MBONED Working Group
Internet Draft
Intended status: BCP
Expires: January 21, 2013

Percy S. Tarapore
Robert Sayko
AT&T
Ram Krishnan
Brocade

October 22, 2012

Multicast Considerations in Support of CDN-I
draft-tarapore-mboned-multicast-cdni-01.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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on January 22, 2013.

Copyright Notice

Abstract

This document examines the current capabilities of multicast to support content distribution in an environment involving multiple Service Providers joining together to form a Content Distribution Network Interconnection (CDN-I) Federation.

Table of Contents

1. Introduction

Content Providers (CP) are experiencing significant growth in demand for all types of internet-based content. A single "over-the-top" CP would require significant resources to deliver content that could be requested from anywhere in the world. Service Providers (SP) are taking advantage of this situation by forming Content Distribution Network (CDN) Federations for the purpose of distributing content on

behalf of Content Providers (CP). There are several advantages to such CDN Federations:

- CPs can simply contract with one or more SPs in a CDN Federation for delivery of their content. This enables CPs to concentrate on their main objective - creation of content.

- SPs can expand their geographic reach via distribution agreements with Federation members without developing costly resources outside their local territories.

Multicast-based delivery mechanisms are a natural fit for content distribution in the proposed CDN Federations. The scope of this document is strictly focused on the interactions between CDN Federation members to support multicast-based content distribution. The purpose of this document is the detailed examination of applicable multicast techniques and the identification of detailed data/metadata/parameters that will have to be exchanged by CDN Federation members to enable multicast-based content distribution.

2. CDN Nomenclature and High Level Overview

Terminology utilized to describe end-to-end user requests is described as follows.

There are many entities involved in distribution of the content from the CP all the way to the End User (EU). Figure 1 is a diagram depicting the basic logical relationships among the various roles involved in content delivery. Besides the CP and EU, the two remaining major entities are the two members of a CDN Federation - the Primary CDN Provider (P-CDN) and the Supporting CDN Provider (S-CDN) [A-0200003]. The relationships between these entities are as follows - see Figure 1.

1. The Content Provider owns the content and specifies conditions of delivery and use. The End User interacts with the CP (link 1 in the figure) for authentication and authorization, and to reach an agreement to obtain specific content (content selection, content purchase, acknowledgement of conditions of use). The CP has the legal right to distribute content and specify conditions for distribution.

2. The CP has an agreement and interacts with the P-CDN for deploying content (link 2). The content is assumed to be retrieved by the P-CDN from the CP Origin server.

3. The P-CDN in turn has an agreement with an S-CDN for deploying and distributing content (link 3).

4. The End User is attached to S-CDN for access and obtains the content from the S-CDN (link 4). The End User also interacts with the CP (link 1) as indicated above.

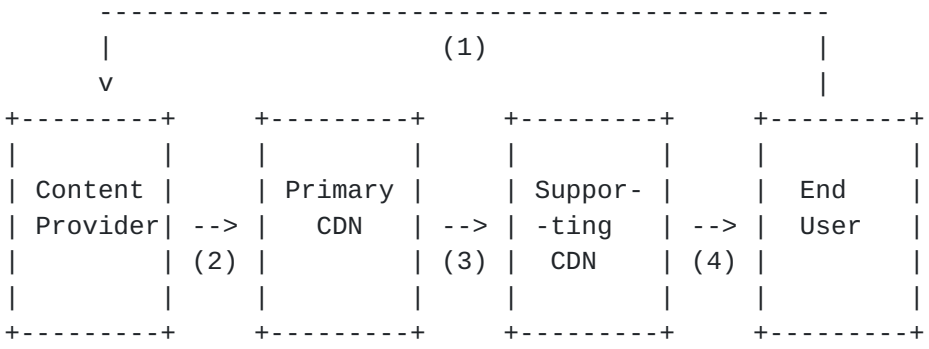

Figure 1 - Relationships in a CDN Federation

Note that all SPs in the CDN Federation can play the role of P-CDN (active relationship with a CP) as well as an S-CDN (attach EUs and distribute content from P-CDN to EUs).

Using the above nomenclature, a CDN Federation is assumed to comprise multiple SPs that agree to function as CDN Providers and distribute content to End Users on behalf of CPs. The term 'SP' in this context is synonymous with the term 'CDN Provider' for the purpose of this document. A CP is assumed to have a contract with one or more SPs for the purpose of content delivery. A high level content request flow is described as follows:

- An End User discovers a specific content on a CP website and requests delivery of the same. The CP selects one CDN Provider with whom it has contracts for content delivery. The rules for CDN Provider selection are part of the CP . CDN Provider contract and are out of scope for this document. The selected CDN Provider functions as the P-CDN.

- The P-CDN then determines the mode of content delivery - either cache/unicast or multicast depending on the type of content and its own capabilities and resources. It then determines whether to deliver the content directly to the End User or utilize another CDN Provider based on rules such as proximity, available resources, common delivery methods (e.g., multicast), etc. These selection rules are out of scope of this document. In case another CDN Provider is selected, then that entity assumes the role of the S-CDN.

- The P-CDN then initiates the delivery process via the S-CDN's domain - the S-CDN acts as the "point of delivery" to the End User.

The simple scenario described here is common to both delivery mechanisms - cache/unicast as well as multicast. The purpose of this document is to address all multicast related aspects in support of multicast-based content delivery.

3. Multicast Use Cases for a CDN-I Federation

Use cases involving multicast methods for distributing content in a CDN Federation have been described in [A-0200004].

3.1. Native Multicast Architectural Use Case

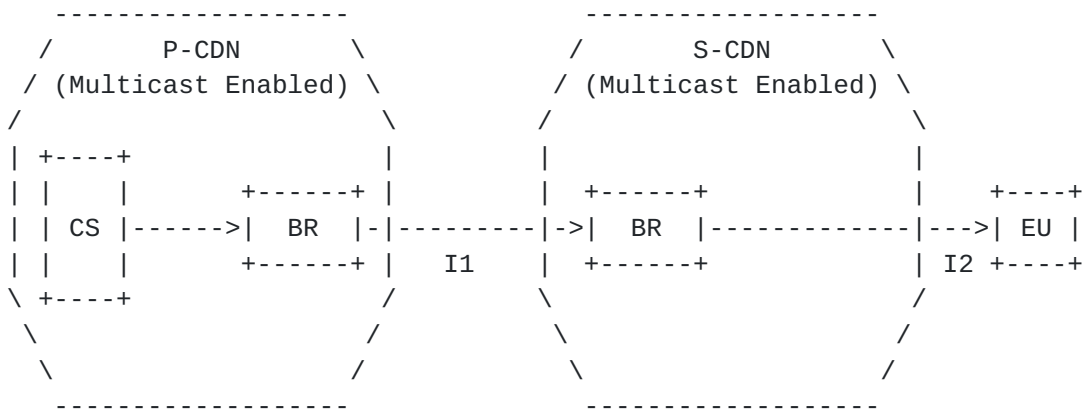

CS = Content Multicast Source (assumes that the Content Source retrieves content from CP Origin)
BR = Border Router
I1 = P-CDN and S-CDN Multicast Interconnection (MBGP or BGMP)
I2 = S-CDN and EU Multicast Connection

Figure 1 - Content Distribution via End to End Native Multicast

This case assumes that both CDN Providers as well as the interconnection between them and the connection between the EU and S-CDN are all multicast enabled.

3.2. A variation of this "pure" Native Multicast case is when the interconnection I1 between the CDNs is multicast enabled via a Generic Routing Encapsulation Tunnel (GRE) [RFC2784] and encapsulating the multicast protocols across the interface. Automatic Multicast Tunneling Architectural Use Cases

In reality, the initial introduction of multicast may not be fully multicast enabled resulting in "Multicast Islands" requiring Automatic Multicast Tunnels (AMT) for enabling multicast connections between them [IETF-ID-AMT].

3.2.1. AMT Interconnection Between P-CDN and S-CDN

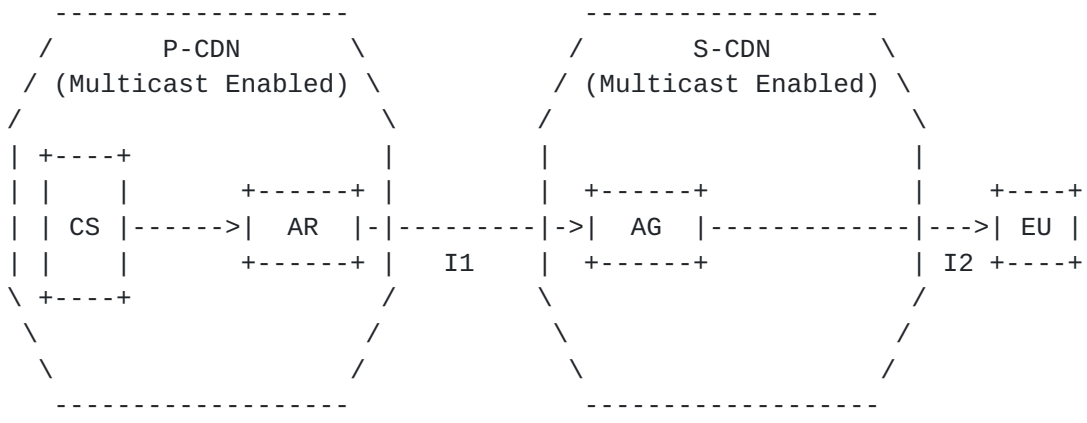

AR = AMT Relay
AG = AMT Gateway
I1 = AMT Interconnection between P-CDN and S-CDN
I2 = S-CDN and EU Multicast Connection

Figure 3 - AMT Interconnection between P-CDN and S-CDN

This configuration assumes both CDN Providers are multicast enabled. Only the interconnection between them is not multicast enabled and hence, an AMT tunnel is established between them as shown in Figure 3.

3.2.2. AMT Tunnel Connecting S-CDN and EU

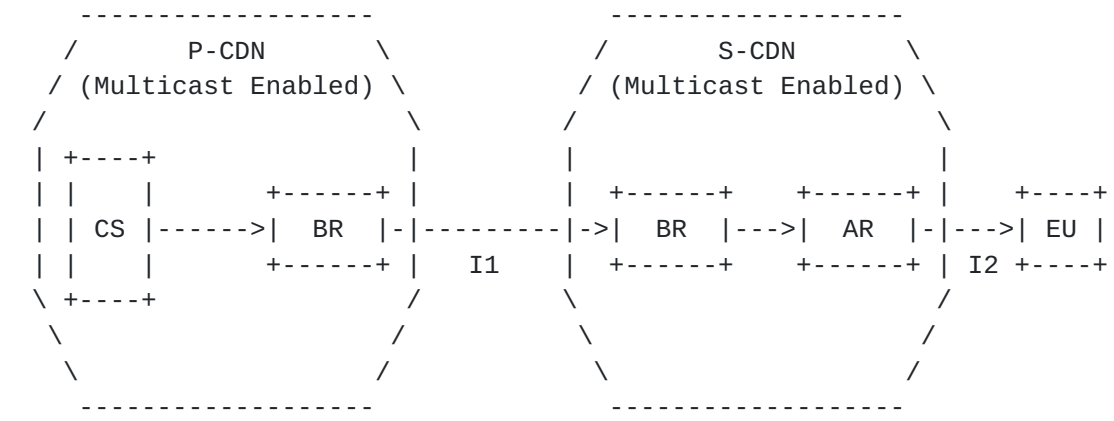

CS = Content Server
BR = Border Router
AR = AMT Relay
I1 = P-CDN and S-CDN Multicast Interconnection (MBGP or BGMP)
I2 = AMT Connection between S-CDN and EU

Figure 4 - AMT Tunnel Connecting S-CDN and EU

This case involves EU devices that are not multicast enabled. Hence an AMT Tunnel is established between the S-CDN AMT Relay and the EU device. This implies one tunnel per EU - potentially several AMT tunnels may need to be setup.
Note that there could be configurations involving both situations described in 3.2.1 and 3.2.2.

3.2.3. AMT Tunnel Connecting EU to P-CDN Through Non-Multicast S-CDN

This Use Case assumes that EU attached to the non-multicast enabled S-CDN has a device populated with a client that establishes an AMT tunnel to the AMT Relay in the P-CDN.

This configuration is needed when the S-CDN is not multicast-enabled. This is the most "extreme" AMT case as the length of the tunnels as well as the number of tunnels can be large.

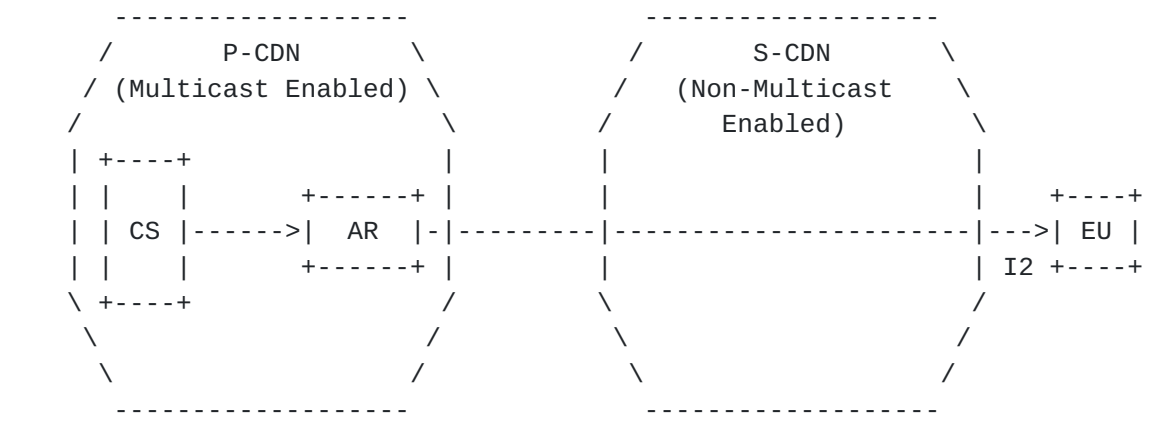

CS = Content Source
AR = AMT Relay
I2 = AMT Tunnel Connecting EU to P-CDN Relay through Non-Multicast Enabled S-CDN.

Figure 5 - AMT Tunnel Connecting P-CDN AMT Relay and EU

4. Content Types Suitable for Multicast-based CDN

This section highlights applications and content types that are suitable for multicast-based delivery in a CDN Federation. Any unique aspects of specific applications/content types that require special attention are duly noted.

Note: Suggest we put in content types in tabular form followed by this text:

For the purpose of this document, live streaming is chosen as a typical multicast application for content distribution via a CDN Federation. Note that a definition of live streaming needs to be provided.

5. Evaluation of Existing Multicast Protocols to Support CDN-I

Use Case 3.1 describes Native Multicast configurations. This is the "simplest" multicast case in that a single standard set of protocols supports end-to-end content delivery from the CP to EU via two or more fully multicast-enabled CDN Providers. It also provides for efficient use of bandwidth and resources.

Use Case 3.2.1 does deploy an AMT Tunnel for interconnecting two CDN Providers; the rest of the configuration is Native Multicast - this assumes that the EU devices are also multicast-enabled.

Thus existing Multicast capabilities need to be examined to determine their ability to fully support content distribution in a CDN Federation. A list of issues requiring examination is as follows:

- Delivery - Identification and communication of {Source, Group} information and DNS information for provisioning across CDNs. Details to be provided.
- Routing/Peering - Identification and acknowledgement of external IP addresses particularly when utilizing a GRE Tunnel for interconnecting CDNs. Details to be provided.
- Back-Office Functions - Identification of appropriate data/metadata collected by Native Multicast to support usage of content for billing, settlements, logging, etc. Details to be provided.
- Security - Determine ability of Native Multicast to deal with security risks such as bot attacks, denial of service, etc. Details to be provided. Also, EU authentication and authorization
- Others - To Be Determined

5.1. Use Case 3.1

In this Use Case, with E-T-E multicast connectivity, the P-CDN provides source, routing and reporting functionality. The S-CDN provides routing and reporting functions.

Though the CP is the origin the interface between the origin and the MC source is out of scope.

Issues:

-what about performance and reliability "tricks" that may be employed do we address here?

-Moving from network to network? Addressed in separate RFC

-IPv4 to IPv6 (assumption) not addressed or in scope - need to reference a standard if there is one

-expand definition of S-CDN???

5.1.1. Delivery

1) If there exists end to end multicast between the EU and the MC Source, then IGMP/MLD, PIM and MBGP protocols will ensure the content stream is delivered to the requesting EU.

2) The (S,G) and in fact all parameters required by the application related to the requested stream can be obtained by the client application in a number of ways. None of which affects the MC protocols and generally neither are the delivery mechanisms of the parameters mutually exclusive. Examples are:

   a. Upon selecting (clicking a link) on a web page for MC content a file (Manifest file) is delivered via unicast to the EU from the web page

   b. An HTTP reply from the web page to the EU request may contain the required parameters

   c. An application may be "programmed" with the (S,G) and other parameters. This method would require that the parameters remain fairly static/stable.

   d. An application may even require an EU to enter the (S,G) and other parameters manually

3) There are a number of ways to control and communicate parameters to the EU that direct the EU to the appropriate MC source. Eg.

   a. The (S,G) can be FQDNs in which case DNS can be used to select specific (S,G) address set.

   b. When the EU selects the stream from the web page the parameters forwarded back to to the EU can be selected based on criteria related to the EU such as source IP, EU SP. Decisions are driven by

      i. Business relationship to the EU and/or EU CDN

ii. Geographic location of SP

iii. On-net or off-net EU

iv. MC network conditions

v. TOD

vi. Etc

c. In ETE MC any intervening (transit) network or SP should not make any attempt to change parameters forwarded to the EU from the serving CDN.

i. Only a selected and participating S-CDN SP should, if required, intercept and change any MC parameters forwarded to the EU. There should be agreement between the participating CDNs that this is permissible.

ii. If the S-CDN changes parameters this should be recorded and reported to the upstream (P-CDN) CDN.

1. Generally reporting to the P-CDN will not be required in real time

2. TBD = situations where it should be Real Time or near RT

iii. An S-CDN should generally avoid unilaterally intercepting and changing parameters values sent from an upstream participating CDN (e.g. P-CDN)because this may cause problems such as the following examples:

1. Performance issues in delivery to the EU

2. Conflicting/erroneous parameter values

3. Conflicting/erroneous authentication, authorization or accounting

4) The MC Source application, which in the case of ETE is the P-CDN, should employ techniques to ensure reliable deliver of the content to the UE. The methodology or combination of methodologies will often depend on the content type being delivered. There are

several effective methods currently available, of which a few examples are described below:

a. Carrouselling

b. FEC

c. Parallel streaming

d. Grid or burst severs

e. Unicast fallback

5.1.2. Routing and Peering

1) Routing between CDNi federation participants should be on a pair-wise agreement and implementation

2) Route/peering preferences of the S-CDN should be communicated to the P-CDN - how can this be done?????? How necessary is it?

a. Allow S-CDN to have some ability to balance peering traffic

b. Route around congestion.

3) CDNi participants must ensure that advertisement of the Source IP address is provided across the peering interface.

a. If an Anycast address is to be used for the Source, then participants must coordinate use of the Anycast address space.

b. Coordination of the address space should be spelled out in the participating members business arrangement/agreement

c. Coordination can occur in a pair-wise manner, but each CDN with multiple arrangements must ensure that there are no conflicts introduced into any individual arrangement.

4) If the peering interface is not multicast enabled a GRE tunnel should be provisioned between the CDNi participants

a. Source IP must be advertised across the GRE tunnel

b. The GRE tunnel may be provisioned to be terminated on participants border routers

c. Alternatively, the P-CDN may desire to terminate the GRE directly to the Source. This will allow the P-CDN to have a separate IP address for external and internal EUs.

5) Is there a way a P-CDN can instruct a S-CDN router to not broadcast its IGMP group memberships across its subnet?????

a. How is that communicated to the router

b. Is it a provisioning only option

c. Or can it be interactively done via some protocol message?

5.1.3. Back Office Functions

5.1.4. Security

5.2. Use Case 3.2.1 (AMT only across the peering interface)

5.2.1. Delivery

All the recommended practices of 5.1.1 are applicable, plus the following additional:

1) The P-CDN and S-CDN should each deploy and provision an AMT Relay at their respective peering interface(s)

2) These AMT Relay must also include and AMT Gateway (GW) function

3) The P-CDN and S-CDN must have an agreement that allows the S-CDN to advertise the Source IP address across the S-CDN multicast enabled network.

4) If the AMT interface fails or is out of service the P-CDN and S-CDN interface agreement should include provision for unicast failover.

5.2.2. Routing and Peering

1) The GW function in the S-CDN AMT Relay can be provisioned with the proper address (this may be an Anycast address) of the P-CDN AMT Relay with connectivity to the source. Alternatively, if the MC source IP (or Anycast address) is not provisioned then a mechanism for the GW function to find the P-CDN Relay with connectivity to the source must be available

   a. A recommended mechanism for the GW to find the correct Relay is documented in "the RFC draft being constructed"

   b. The P-CDN and S-CDN agreements should cover the methodology used

   c. The agreement should specify how notification of MC Source IP is communicated (including timing)

5.2.3.

5.2.4.

6. Security Considerations

TBD

7. IANA Considerations

TBD

8. Conclusions

TBD

9. References

9.1. 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

9.2. Informative References

[A-0200003] P. Tarapore, "CDN Interconnection Use Case Specifications and High Level Requirements", ATIS Standard A-0200003, June 2011 (contact Nicole Butler at nbutler@atis.org using code IETF12 to receive a free copy before September 30, 2012)

[A-0200004] P. Tarapore and R. Sayko, "CDN Interconnection Use Cases and Requirements for Multicast-Based Content Distribution", ATIS Standard A-0200004, January 2012 (contact Nicole Butler at nbutler@atis.org using code IETF12 to receive a free copy before September 30, 2012)

10. 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

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