

MBONED
Internet-Draft
Intended status: Informational
Expires: April 30, 2015

W. Atwood
Concordia University/CSE
B. Li
Normal College of Shenzhen University
S. Islam
North South University
October 27, 2014

Receiver Access Control using PANA in IP Multicast
draft-atwood-mboned-mrac-pana-00

Abstract

Multicast Receiver Access Control must be enforced at both the application level and at the network level. The control at the two levels must be correlated, to ensure that only a legitimate group member at the application level is permitted to join group at the network level. We assume that authentication and authorization at the application level are provided by Extensible Authentication Protocol (EAP) exchanges. We describe how to use Protocol for carrying Authentication for Network Access (PANA) to transport the EAP packets in such a way that authentication and authorization can be easily achieved at the network level and that the necessary coordination is achieved.

Status of This Memo

This Internet-Draft is submitted to IETF 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 30, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	2
1.1. Terminology	3
2. Framework for Multicast Receiver Access Control	3
2.1. PANA Role Description	3
2.2. MRAC Role Description	4
2.3. Framework	5
3. Receiver Access Control Process in IP Multicast	6
3.1. Handshake Phase	6
3.2. Authentication and Authorization Phase	6
3.3. Access Phase	7
3.4. Re-authentication Phase	7
3.5. Termination Phase	7
4. Cryptographic Keys	7
5. IANA Considerations	8
6. References	8
6.1. Normative References	8
6.2. Informative References	8
Authors' Addresses	9

1. Introduction

This document is part of a series that proposes a solution for Multicast Receiver Access Control (MRAC). The reader is encouraged to read the other documents in the series to gain insight into how the various components of the solution interact.

A list of desirable properties for MRAC is presented in [I-D.atwood-mboned-mrac-req]. A possible architecture for achieving these properties is presented in [I-D.atwood-mboned-mrac-arch]. The use of the keys described in this document by a secure version of IGMP is presented in [I-D.atwood-pim-sigmp]. The corresponding use of the keys for a secure version of MLD will be presented in draft-atwood-pim-smld (not yet published). The required coordination of keys and security associations among the End User(s) and the router(s) on a network segment is described in [I-D.atwood-pim-gsam].

MRAC can be viewed at two levels: the application level and the network level. At the application level, an End User will obtain permission to subscribe to a group session. This permission will

contain at least two components: a description of how the session is to be accessed and a certification that the End User is authorized to access the session.

The certification will be presented at the application level. If it is valid the End User will be permitted to join the group.

At the network level, the session descriptor will be used to issue the network level join, which allows the session data to flow to the End User device.

To prevent the End User from presenting an arbitrary session descriptor, it is necessary to coordinate the application level join and the network level join.

This draft describes how to achieve receiver access control at the application level, using Protocol for carrying Authentication for Network Access (PANA) [RFC5191] in IP multicast. The approach uncouples the receiver access control from the process of joining a multicast group. An End User is authenticated and authorized at the application level while he/she shows his/her interest in a multicast group at the network level.

This draft does not conflict with or intend to replace [RFC3740] published by the Multicast Security (MSEC) working group. Encryption for multicast data could also be implemented in addition to receiver access control if the multicast application requires it.

1.1. Terminology

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

2. Framework for Multicast Receiver Access Control

2.1. PANA Role Description

There are four roles in a PANA environment: the PANA client (PaC), the PANA authentication agent (PAA), the AAA server (AAAS) and the enforcement point (EP). They are briefly described as follows:

PaC: It is the client implementation of PANA. A PaC runs on a device operated by an End User that wishes to be authorized to perform some action.

PAA: It is the server implementation of PANA. The PAA interacts with a PaC to convey the EAP exchanges that are necessary for

authentication, authorization and accounting. The location of a PAA may be one or more IP hops away from the PaCs that it is responsible for.

AAAS: It is a server that handles authentication, authorization and accounting service. The AAAS interprets the certification presented by a PaC. According to the interpreted information, the AAAS authenticates and authorizes a PaC to perform the action that it has requested.

EP: It is the point in the network where the limitations on the desired action are enforced.

2.2. MRAC Role Description

In the architecture described in [I-D.atwood-mboned-mrac-arch], the End User is required to obtain a ticket from a multicast service provider, which authorizes him/her to participate in the multicast group. The ticket contains the certification and the session descriptor mentioned in Section 1. The certification is carried as a payload by EAP, and the EAP message is presented to the PaC for transmission to the PAA. The certification will be (in most cases) validated by the AAAS.

The PAA is responsible for managing the negotiation with a PaC, usually with the help of the AAAS, that will authorize the End User to join the multicast group. Once the authorization has been achieved, the PAA is responsible for informing the EP that the access to the multicast group can be permitted, and what its responsibilities are with regard to accounting.

The AAAS is responsible for validating the certification, and determining what accounting information must be collected related to the received multicast traffic.

The EP is responsible for allowing authorized PaCs to receive the multicast data while preventing others from doing so. In the case of controlling access to multicast groups, the EP is actually the access router (AR) that is one hop away from the PaC. In a multicast network, the multicast routing protocol designates one AR, called the Designated Router (DR), to join multicast groups on behalf of an End User device. It is clear that the DR takes the role of the EP to enforce the access to multicast groups. Since any AR is potentially designated as a DR, all ARs are considered as (potential) EPs.

To distinguish the role of the EP in this MRAC case from the role of EPs in general in the PANA environment, we will use the designator

mEP. In the simple case, there is only one mEP in the network segment (on the DR) and the PAA resides on the DR.

2.3. Framework




Figure 1: Receiver Access Control Framework

In an IP multicast network, the MRAC framework is as shown in Figure 1. A PaC interacts with a PAA using PANA, which carries the EAP authentication method to request the authorization for a multicast group. A PAA consults a AAAS for authentication and authorization of a PaC according to the EAP method carried in PANA. Moreover a PAA also communicates with the pertinent mEPs to forward the authorization attributes (mainly including the key derived from the EAP authentication method) of a PaC to the pertinent mEPs if a PaC is authorized. The mEPs establish IPsec SAs [RFC4301] with the authorized PaC according to the authorized attributes using a protocol called group security association management (GSAM) [I-D.atwood-pim-gsam]. A PaC then shows its interest in a multicast group to its mEP using the Secure IGMP (SIGMP) protocol [I-D.atwood-pim-sigmp] in an IPv4 network or using the Secure MLD (SMLD) protocol in an IPv6 network. The SIGMP packet or the SMLD packet is secured by an IPsec SA established during the GSAM negotiations. The mEP will join the multicast group on behalf of the authenticated and authorized PaCs and then send the multicast data to them. Usually the PaC is not allowed to forward the multicast data to any other device. The way to prevent forwarding is out of scope for this document.

If the PAA and the AAAS reside in the same device, an API is used to communicate between the PAA and the AAAS. Otherwise, a AAA protocol is usually needed, e.g., Diameter. Similarly, if the PAA and the mEP reside in the same device, an API is used to communicate between the PAA and the mEP. Otherwise, a PAA-to-EP protocol is needed. Possible candidates include, but are not limited to, COPS, SNMP, Diameter, etc.

This draft shows how to share the authorization attributions between a PaC and its mEP. The creation of the IPsec SAs and the protocols SIGMP and SMLD are out-of-scope for this draft. These topics are discussed in [I-D.atwood-pim-gsam], [I-D.atwood-pim-sigmp] and draft-atwood-pim-smld (not yet published) respectively.

3. Receiver Access Control Process in IP Multicast

As defined in [RFC5191], a PANA session has five phases. In our MRAC system, the five phases are explained as follows.

3.1. Handshake Phase

The handshake phase is triggered when a PaC receives the request to establish IPsec SAs in its local GSAM instance. In this phase, a PaC sends a PANA-Client-Initiation message to the PAA to initiate a PANA session. In MRAC, only PaC may initiate a PANA session rather than both a PAA and a PaC as described in [RFC5191]. A PaC may use its local configuration or DHCP to discover its PAA.

3.2. Authentication and Authorization Phase

Immediately following the handshake phase, a PAA and a PaC interact using both PANA-Auth-Request message and PANA-Auth-Answer message in the authentication and authorization phase. The EAP method is carried in the PANA message. The certification that the PaC possesses is encapsulated in one of the EAP packets and is delivered from a PaC to a PAA. A PAA will consult the AAAS using an API or a AAA protocol for authentication and authorization based on the EAP method and then convey the result to a PaC.

On successful authentication, a PANA Master Session Key (PANA MSK) becomes known to the PAA and the PaC as a result of the EAP exchanges. At the network side, the PAA must combine the PANA MSK with EP-specific information to produce the PaC-EP Master Key (PEMK), which is then forwarded (securely) to the pertinent mEP(s) using an API (when the PAA and the mEP are located in the same device) or a PAA-to-EP protocol (when the PAA and mEP are located in different devices). The rules for doing this are specified in [RFC5807]. The mEP must, in turn, combine this PEMK with group-specific information to produce the Multicast Session-Specific Key (MSSK), which will be used in GSAM to establish an IPsec SA to permit the network-level joining of the End User. On the End User device, the PaC must store the PANA MSK itself, since it does not know the identification of its mEP at this time. On the End User device, the work to calculate the MSSK based on the PANA MSK is assigned to GSAM. The details of how to calculate the PEMK and the MSSK will be shown in Section 4.

In order to achieve strong security, the EAP method carried in the PANA messages is required to provide the function of dynamic key exchange. Here the EAP-FAST method is recommended.

At the end of this phase, both a PaC and its mEP have calculated the keys (PaC calculates PANA MSK and mEP calculates MSSK) for authentication and authorization at the network layer. For the network layer join, the PaC will use an SIGMP message protected by an IPsec SA to show its interest in a specific group that has been authorized at the application layer. The details of the network layer interactions may be found in [I-D.atwood-pim-sigmp] and [I-D.atwood-pim-gsam].

3.3. Access Phase

On the one hand, a PaC and a PAA use the PANA message to test peer liveness in a PANA session. On the other hand, the multicast data is distributed from the mEP to the network on which the PaC resides.

3.4. Re-authentication Phase

The re-authentication phase is triggered when the access lifetime specified as the PANA session lifetime needs to be extended. In this phase, the EAP authentication carried in PANA messages is re-executed between a PaC and a PAA.

Upon successful re-authentication, access control re-enters the access phase. The lifetime of the PANA session is extended. Moreover a PAA notifies the pertinent mEPs to extend the lifetime of the authentication attributes of the PaC. However, if re-authentication fails, the PANA session must be terminated. Moreover, a PAA notifies mEPs to revoke the access authorization of a PaC.

3.5. Termination Phase

Either a PaC (i.e., disconnect indication) or a PAA (i.e., session revocation) may initiate the termination of the access authorization at any time. On the one hand, they use PANA-Termination-Request and PANA-Termination-Answer message exchanges to terminate the PANA session between them. On the the other hand, a PAA uses an API or a PAA-to-EP protocol to notify mEPs to revoke the access authentication of a PaC.

4. Cryptographic Keys

At the end of the authentication and authorization phase, a PaC and a PAA share the same secret key (PANA MSK). A PAA will derive a separate PaC-EP-Master-Key (PEMK) as follows [RFC5807]:

PEMK = prf+(MSK, "IETF PEMK" | SID | KID | mEPID)

Here, "|" means concatenation of different fields and prf+ is a pseudo-random function defined in [RFC5996]. "IETF PEMK" is the ASCII code representation, SID is a four-octet Session Identifier, KID is associated with the MSK and mEPID is the identifier of the mEP.

A PaC may be authorized to join more than one multicast group in one PANA session. For each multicast group, a separate Multicast Group-Specific Key (MSSK) is derived from the PEMK using the following method:

MSSK = HMAC-SHA-1 (PEMK, "MSSK" | MSSInf | SID | KID | mEPID).

Here, "MSSK" is the ASCII code representation, SID is a four-octet Session Identifier, KID is associated with the PEMK and mEPID is the identifier of the mEP. MSSInf is the multicast group-specific information.

5. IANA Considerations

This document has no actions for IANA.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, May 2008.
- [RFC5807] Ohba, Y. and A. Yegin, "Definition of Master Key between PANA Client and Enforcement Point", RFC 5807, March 2010.

6.2. Informative References

- [I-D.atwood-mboned-mrac-arch] Atwood, B., Li, B., and S. Islam, "Architecture for IP Multicast Receiver Access Control", draft-atwood-mboned-mrac-arch-01 (work in progress), July 2014.

- [I-D.atwood-mboned-mrac-req]
Atwood, B., Islam, S., and B. Li, "Requirements for IP Multicast Receiver Access Control", draft-atwood-mboned-mrac-req-01 (work in progress), July 2014.
- [I-D.atwood-pim-gsam]
Atwood, B. and B. Li, "Group Security Association Management Protocol", draft-atwood-pim-gsam-00 (work in progress), July 2014.
- [I-D.atwood-pim-sigmp]
Atwood, B. and B. Li, "Secure Internet Group Management Protocol", draft-atwood-pim-sigmp-01 (work in progress), July 2014.
- [RFC3740] Hardjono, T. and B. Weis, "The Multicast Group Security Architecture", RFC 3740, March 2004.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

Authors' Addresses

William Atwood
Concordia University/CSE
1455 de Maisonneuve Blvd, West
Montreal, QC H3G 1M8
Canada

Phone: +1(514)848-2424 ext3046
Email: william.atwood@concordia.ca
URI: <http://users.encs.concordia.ca/~bill>

Bing Li
Normal College of Shenzhen University
Nanhai Ave 3688
Shenzhen, Guangdong 518060
China

Phone: +86(0755)26558364
Email: libingice@szu.edu.cn

Salekul Islam
North South University
House 80, Road 8/A, Mirza Golam Hafiz Road
Dhanmondi, Dhaka 1209
Bangladesh

Email: salekul@northsouth.edu

MBONED Working Group
Internet Draft
Intended status: BCP
Expires: April 27, 2015

Percy S. Tarapore
Robert Sayko
AT&T
Greg Shepherd
Toerless Eckert
Cisco
Ram Krishnan
Brocade
October 27, 2014

Multicasting Applications Across Inter-Domain Peering Points
draft-tarapore-mboned-multicast-cdni-07.txt

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 27, 2015.

Copyright Notice

Copyright (c) 2014 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.

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.

Abstract

This document examines the process of transporting applications via multicast across inter-domain peering points. The objective is to describe the setup process for multicast-based delivery across administrative domains and document supporting functionality to enable this process.

Table of Contents

- 1. Introduction.....3
- 2. Overview of Inter-domain Multicast Application Transport.....4
- 3. Inter-domain Peering Point Requirements for Multicast.....5
 - 3.1. Native Multicast.....5
 - 3.2. Peering Point Enabled with GRE Tunnel.....7
 - 3.3. Peering Point Enabled with an AMT - Both Domains Multicast Enabled.....8
 - 3.4. Peering Point Enabled with an AMT - AD-2 Not Multicast Enabled.....9
 - 3.5. AD-2 Not Multicast Enabled - Multiple AMT Tunnels Through AD-2.....11
- 4. Supporting Functionality.....13
 - 4.1. Network Interconnection Transport and Security Guidelines14
 - 4.2. Routing Aspects and Related Guidelines.....15
 - 4.2.1 Native Multicast Routing Aspects.....15
 - 4.2.2 GRE Tunnel over Interconnecting Peering Point.....16
 - 4.2.3 Routing Aspects with AMT Tunnels.....16
 - 4.3. Back Office Functions - Billing and Logging Guidelines...19
 - 4.3.1 Provisioning Guidelines.....19
 - 4.3.2 Application Accounting Billing Guidelines.....20
 - 4.3.3 Log Management Guidelines.....21
 - 4.3.4 Settlement Guidelines.....21
 - 4.4. Operations - Service Performance and Monitoring Guidelines22
 - 4.5. Client Reliability Models/Service Assurance Guidelines...24

5. Security Considerations.....25
6. IANA Considerations.....25
7. Conclusions.....25
8. References.....26
 8.1. Normative References.....26
 8.2. Informative References.....26
9. Acknowledgments.....26

1. Introduction

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 applications offers significant savings for utilization of resources in any given administrative domain. End user demand for such applications is growing. Often, this requires transporting such applications across administrative domains via inter-domain peering points.

The objective of this Best Current Practices document is twofold:

- o Describe the process and establish guidelines for setting up multicast-based delivery of applications across inter-domain peering points, and
- o Catalog all required information exchange between the administrative domains to support multicast-based delivery.

While there are several multicast protocols available for use, this BCP will focus the discussion to those that are applicable and recommended for the peering requirements of today's service model, including:

- o Protocol Independent Multicast - Source Specific Multicast (PIM-SSM) [RFC4607]
- o Internet Group Management Protocol (IGMP) v3 [RFC4604]
- o Multicast Listener Discovery (MLD) [RFC4604]

This BCP is independent of the choice of multicast protocol; it focuses solely on the implications for the inter-domain peering points.

This document therefore serves the purpose of a "Gap Analysis" exercise for this process. The rectification of any gaps identified - whether they involve protocol extension development or otherwise - is beyond the scope of this document and is for further study.

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) [IETF-ID-AMT].
- o The application stream originates at a source in Domain 1.
- o An End User associated with Domain 2 requests the application. It is assumed that the application is suitable for delivery via multicast means (e.g., live steaming of major events, software downloads to large numbers of end user devices, etc.)
- o The request is communicated to the application source which provides the relevant multicast delivery information to the EU device via a "manifest file". At a minimum, this file contains the {Source, Group} or (S,G) information relevant to the multicast stream.
- o The application client in the EU device then joins the multicast stream distributed by the application source in domain 1 utilizing the (S,G) information provided in the manifest file. The manifest file may also contain additional information that the application client can use to locate the source and join the stream.

It should be noted that the second administrative domain - domain 2 - may be an independent network domain (e.g., Tier 1 network operator domain) or it could also be an Enterprise network 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. A comprehensive list of pertinent information that needs to be exchanged between the two domains to support various functions enabling the application transport is provided in section 4.

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 three possible Use Cases for consideration.

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 (MBGP or BGMP)
 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 unicast transmissions.

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:

- o Dual homing for peering points between domains is recommended as a way to ensure reliability with full BGP table visibility.
- o 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.
- o 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.
- o 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.

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 unicast transmissions.
- o Ability to support only partial IP multicast deployments in AD-1 and/or 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 GRE must be in place prior to stream starting.
- o GRE 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.

- o GRE tunnels are typically configured manually between peering points to support multicast delivery between domains.

- o 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 bytes transmitted.
- 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 to select the AR from the AG automatically.


Architectural guidelines for this configuration are as follows:

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

- 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. This implies that the interconnection between AD-2 and the End User is also not multicast enabled as 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.

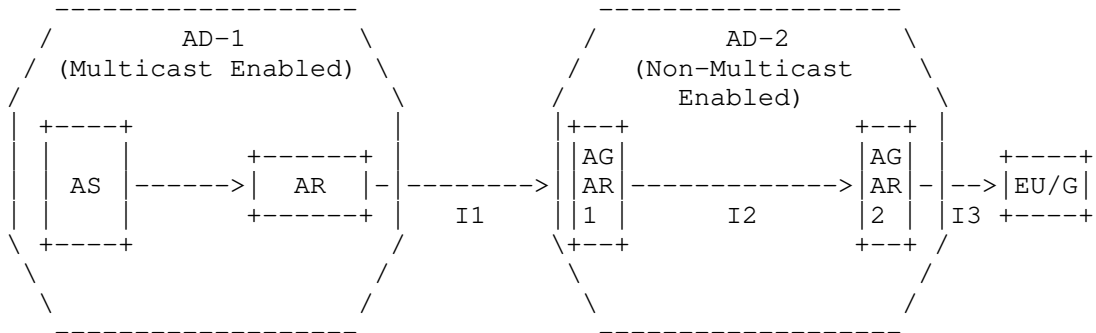
Architectural guidelines for this configuration are as follows:

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:



(Note: Diff-marks for the figure have been removed to improve viewing)

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 AD-1 AMT Relay and EU Gateway

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. Supporting Functionality

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, BGMP, and MBGP.
- o Bandwidth Allocation - If shared with other services, then there needs to be a determination of the share of bandwidth reserved for multicast delivery.
- 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.

From a security perspective, it is expected that normal/typical security procedures will be followed by each AD to facilitate multicast delivery to registered and authenticated end users. Some security aspects for consideration are:

- o Encryption - Peering point links may be encrypted per agreement if dedicated 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 alternated peering points. It is also expected that appropriate information will be shared for the purpose of securing the identified breach.

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 to be:

- o Scalable
- o Avoid/minimize new protocol development or modifications, and
- o 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, but 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 is usually done via an appropriate file transfer - this file is typically known as the manifest file. It contains instructions directing the EU

client to launch an appropriate application if necessary, and also 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 file 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 [RFC2236].

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 (Other??)

4.2.2 GRE Tunnel over Interconnecting Peering Point

If the interconnecting peering point is not multicast enabled and both ADs are multicast enabled, then a simple solution is to provision a GRE tunnel between the two ADs - 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 (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 meet:

- 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:

- o 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 endpoint within the unicast-only network.
- o AMT Gateway (GW): The Gateway will reside on an on 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 endpoint 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 ADs. 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 ADs can coordinate and advertise special AMT Relay Anycast addresses with each other - though they may alternately decide to simply provision Relay addresses, though this would not be a 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. 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. An example of a basic delivery via an AMT Multicast process for these two Use Cases is as follows:

- o The manifest file is obtained by the EU client application. This file 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 manifest file 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 in the manifest file, 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.

Note: Further routing discussion on optimal method to find "best AMT Relay/GW combination" and information exchange between AD's to be provided.

4.3. Back Office Functions - Billing 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 ADs.
- o Functionality associated with logging, reporting, ordering, provisioning, maintenance, service assurance, settlement, etc.

4.3.1 Provisioning Guidelines

Resources for basic connectivity between ADs Providers need to be provisioned as follows:

- o Sufficient capacity must be provisioned to support multicast-based delivery across ADs.
- o Sufficient capacity must be provisioned for connectivity between all supporting back-offices of the ADs 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 the manifest file. Hence, this file must be provided to the EU regarding multicast URL - and unicast fallback if applicable. AD-2 must build manifest and provision capability to provide the file to the EU.

Native multicast functionality is assumed to be available in 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-2s to support this option.

4.3.2 Application Accounting Billing Guidelines

All interactions between pairs of ADs 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 ADs 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 ADs 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)

4.3.4 Settlement Guidelines

Settlements between the ADs relate to (1) billing and reimbursement aspects for delivery of applications, and (2) aggregation, transport, and collection of data in preparation for the billing and

reimbursement aspects for delivery of applications for the Application Provider. At a high level:

- o AD-2 collects "usage" data for AD-1 related to application delivery to End Users, and submits invoices to AD-1 based on this usage data. The data may include information related to the type of content delivered, total bandwidth utilized, storage utilized, features supported, etc.
- o AD-1 collects all available data from partner AD-2 and creates aggregate reports pertaining to responsible Application Providers, and submits subsequent reports to these Providers for reimbursements.
- o AD-1 may convey charging values or charging rules to the AD-2, proactively or in response to a query, especially in cases where these may change.
- o AD-2 may convey prices/rates to AD-1, proactively or in response to a query, especially in cases where these may change.
- o Usage data may be collected per end user or on an aggregated basis; the method of collection will depend on the application delivered and/or the agreements with the source provider. In all cases, usage volume is expected to be in terms of delivered packet bits or bytes.

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. EU's 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. Security Considerations

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.

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. Details will have to be worked out during implementation (e.g., test tokens or trace token exchange process).

6. IANA Considerations

7. 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 4.3 provides an overview of one method that finds the optimal Relay-Gateway combination via the use of an Anycast IP address for AMT Relays.

8. References

8.1. Normative References

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

[IETF-ID-AMT] G. Bumgardner, "Automatic Multicast Tunneling", draft-ietf-mboned-auto-multicast-13, April 2012, Work in progress

[RFC4604] H. Holbrook, et al, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source Specific Multicast", RFC 4604, August 2006

[RFC4607] H. Holbrook, et al, "Source Specific Multicast", RFC 4607, August 2006

8.2. Informative References

[INF_ATIS_10] "CDN Interconnection Use Cases and Requirements in a Multi-Party Federation Environment", ATIS Standard A-0200010, December 2012

9. Acknowledgments

Authors' Addresses

Percy S. Tarapore
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
Phone:
Email: shep@cisco.com

Toerless Eckert
Cisco
Phone:
Email: eckert@cisco.com

Ram Krishnan
Brocade
Phone:
Email: ramk@brocade.com

