Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 4, 2017

G. Eriksson
C. Holmberg
Z. Sarker
Ericsson
J. Reschke
greenbytes

October 31, 2016

Resource Maps
draft-eriksson-http-resource-map-00

Abstract

When the 'out-of-band' content coding ('OOB') is used for delivering a number of resources from a primary server via a secondary server, the additional round trips for OOB responses and load on the primary server can be a significant nuisance.

In such situations, it is useful for the primary server to be able to provide the client with OOB response information for several resources in one go anticipating future client requests.

This document describes a format for providing the client with the information, called a resource map, and how the resource map could be delivered to a client.

Editorial Note (To be removed by RFC Editor before publication)

Distribution of this document is unlimited. Although this is not a work item of the HTTPbis Working Group, comments should be sent to the Hypertext Transfer Protocol (HTTP) mailing list at ietf-http-wg@w3.org [1], which may be joined by sending a message with subject "subscribe" to ietf-http-wg-request@w3.org [2].

Discussions of the HTTPbis Working Group are archived at <http://lists.w3.org/Archives/Public/ietf-http-wg/>.

XML versions, latest edits, and issue tracking for this document are available from <https://github.com/EricssonResearch/Blind-Cache-Drafts>.

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 May 4, 2017.

Copyright Notice

Table of Contents

1. Introduction

The mechanisms outlined in "An Architecture for Secure Content Delegation using HTTP" [SCD] and "Caching Secure HTTP Content using Blind Caches" [BC] use the 'out-of-band' content coding ('OOB') mechanism [OOBENC] to delegate the delivery of requested resource from a primary server to one or several secondary servers.

A primary server might decide to delegate the delivery of response payload for a set of resources, for instance individual video segments, or a set of images or parts of a large file.

In one approach the client sends individual requests for each of the resources to the primary server which provides the OOB content coding response information for the requested resources as meta data in each and every request to the primary server. This approach adds a minimum of one extra RTT (round trip time) for each request before the client can send the request to the desired secondary server.

In another approach the primary server anticipates a client's requests, for instance leveraging object dependency graph information, to provide the client with OOB content coding information for the subsequent requests in advance of requests.

This document describes a format for providing the client with the information for multipe OOB responses, called a resource map, and how the information is delivered to a client.

1.1. Notational Conventions

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. Basic procedure for providing a Resource Map

A primary server creates a resource map containing information about how a set of resources on the origin it serves maps to a set of resource locations on one or multiple secondary servers.

[OOBENC] describes the basic procedure for providing the client with the OOB response meta data for an individual request, as examplified below:

Client request of a resource:

```
GET /test HTTP/1.1
Host: www.example.com
Accept-Encoding: gzip, out-of-band
```

Response from primary server:

```
HTTP/1.1 200 OK
Date: Thu, 14 May 2015 18:52:00 GMT
Content-Type: text/plain
Cache-Control: max-age=10, public
Content-Encoding: out-of-band
Content-Length: 165
Vary: Accept-Encoding

{
  "sr": [
    { "r" :
      "http://example.net/bae27c36-fa6a-11e4-ae5d-00059a3c7a00"},
    { "r" :
      "/c/bae27c36-fa6a-11e4-ae5d-00059a3c7a00"}
  ]
}
```

A primary server can provide a Resource Map location in any response to a request for any resource using a Link header field [RFC5988] -- this can be an OOB response like the one above, but doesn't need to.

```
Link: </map>; rel="http://purl.org/NET/linkrel/resource-map"
```

Note that in this example, "/map" is a relative reference identifying the resource map (relative to the request URI), and "http://purl.org/NET/linkrel/resource-map" is the identifier for the link relation "resource map".

A primary server can speculatively anticipate a client requesting a Resource Map and MAY use HTTP/2 [RFC7540] to push the response with the Resource Map to the client to decrease the extra delay a request for resource map otherwise would incur.

The basic operation for providing the client with a Resource Map is as follows:

1. Client requests resource from primary. Request includes Accept-Encoding header field including "out-of-band".
2. Primary server response includes Link header field identifying the resource map.

3. Client retrieves the resource map.
4. Resource Map becomes stale and will be retrieved again.

3. Updating the Resource Map

3.1. Primary server need to update Resource Map

Depending on the nature of the application and the resources, the primary server's need to provide the client with updates of the resource map information vary. In some cases there is no or little need for updates, in others there may be reasons for more frequent updates.

For instance, delegation information for static resources such as banners, images and script libraries for analytics might not change. In such a situation, the primary server does not need to stay in touch with the client to provide it with updates.

In another use case, the primary server provides the client with updates, for instance changing the secondary servers to be used by a mobile client.

Basically there are four approaches for a primary server to provide a client with updates to a resource map:

- Piggyback signal about update on response to request from client
- Client subscribes to Web Push updates
- Client receives continuous updates
- Resource Map validity time-out

3.2. Piggyback

A primary server MAY add an Link header field to any response to a client. A client receiving a response with a Link header MUST retrieve the Resource Map web resource and process it.

3.3. Using Web Push

A primary server MAY add an indication in the Resource Map to the client to use a notification service to receive updates to a Resource Map. When present a client SHOULD register for notifications.

A client receiving a Resource Map update notification MUST retrieve the Resource Map using the location information in the notification.

3.4. Continuous Updates

In some situations, it is desirable to provide the client with a stream of Resource Map updates. Depending on the client type, either a secure bi-directional transport protocol is used, such as WebSocket, or the client is instructed to poll for changes.

A primary server MAY add an indication to the client to either poll or use a stream oriented protocol for receiving intermittent, frequent updates.

3.5. Resource Map Timer

The freshness of a Resource Map can be given by the 'ma' timer. When expired, the OOB delegation is not valid and the client MUST request an update from the primary server (this duplicates HTTP cache information, but might become important when the resource map is transmitted over a protocol other than HTTP).

3.6. Resource Map compression

3.6.1. gzip

A resource map might be large. For this reason a client can use compression content codings such as 'gzip' to decrease the size.

3.6.2. Formulas

When using the resource map, the primary server could use a template or formula to optimize the size of the resource map information. For example, using URI templates [RFC6570], and an agreed upon HMAC of the path postfix, it would be possible to specify a mapping for many URIs at once.

```
originURI =
  "https://{origin}/images/{postfix}
mappedURI =
  "https://cch.example.com/{origin}/images/{hmac-of-postfix}
```

...would indicate a mapping for any https URI on "origin" below a root path of "/images/", where the mapped URI would be constructed by concatenating "https://cch.example.com/", the origin's host name, "/images/", and a base64- or hex-encoded opaque part, computed based on the remainder of the origin URI's path.

4. The Resource Map format

4.1. Definitions

A resource map will use and extend attributes outlined in [OOBENC] as well as additional extension attributes to be defined later.

The format of the resource map uses JavaScript Object Notation (JSON, [RFC7159]) describing a set of objects:

- Objects for describing handling of updates of a resource map.
- Objects for describing secondary servers and which resources that are mapped.
- Objects for describing the resources on the secondary servers.

```
{
    "resource-map-updates" : {

        //Information about resource map updates.

    },

    "secondary-servers" : {

        //Information about secondary servers.

    },

    "resources" : [

        //Information about resources on secondary servers.

    ]
}
```

4.2. Updates for Resource Map

[[anchor15: Details to be done. Example below illustrates how this could be used.]]

```
{
    "resource-map-updates" : {
        "max-age" : [ ],
        "web push" : {
            "notification service" : [ ]    // URL
            ...                             // Other Web Push details
        },
        "online" : {
            "poll" : {
                "intervall" : [ ]   //Poll
            }
            "connect" : [  ]  // TODO: Relation to URL in Link <map>?
        },
}
```

4.3. Secondary Servers and Mapping of Resources

[[anchor17: Details to be done.]]

5. Security Considerations

All the considerations of [OOBENC] and [SCD] apply.

A resource map can be used to cause the client to request malicious content or perform DoS attacks on a victim secondary server. Clients MUST verify the identity of server providing the Resource Map.

In the case the resource map is delivered as a separate Web resource, the client MUST verify that the server providing the resource map belong to the same authority as the primary server.

The secondary server SHOULD take actions to detect and manage request loops caused by erroneous request from a client.

[[anchor19: TODO: Using tokens for access control to secondary.]]
[[anchor20: Issue: A resource map attribute that is applicable for a set of resources may open up for mixed attacker-controlled data and secrets.]]

6. IANA Considerations

This document currently has no actions for IANA.

7. References

7.1. Normative References

[OOBENC] Reschke, J. and S. Loreto, "'Out-Of-Band' Content Coding for HTTP", draft-reschke-http-oob-09 (work in progress), October 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC5988] Nottingham, M., "Web Linking", RFC 5988, DOI 10.17487/RFC5988, October 2010, <http://www.rfc-editor.org/info/rfc5988>.

[RFC7159] Bray, T., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <http://www.rfc-editor.org/info/rfc7159>.

7.2. Informative References

[BC] Thomson, M., Eriksson, G., and C. Holmberg, "Caching Secure HTTP Content using Blind Caches", draft-thomson-http-bc-01 (work in progress), October 2016.

[RFC6570] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., and D. Orchard, "URI Template", RFC 6570, DOI 10.17487/RFC6570, March 2012, <http://www.rfc-editor.org/info/rfc6570>.

[RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol version 2", RFC 7540, DOI 10.17487/RFC7540, May 2015, <http://www.rfc-editor.org/info/rfc7540>.

[SCD] Thomson, M., Eriksson, G., and C. Holmberg, "An Architecture for Secure Content Delegation using HTTP", draft-thomson-http-scd-02 (work in progress), October 2016.

URIs

[1] <mailto:ietf-http-wg@w3.org>

[2] <mailto:ietf-http-wg-request@w3.org?subject=subscribe>

Appendix A. Multiple Secondary Servers

A primary server can provide multiple secondary servers for retrieving a resource or set of resources.

A Resource Map for the case when resources on "https://origin.example.com:8080" are mapped to two different secondary servers:

```
{
  "secondary-servers": {
    "server1": {
      "name": "blind cache 1",
      "address": "bc.example.com",
      "protocol": "https",
      "port": 8083,
      "description": "Some text for debugging"
    },
    "server2": {
      "name": "blind cache 2",
      "address": "bc2.example.com",
      "protocol": "https",
      "port": 8082
    }
  },

  "resources": [
    {
      "resource-origin": "/ex_jsl.js",
      "mapped-path":
          "/origin.example.com%3A8080/j39jl3jaac/29jfnf0f",
      "attributes": {
        "Content-Type": "application/javascript",
        "Content-Encoding": "aesgcm",
        "MI": ".....",
        "Crypto-Key": "....."
      },
      "mapped": [
        {
          "server": "server1",
          "attributes": {
            "Max-Age": 1800
          }
        },
        {
          "server": "server2",
          "attributes": {
            "Max-Age": 1800
```

```
            }
          }
        ]
      },
      {
        "resource-origin": "/another_resource.txt",
        "mapped-path":
            "/origin.example.com%3A8080/i39jfu2/1njknbs3",
        "attributes": {
          "Content-Type": "text/plain",
          "Content-Encoding": "aesgcm",
          "MI": ".....",
          "Crypto-Key": "....."
        },
        "mapped": [
          {
            "server": "server2",
            "attributes": {
              "Max-Age": 1800
            }
          }
        ]
      }
    ]
  }
```

[[anchor25: in the above example, why does resource-origin include scheme/host/port?]]

Appendix B. Change Log (To be removed before publication as RFC

Nothing yet.

Authors' Addresses

Goran AP Eriksson
Ericsson
Farogatan 6
Stockholm 16480
Sweden

Email: goran.ap.eriksson@ericsson.com

Christer Holmberg
Ericsson

Email: christer.holmberg@ericsson.com

Zaheduzzaman Sarker
Ericsson

Email: zaheduzzaman.sarker@ericsson.com

Julian F. Reschke
greenbytes GmbH
Hafenweg 16
Muenster, NW 48155
Germany

Email: julian.reschke@greenbytes.de
URI: http://greenbytes.de/tech/webdav/