Gateways to find the "closest" AMT Relay - the nearest edge of the multicast topology of the source. Note that this is strictly illustrative; the choice of the method is up to the network operators. The basic process is as follows:

- o Appropriate metadata is obtained by the EU client application. The metadata contains instructions directing the EU client to an ordered list of particular destinations to seek the requested stream and, for multicast, specifies the source location and the group (or stream) ID in the form of the "S,G" data. The "S" portion provides the URI (name or IP address) of the source of the multicast stream and the "G" identifies the particular stream originated by that source. The metadata may also contain alternate delivery information such as the address of the unicast form of the content to be used, for example, if the multicast stream becomes unavailable.
- o Using the information from the metadata, and possibly information provisioned directly in the EU client, a DNS query is initiated in order to connect the EU client/AMT Gateway to an AMT Relay.
- o Query results are obtained, and may return an Anycast address or a specific unicast address of a relay. Multiple relays will typically exist. The Anycast address is a routable "pseudo-address" shared among the relays that can gain multicast access to the source.
- o If a specific IP address unique to a relay was not obtained, the AMT Gateway then sends a message (e.g., the discovery message) to the Anycast address such that the network is making the routing choice of particular relay - e.g., closest relay to the EU. (Note that in IPv6 there is a specific Anycast format and Anycast is inherent in IPv6 routing, whereas in IPv4 Anycast is handled via provisioning in the network. Details are out of scope for this document.)

- o The contacted AMT Relay then returns its specific unicast IP address (after which the Anycast address is no longer required). Variations may exist as well.
- o The AMT Gateway uses that unicast IP address to initiate a three-way handshake with the AMT Relay.
- o AMT Gateway provides "S,G" to the AMT Relay (embedded in AMT protocol messages).
- o AMT Relay receives the "S,G" information and uses the S,G to join the appropriate multicast stream, if it has not already subscribed to that stream.
- o AMT Relay encapsulates the multicast stream into the tunnel between the Relay and the Gateway, providing the requested content to the EU.

4.3. Back Office Functions - Provisioning and Logging Guidelines

Back Office refers to the following:

- o Servers and Content Management systems that support the delivery of applications via multicast and interactions between AD's.
- o Functionality associated with logging, reporting, ordering, provisioning, maintenance, service assurance, settlement, etc.

4.3.1. Provisioning Guidelines

Resources for basic connectivity between AD's Providers need to be provisioned as follows:

- o Sufficient capacity must be provisioned to support multicast-based delivery across AD's.
- o Sufficient capacity must be provisioned for connectivity between all supporting back-offices of the AD's as appropriate. This includes activating proper security treatment for these back-office connections (gateways, firewalls, etc) as appropriate.
- o Routing protocols as needed, e.g. configuring routers to support these.

Provisioning aspects related to Multicast-Based inter-domain delivery are as follows.

The ability to receive requested application via multicast is triggered via receipt of the necessary metadata. Hence, this metadata must be provided to the EU regarding multicast URL - and unicast fallback if applicable. AD-2 must enable the delivery of this metadata to the EU and provision appropriate resources for this purpose.

Native multicast functionality is assumed to be available across many ISP backbones, peering and access networks. If, however, native multicast is not an option (Use Cases 3.4 and 3.5), then:

- o EU must have multicast client to use AMT multicast obtained either from Application Source (per agreement with AD-1) or from AD-1 or AD-2 (if delegated by the Application Source).
- o If provided by AD-1/AD-2, then the EU could be redirected to a client download site (note: this could be an Application Source site). If provided by the Application Source, then this Source would have to coordinate with AD-1 to ensure the proper client is provided (assuming multiple possible clients).
- o Where AMT Gateways support different application sets, all AD-2 AMT Relays need to be provisioned with all source & group addresses for streams it is allowed to join.
- o DNS across each AD must be provisioned to enable a client GW to locate the optimal AMT Relay (i.e. longest multicast path and shortest unicast tunnel) with connectivity to the content's multicast source.

Provisioning Aspects Related to Operations and Customer Care are stated as follows.

Each AD provider is assumed to provision operations and customer care access to their own systems.

AD-1's operations and customer care functions must have visibility to what is happening in AD-2's network or to the service provided by AD-2, sufficient to verify their mutual goals and operations, e.g. to know how the EU's are being served. This can be done in two ways:

- o Automated interfaces are built between AD-1 and AD-2 such that operations and customer care continue using their own systems. This requires coordination between the two AD's with appropriate provisioning of necessary resources.
- o AD-1's operations and customer care personnel are provided access directly to AD-2's system. In this scenario, additional

provisioning in these systems will be needed to provide necessary access. Additional provisioning must be agreed to by the two AD's to support this option.

4.3.2. Application Accounting Guidelines

All interactions between pairs of AD's can be discovered and/or be associated with the account(s) utilized for delivered applications. Supporting guidelines are as follows:

- o A unique identifier is recommended to designate each master account.
- o AD-2 is expected to set up "accounts" (logical facility generally protected by login/password/credentials) for use by AD-1. Multiple accounts and multiple types/partitions of accounts can apply, e.g. customer accounts, security accounts, etc.

4.3.3. Log Management Guidelines

Successful delivery of applications via multicast between pairs of interconnecting AD's requires that appropriate logs will be exchanged between them in support. Associated guidelines are as follows.

AD-2 needs to supply logs to AD-1 per existing contract(s). Examples of log types include the following:

- o Usage information logs at aggregate level.
- o Usage failure instances at an aggregate level.
- o Grouped or sequenced application access. performance/behavior/failure at an aggregate level to support potential Application Provider-driven strategies. Examples of aggregate levels include grouped video clips, web pages, and sets of software download.
- o Security logs, aggregated or summarized according to agreement (with additional detail potentially provided during security events, by agreement).
- o Access logs (EU), when needed for troubleshooting.
- o Application logs (what is the application doing), when needed for shared troubleshooting.
- o Syslogs (network management), when needed for shared troubleshooting.

The two AD's may supply additional security logs to each other as agreed to by contract(s). Examples include the following:

- o Information related to general security-relevant activity which may be of use from a protective or response perspective, such as types and counts of attacks detected, related source information, related target information, etc.
- o Aggregated or summarized logs according to agreement (with additional detail potentially provided during security events, by agreement).

4.4. Operations - Service Performance and Monitoring Guidelines

Service Performance refers to monitoring metrics related to multicast delivery via probes. The focus is on the service provided by AD-2 to AD-1 on behalf of all multicast application sources (metrics may be specified for SLA use or otherwise). Associated guidelines are as follows:

- o Both AD's are expected to monitor, collect, and analyze service performance metrics for multicast applications. AD-2 provides relevant performance information to AD-1; this enables AD-1 to create an end-to-end performance view on behalf of the multicast application source.
- o Both AD's are expected to agree on the type of probes to be used to monitor multicast delivery performance. For example, AD-2 may permit AD-1's probes to be utilized in the AD-2 multicast service footprint. Alternately, AD-2 may deploy its own probes and relay performance information back to AD-1.
- o In the event of performance degradation (SLA violation), AD-1 may have to compensate the multicast application source per SLA agreement. As appropriate, AD-1 may seek compensation from AD-2 if the cause of the degradation is in AD-2's network.

Service Monitoring generally refers to a service (as a whole) provided on behalf of a particular multicast application source provider. It thus involves complaints from End Users when service problems occur. EUs direct their complaints to the source provider; in turn the source provider submits these complaints to AD-1. The responsibility for service delivery lies with AD-1; as such AD-1 will need to determine where the service problem is occurring - its own network or in AD-2. It is expected that each AD will have tools to monitor multicast service status in its own network.

- o Both AD's will determine how best to deploy multicast service monitoring tools. Typically, each AD will deploy its own set of monitoring tools; in which case, both AD's are expected to inform each other when multicast delivery problems are detected.
- o AD-2 may experience some problems in its network. For example, for the AMT Use Cases, one or more AMT Relays may be experiencing difficulties. AD-2 may be able to fix the problem by rerouting the multicast streams via alternate AMT Relays. If the fix is not successful and multicast service delivery degrades, then AD-2 needs to report the issue to AD-1.
- o When problem notification is received from a multicast application source, AD-1 determines whether the cause of the problem is within its own network or within the AD-2 domain. If the cause is within the AD-2 domain, then AD-1 supplies all necessary information to AD-2. Examples of supporting information include the following:
 - o Kind of problem(s).
 - o Starting point & duration of problem(s).
 - o Conditions in which problem(s) occur.
 - o IP address blocks of affected users.
 - o ISPs of affected users.
 - o Type of access e.g., mobile versus desktop.
 - o Locations of affected EUs.
- o Both AD's conduct some form of root cause analysis for multicast service delivery problems. Examples of various factors for consideration include:
 - o Verification that the service configuration matches the product features.
 - o Correlation and consolidation of the various customer problems and resource troubles into a single root service problem.
 - o Prioritization of currently open service problems, giving consideration to problem impact, service level agreement, etc.
 - o Conduction of service tests, including one time tests or a series of tests over a period of time.

- o Analysis of test results.
- o Analysis of relevant network fault or performance data.
- o Analysis of the problem information provided by the customer (CP).
- o Once the cause of the problem has been determined and the problem has been fixed, both AD's need to work jointly to verify and validate the success of the fix.
- o Faults in service could lead to SLA violation for which the multicast application source provider may have to be compensated by AD-1. Subsequently, AD-1 may have to be compensated by AD-2 based on the contract.

4.5. Client Reliability Models/Service Assurance Guidelines

There are multiple options for instituting reliability architectures, most are at the application level. Both AD's should work those out with their contract/agreement and with the multicast application source providers.

Network reliability can also be enhanced by the two AD's by provisioning alternate delivery mechanisms via unicast means.

5. Troubleshooting and Diagnostics

Any service provider supporting multicast delivery of content should have the capability to collect diagnostics as part of multicast troubleshooting practices and resolve network issues accordingly. Issues may become apparent or identified either through network monitoring functions or by customer reported problems as described in [section 4.4](#).

It is expected that multicast diagnostics will be collected according to currently established practices [MDH-04]. However, given that inter-domain multicast creates a significant interdependence of proper networking functionality between providers there does exist a need for providers to be able to signal/alert each other if there are any issues noted by either one.

Service providers may also wish to allow limited read-only administrative access to their routers via a looking-glass style router proxy to facilitate the debugging of multicast control state and peering status. Software implementations for this purpose is readily available [[Traceroute](#)], [[I-D.ietf-mboned-mtrace-v2](#)] and can

be easily extended to provide access to commonly-used multicast troubleshooting commands in a secure manner.

The specifics of the notification and alerts are beyond the scope of this document, but general guidelines are similar to those described in [section 4.4](#) (Service Performance and Monitoring). Some general communications issues are stated as follows.

- o Appropriate communications channels will be established between the customer service and operations groups from both AD's to facilitate information sharing related to diagnostic troubleshooting.
- o A default resolution period may be considered to resolve open issues. Alternately, mutually acceptable resolution periods could be established depending on the severity of the identified trouble.

6. Security Considerations

From a security perspective, normal security procedures are expected to be followed by each AD to facilitate multicast delivery to registered and authenticated end users. Additionally:

- o Encryption - Peering point links may be encrypted per agreement for multicast delivery.
- o Security Breach Mitigation Plan - In the event of a security breach, the two AD's are expected to have a mitigation plan for shutting down the peering point and directing multicast traffic over alternative peering points. It is also expected that appropriate information will be shared for the purpose of securing the identified breach.

DRM and Application Accounting, Authorization and Authentication should be the responsibility of the multicast application source provider and/or AD-1. AD-1 needs to work out the appropriate agreements with the source provider.

Network has no DRM responsibilities, but might have authentication and authorization obligations. These though are consistent with normal operations of a CDN to insure end user reliability, security and network security.

AD-1 and AD-2 should have mechanisms in place to ensure proper accounting for the volume of bytes delivered through the peering point and separately the number of bytes delivered to EUs. For example, [[BCP38](#)] style filtering could be deployed by both AD's to

ensure that only legitimately sourced multicast content is exchanged between them.

Authentication and authorization of EU to receive multicast content is done at the application layer between the client application and the source. This may involve some kind of token authentication and is done at the application layer independently of the two AD's. If there are problems related to failure of token authentication when end-users are supported by AD-2, then some means of validating proper working of the token authentication process (e.g., back-end servers querying the multicast application source provider's token authentication server are communicating properly) should be considered. Implementation details are beyond the scope of this document.

7. IANA Considerations

No considerations identified in this document

8. Conclusions

This Best Current Practice document provides detailed Use Case scenarios for the transmission of applications via multicast across peering points between two Administrative Domains. A detailed set of guidelines supporting the delivery is provided for all Use Cases.

For Use Cases involving AMT tunnels (cases 3.4 and 3.5), it is recommended that proper procedures are implemented such that the various AMT Gateways (at the End User devices and the AMT nodes in AD-2) are able to find the correct AMT Relay in other AMT nodes as appropriate. [Section 4.2](#) provides an overview of one method that finds the optimal Relay-Gateway combination via the use of an Anycast IP address for AMT Relays.

9. Acknowledgments

The authors would like to thank the following individuals for their suggestions, comments, and corrections:

Mikael Abrahamsson

Hitoshi Asaeda

Dale Carder

Tim Chown

Leonard Giuliano

Jake Holland

Joel Jaeggli

Albert Manfredi

Stig Venaas

Henrik Levkowetz

10. Change log [RFC Editor: Please remove]

Please see discussion on mailing list for changes before -111.

-11: version in IESG review.

-12: XML'ified version of -11, committed solely to make rfcdiff easier. XML versions hosted on <https://www.github.com/toerless/peering-bcp>

11. References

11.1. Normative References

- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", [RFC 2784](#), DOI 10.17487/RFC2784, March 2000, <<https://www.rfc-editor.org/info/rfc2784>>.
- [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>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.

- [RFC4604] Holbrook, H., Cain, B., and B. Haberman, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", [RFC 4604](#), DOI 10.17487/RFC4604, August 2006, <<https://www.rfc-editor.org/info/rfc4604>>.
- [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, <<https://www.rfc-editor.org/info/rfc4609>>.
- [RFC7450] Bumgardner, G., "Automatic Multicast Tunneling", [RFC 7450](#), DOI 10.17487/RFC7450, February 2015, <<https://www.rfc-editor.org/info/rfc7450>>.
- [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, <<https://www.rfc-editor.org/info/rfc7761>>.
- [BCP38] Ferguson, P., et al, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP: 38, May 2000.
- [BCP41] Floyd, S., "Congestion Control Principles", [BCP 41](#), September 2000.

11.2. Informative References

- [INF_ATIS_10] "CDN Interconnection Use Cases and Requirements in a Multi-Party Federation Environment", ATIS Standard A-0200010, December 2012.
- [Traceroute] , <<http://traceroute.org/#source%20code>>.
- [I-D.ietf-mboned-mtrace-v2] Asaeda, H., Meyer, K., and W. Lee, "Mtrace Version 2: Traceroute Facility for IP Multicast", [draft-ietf-mboned-mtrace-v2-20](#) (work in progress), October 2017.

Authors' Addresses

Percy S. Tarapore (editor)
AT&T

Phone: 1-732-420-4172
Email: tarapore@att.com

Robert Sayko
AT&T

Phone: 1-732-420-3292
Email: rs1983@att.com

Greg Shepherd
Cisco

Email: shep@cisco.com

Toerless Eckert (editor)
Futurewei Technologies Inc.

Email: tte@cs.fau.de

Ram Krishnan
SupportVectors

Email: ramkri123@gmail.com

