

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 30, 2018

D. Purkayastha
A. Rahman
D. Trossen
InterDigital Communications, LLC
October 27, 2017

Multicast HTTP using BIER
draft-purkayastha-multicast-http-using-bier-00

Abstract

HTTP Level multicast, using BIER, is described as a use case in [I-D.ietf-bier-use-cases]. In order to enable the use case, the document describes additional functions in the ingress and egress nodes to the BIER network. These functions are assumed to be part of the BIER multicast overlay.

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.

Table of Contents

1. Introduction	2
2. Requirements	2
3. HTTP Multicast Overlay Components	3
4. HTTP Multicast Overlay Operations	4
5. Required Protocol Changes	5
6. IANA Considerations	5
7. Security Considerations	6
8. Informative References	6
Authors' Addresses	6

1. Introduction

BIER Use Cases document [I-D.ietf-bier-use-cases] describes an "HTTP Level Multicast" scenario, where HTTP-level clients benefit from the dynamic multicast group formation enabled by BIER. HTTP multicast means aggregating individual HTTP Responses (e.g. for the same segment of a viral video) and mapping it onto the BIER multicast overlay. For this, the server NAP (network attachment point), creates a list of outstanding client NAP requests to the same HTTP request URI. When a response is available, BIER forwarding information is retrieved and used to send the HTTP response.

In this draft, we introduce the requirements for a BIER multicast overlay realizing this use case. It also describes the necessary functions that form the BIER multicast overlay and the operations that enable the desired "HTTP Level Multicast" behavior. We describe a list of protocol changes needed for the realization of the individual operations.

2. Requirements

A realization for the "HTTP multicast" use case may have the following requirements:

- o MUST support multiple FQDN-based service endpoints to exist in the overlay
- o MUST send FQDN-based service requests at the network level to a suitable FQDN-based service endpoint via policy-based selection of appropriate path information
- o MUST allow for multicast delivery of HTTP response to same HTTP request URI

- o MUST provide direct path mobility, where the path between the egress and ingress NAPs can be determined as being optimal (e.g., shortest path or direct path to a selected instance), is needed to avoid the use of anchor points and further reduce service-level latency

3. HTTP Multicast Overlay Components

Let us formulate the architecture of the BIER multicast overlay for the scenario outlined in [I-D.ietf-bier-use-cases]. This overlay is shown in Figure 1 below.

The multicast overlay is formed by the BFIR and BFER of the BIER layer and the additional NAP and PCE elements shown in the figure. When connecting to a standard IP routed peering network, a special NAP is utilized, shown as the border GW in the figure.

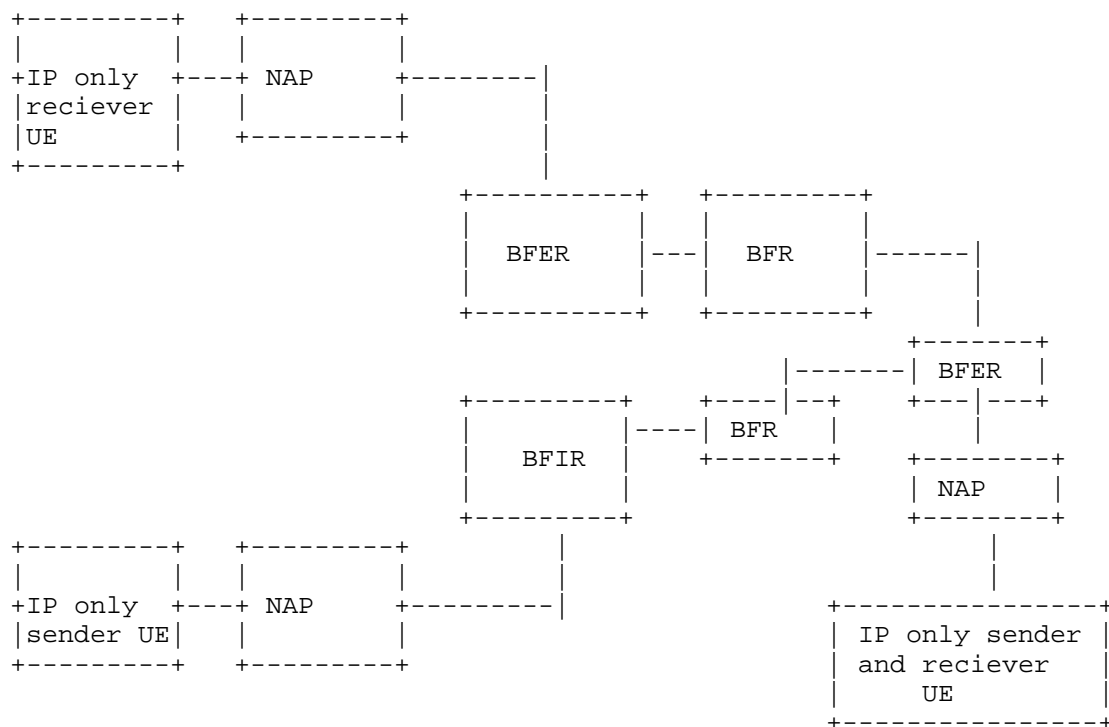


Figure 1: BIER Multicast Overlay for HTTP Multicast Use case

4. HTTP Multicast Overlay Operations

As shown in Figure 1, the multicast overlay includes a function called PCE (Path Computation Element function), which is responsible for selecting the correct multicast end point and possibly realizing path policy enforcement. The result of the selection is a BIER path identifier, which is delivered to the NAP upon initial path computation request (i.e., when sending a request to or response for a specific URL for the first time). The path identifier is utilized for any future request for a given URL-based request. All service end points indicate availability to the PCE through a registration procedure, the PCE will instruct all NAPs to invalidate previous path identifiers to the specific URL. This may result in an initial path computation request at the next service request forwarding. Through this, the newly registered service endpoint might be utilized if the policy-governed path computation selects said service instance.

In the architecture of Figure 1, an HTTP request is sent by an IP-based device towards the FQDN of the server defined in the HTTP request.

At the client facing NAP, the HTTP request is terminated at the HTTP level at a local HTTP proxy. We assume termination on the client side at Layer 3 and above protocols, such as TCP. Server NAP at the egress, terminates any transport protocol on the outgoing (server) side. These terminating functions are assumed to be part of the client/server NAP.

If no local BIER forwarding information exists to the server (NAP), a path computation entity (PCE) is consulted, which calculates a unicast path from the BFIR to which the client NAP is connect to the BFER to which the server NAP is connected. The PCE provides the forwarding information to the client NAP, which in turn caches the result.

Ultimately, the HTTP request is forwarded by the client NAP towards the server-facing NAP via the local BFIR. We assume a (TCP-friendly) transport protocol being used for the transmission between client and server NAP while not mandating the use of TCP for this transmission.

Upon arrival of an HTTP request at the server NAP, the server NAP proxy forwards the HTTP request as a well-formed HTTP request locally to the server.

If no BIER forwarding information exists for the reverse direction towards the requesting client NAP, this information is requested from the PCE, similar to the operation in forward direction.

Upon arrival of any further client NAP request at the server NAP to an HTTP request whose response is still outstanding, the client NAP is added to an internal request table. Optionally, the request is suppressed from being sent to the server.

Upon arrival of an HTTP response at the server NAP, the server NAP consults its internal request table for any outstanding HTTP requests to the same request. The server NAP retrieves the stored BIER forwarding information for the reverse direction for all outstanding HTTP requests and determines the path information to all client NAPs through a binary OR over all BIER forwarding identifiers with the same SI field. This newly formed joint BIER multicast response identifier is used to send the HTTP response across the network.

5. Required Protocol Changes

For the operations outlined in the previous section, we foresee the following protocol changes may be required:

- o NAP-to-NAP protocol for HTTP: Map HTTP to BIER message exchange between client and server NAPs
- o NAP-PCE protocol: Used for path computation and delivery of BIER routing information as well as path updates
- o Overlay transport protocol: Used for transport-level exchange over BIER layer
- o Registration protocol: Used to register FQDN service endpoints
- o Content certificate distribution protocol: Used for HTTPS support

There is a similar ongoing work in SFC WG, which handles HTTP redirection [I-D.purkayastha-sfc-service-indirection]. The lower layers for the NAPs and PCE infrastructure is similar between the two approaches. Does the WG see value in supporting the requirements for BIER to enable HTTP Multicast Use case as defined in [I-D.ietf-bier-use-cases]? It also raises a relevant question, where shall the protocol work be done?

6. IANA Considerations

This document requests no IANA actions.

7. Security Considerations

TBD.

8. Informative References

[I-D.ietf-bier-use-cases]

Kumar, N., Asati, R., Chen, M., Xu, X., Dolganow, A., Przygienda, T., arkadiy.gulko@thomsonreuters.com, a., Robinson, D., Arya, V., and C. Bestler, "BIER Use Cases", draft-ietf-bier-use-cases-05 (work in progress), July 2017.

[I-D.purkayastha-sfc-service-indirection]

Purkayastha, D., Rahman, A., and D. Trossen, "USE CASE FOR HANDLING DYNAMIC CHAINING AND SERVICE INDIRECTION", draft-purkayastha-sfc-service-indirection-00 (work in progress), July 2017.

Authors' Addresses

Debashish Purkayastha
InterDigital Communications, LLC
Conshohocken
USA

Email: Debashish.Purkayastha@InterDigital.com

Akbar Rahman
InterDigital Communications, LLC
Montreal
Canada

Email: Akbar.Rahman@InterDigital.com

Dirk Trossen
InterDigital Communications, LLC
64 Great Eastern Street, 1st Floor
London EC2A 3QR
United Kingdom

Email: Dirk.Trossen@InterDigital.com
URI: <http://www.InterDigital.com/>