

Network Working Group
Internet-Draft
Intended status: Informational
Expires: June 24, 2013

M. Nottingham
M. Kelly
December 21, 2012

Linked Cache Invalidation
draft-nottingham-linked-cache-inv-04

Abstract

This memo defines two new link types that indicate relationships between resources in terms of cache invalidation, along with a HTTP cache-control extension that takes advantage of those relationships to use them to extend response freshness. Collectively, this is referred to as Linked Cache Invalidation (LCI).

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 June 24, 2013.

Copyright Notice

Copyright (c) 2012 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	3
1.1.	Example	3
2.	Notational Conventions	5
3.	The 'invalidates' Link Relation Type	5
4.	The 'inv-by' Link Relation Type	5
5.	The 'inv-maxage' Response Cache-Control Extension	6
5.1.	inv-maxage Syntax	6
5.2.	inv-maxage Semantics	6
6.	Security Considerations	7
7.	IANA Considerations	8
8.	References	8
8.1.	Normative References	8
8.2.	Informative References	8
Appendix A.	Acknowledgements	8
	Authors' Addresses	8

1. Introduction

In normal operation, a HTTP [[RFC2616](#)] cache will invalidate a stored response if an unsafe request (e.g., POST, PUT or DELETE) is made to its URI. HTTP also provides for such a state-changing request to invalidate related resources (using the Location and Content-Location headers in the response), but this is of limited utility, because those headers have defined semantics, and can only occur once each.

Because of this, it is not practical to make a response that depends on the state of another resource cacheable. For example, an update to a blog entry might change several different resources, such as the user's summary page, the blog's "front" page, the blog's Atom feed, and of course the blog entry itself. If any of these resources is made cacheable, it will not reflect those changes, causing confusion if the user tries to verify that their changes have been correctly applied.

This memo introduces new link relation types [[RFC5988](#)] that allow more fine-grained relationships between resources to be defined, so that caches can invalidate all related representations when the state of one changes. It also introduces a cache-control response extension, so that responses using the relations can be cached by implementations that understand these relations.

1.1. Example

Taking the blog use case described above, imagine that we have the following related resources:

- o `http://example.com/blog/2012/05/04/hi` [the blog entry]
- o `http://example.com/blog/2012/05/04/hi/comments` [full comments for the entry]
- o `http://example.com/blog/` {the blog "home"}
- o `http://example.com/users/bob/` [the user page, listing his entries]

When someone comments on Bob's blog entry, they might send a request like this:

```
POST /cgi-bin/blog.cgi HTTP/1.1
Host: example.com
Content-Type: application/x-www-form-urlencoded
Content-Length: 7890
```

[...]

When the comment is successful, it's typical to redirect the client back to the original blog page, with a response like this:


```
HTTP/1.1 302 Moved Temporarily
Location: http://example.com/blog/2012/05/04/hi
Content-Length: 0
```

Which would invalidate the blog entry URI, as per HTTP's normal operation.

To invalidate the full comments page for the entry, the relationship can be described in that page's response headers:

```
HTTP/1.1 200 OK
Content-Type: text/html
Content-Length: 5555
Link: </blog/2012/05/04/hi>; rel="inv-by"
Cache-Control: no-cache, inv-maxage=600
```

[...]

This declares that whenever the entry page (the target of the link header) changes, this response (the full comments page) changes as well; it's invalidated by the link target.

Note that the full comments page also carries a Cache-Control header that instructs "normal" caches not to reuse this response, but allows those caches that are aware of LCI to consider it fresh for ten minutes.

To invalidate the blog home page and user page, it's impractical to list all of the resources that might change if a new entry is posted; not only are there many of them, but their URLs might not be known when the pages are cached. To address this, the POST response itself (i.e., when the comment is made) can nominate resources to invalidate, using the 'invalidates' relation, making that response:

```
HTTP/1.1 302 Moved Temporarily
Location: http://example.com/blog/2012/05/04/hi
Link: <http://example.com/blog/>; rel="invalidates",
      <http://example.com/users/bob/>; rel="invalidates"
Content-Length: 0
```

Depending on how important it is to see updates on the home page and user page, those responses can either allow caching regardless of support for LCI, like this:

```
Cache-Control: max-age=300
```


... or they can only allow caching by LCI-aware caches, like this:

```
Cache-Control: no-cache, inv-maxage=300
```

Together, these techniques can be used to invalidate a variety of related responses.

It is important to note that the invalidations are only effective in the caches that the client's request stream travels through. This means other caches will continue to serve the "old" content until the invalidation event is propagated to them (see below) or the cached responses become stale.

When multiple caches are close together, the HyperText Caching Protocol (HTCP) [[RFC2756](#)] can be used to propagate invalidation events between them.

2. 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 [RFC 2119](#) [[RFC2119](#)].

This document uses the Augmented Backus-Naur Form (ABNF) notation of [[RFC2616](#)], and explicitly includes the following rules from it: delta-seconds.

3. The 'invalidates' Link Relation Type

The 'invalidates' link relation type allows a response that signifies a state change on the server to indicate one or more associated URIs whose states have also changed.

- o Relation name: invalidates
- o Description: Indicates that when the link context changes, the link target also has changed.
- o Reference: [this document]
- o Notes:

4. The 'inv-by' Link Relation Type

The 'inv-by' link relation type allows a response to nominate one or more other resources that affect the state of the resource it's associated with. That is, when one of the nominated resources changes, it also changes the state of this response's resource.

- o Relation name: inv-by
- o Description: Indicates that when the link target changes, the link's context has also changed.
- o Reference: [this document]
- o Notes:

5. The 'inv-maxage' Response Cache-Control Extension

When present, the 'inv-maxage' cache-control extension indicates the number of seconds that caches who implement Linked Cache invalidation can consider responses fresh for, provided they are not invalidated.

5.1. inv-maxage Syntax

The inv-maxage cache-control extension is parameterised with a numeric argument:

```
"inv-maxage" "=" ( delta-seconds | ( "<" delta-seconds ">" ) )
```

Note that the argument MAY occur in either token or quoted-string form.

If the argument is missing or otherwise does not conform to the BNF rule, it is invalid and MUST be ignored by caches.

If the directive appears more than once in a response, each instance is invalid and MUST be ignored by caches.

5.2. inv-maxage Semantics

HTTP caches MAY, if they fully implement this specification, disregard the HTTP response cache-control directives 'no-cache', 'max-age' and 's-maxage' when a valid 'inv-maxage' cache-control directive is present in a response, using its value as a replacement for max-age.

HTTP caches using inv-maxage to calculate freshness MUST invalidate all stored responses whose request-URIs (after normalisation) are the target of a 'invalidates' link relation contained in a successful response to a state-changing request, provided that they are allowed.

HTTP caches using inv-maxage to calculate freshness MUST invalidate all stored responses containing a 'inv-by' link relation whose target is the current request-URI (after normalisation) upon receipt of a successful response to a state-changing request.

Likewise, HTTP caches using inv-maxage to calculate freshness MUST

invalidate all stored responses containing a 'inv-by' link relation whose target is the content of either the Location or Content-Location response headers (after normalisation) upon receipt of a successful response to a state-changing request.

Here, a response is considered to "contain" a link relation if it is carried in the Link HTTP header [[RFC5988](#)]. I.e., it is not necessary to look at the response body.

"Invalidate" means that the cache will either remove all stored responses related to the effective request URI, or will mark these as "invalid" and in need of a mandatory validation before they can be returned in response to a subsequent request.

A "successful" response is one with a 2xx or redirecting 3xx (e.g., 301, 302, 303, 307) status code.

A "state-changing" request is one with an unsafe method (e.g., POST, PUT, DELETE, PATCH), or one that is not known to be safe.

In this context, "normalisation" means, in the case of a relative request-URI, that it is absolutised using the value of the Host request header and the appropriate protocol scheme.

Finally, an invalidation based upon "invalidates" is "allowed" if the host part of the request-URI (if absolute) or Host request header (if the request-URI is relative) matches the host part of the target URI. This prevents some types of denial-of-service attacks.

Implementations SHOULD effect invalidations when they become aware of changes through other means; e.g., HTCP [[RFC2756](#)] CLR messages, upon invalidations caused by other links (i.e., chained "cascades" of linked invalidations), or when a changed response is seen (such as when HTTP validation is unsuccessful).

6. Security Considerations

Linked Cache Invalidation does not guarantee that invalidations will be effected; e.g., they can be lost due to network issues or cache downtime. Furthermore, it does not guarantee that all caches that understand LCI will be made aware of invalidations that happen, because of how they originate.

Therefore, care should be taken that LCI invalidations are not relied upon (e.g., to purge sensitive content).

Furthermore, while some care is taken to avoid denial-of-service

attacks through invalidation, cache efficiency may still be impaired under certain circumstances (e.g., arranging for one request to invalidate a large number of responses), leading to a reduction in service quality.

7. IANA Considerations

This document registers two entries in the Link Relation Type Registry; see [Section 3](#) and [Section 4](#).

It also registers a HTTP Cache Directive, "inv-maxage"; see [Section 5](#).

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [RFC5988] Nottingham, M., "Web Linking", [RFC 5988](#), October 2010.

8.2. Informative References

- [RFC2756] Vixie, P. and D. Wessels, "Hyper Text Caching Protocol (HTCP/0.0)", [RFC 2756](#), January 2000.

Appendix A. Acknowledgements

Thanks to Michael Hausenblas for his input.

Authors' Addresses

Mark Nottingham

Email: mnot@mnot.net

URI: <http://www.mnot.net/>

Mike Kelly

Email: mike@stateless.co

URI: <http://stateless.co/>