

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: November 28, 2011

M. Nottingham  
May 27, 2011

HTTP Browser Hints  
draft-nottingham-http-browser-hints-01

## Abstract

Over time, Web browsers have adapted how they use HTTP based upon common server configurations and behaviours. While this is necessary in the common case, it can be detrimental for performance and interoperability.

This document establishes a mechanism whereby origin servers can make available hints for browsers about their preferences and capabilities, without imposing overhead on their interactions or requiring support for them.

This is intended to allow browsers to safely optimise connections to servers.

## 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 November 28, 2011.

## Copyright Notice

Copyright (c) 2011 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.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Requirements . . . . .	<a href="#">3</a>
<a href="#">3.</a>	A file format for Browser Hints . . . . .	<a href="#">3</a>
<a href="#">4.</a>	Discovering Browser Hints for a Web site . . . . .	<a href="#">4</a>
<a href="#">5.</a>	Pre-defined Browser Hints . . . . .	<a href="#">4</a>
<a href="#">5.1.</a>	max-conns . . . . .	<a href="#">4</a>
<a href="#">5.2.</a>	pconn-ip . . . . .	<a href="#">5</a>
<a href="#">5.3.</a>	max-pipeline-depth . . . . .	<a href="#">5</a>
<a href="#">5.4.</a>	small-hdrs . . . . .	<a href="#">5</a>
<a href="#">5.5.</a>	no-referer . . . . .	<a href="#">5</a>
<a href="#">5.6.</a>	send-ref . . . . .	<a href="#">6</a>
<a href="#">5.7.</a>	chunk-req-bodies . . . . .	<a href="#">6</a>
<a href="#">5.8.</a>	omit-cookies . . . . .	<a href="#">6</a>
<a href="#">6.</a>	The Ref HTTP Request Header . . . . .	<a href="#">6</a>
<a href="#">7.</a>	Security Considerations . . . . .	<a href="#">6</a>
<a href="#">8.</a>	IANA Considerations . . . . .	<a href="#">6</a>
<a href="#">8.1.</a>	The 'browser-hints' Well-Known URI . . . . .	<a href="#">7</a>
<a href="#">8.2.</a>	The HTTP Browser Hints Registry . . . . .	<a href="#">7</a>
<a href="#">9.</a>	References . . . . .	<a href="#">7</a>
<a href="#">9.1.</a>	Normative References . . . . .	<a href="#">7</a>
<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">8</a>
<a href="#">Appendix A.</a>	Acknowledgements . . . . .	<a href="#">8</a>
	Author's Address . . . . .	<a href="#">8</a>

## 1. Introduction

HTTP [[RFC2616](#)] clients -- especially browsers -- typically use hardcoded values or heuristics to determine how many TCP connections to use to a server, based on common-case server behaviours and limitations.

Likewise, they often send voluminous request headers (e.g., in User-Agent and Allow) because they fear that changing those headers' values will break some sites that depend upon specific values.

These are just two examples of common, conservative behaviour by browsers that is good for interoperability, but potentially bad for performance in certain circumstances.

This memo proposes a mechanism whereby a HTTP server can advertise hints for browsers (and other clients), so that communication with them can be optimised.

It does so by defining a file format for such Browser Hints [Section 3](#), and defining how clients can discover it for a given Web site [Section 4](#). Finally, an extensible vocabulary of hints is defined [Section 5](#).

Feedback for this draft should take place on the [apps-discuss@ietf.org](mailto:apps-discuss@ietf.org) mailing list  
<<https://www.ietf.org/mailman/listinfo/apps-discuss>>.

## 2. Requirements

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](#)].

### 3. A file format for Browser Hints

Browser hints are indicating using a JSON [[RFC4627](#)] formatted file, containing a single object whose member's names are browser hints, as defined by the registry [Section 8.2](#).

For example;

```
{  
  "max-conns": 5,  
  "small-hdrs": true  
}
```

Nottingham

Expires November 28, 2011

[Page 3]

---

Internet-Draft

HTTP Browser Hints

May 2011

By their nature, all browser hints are optional; i.e., browsers are free to ignore them.

### 4. Discovering Browser Hints for a Web site

The hints relevant to a given site can be determined by fetching the URI path `"/.well-known/browser-hints"` for that site.

Typically, clients (especially browsers) will not block other requests to a site while fetching the browser hints (because they're optional); instead, it will usually be done concurrently with other requests, or on idle connections for future use.

In this specification, "site" is scoped by the URI scheme and authority. As such, all of the following are considered to be different sites, and therefore have different browser hints:

- o <http://foo.com/>
- o <https://foo.com/>
- o <http://foo.com:8000/>
- o <http://www.foo.com/>

Clients SHOULD follow HTTP 3xx redirects when retrieving hints.

A successful response is valid for its associated site for as long as it can be cached in HTTP.

If the response has a 200 status code but no explicit freshness

(e.g., a `Cache-Control: max-age` or `Expires:` header), browsers SHOULD cache the response heuristically for a generous fixed period (e.g., 30 days).

## [5.](#) Pre-defined Browser Hints

### [5.1.](#) max-conns

- o Browser Hint Name: max-conns
- o Description: When present, this hint indicates the maximum number of concurrent persistent connections that the site would like clients to use.
- o Value Type: number
- o Contact: mnot@mnot.net

Nottingham

Expires November 28, 2011

[Page 4]

---

Internet-Draft

HTTP Browser Hints

May 2011

### [5.2.](#) pconn-ip

- o Browser Hint Name: pconn-ip
- o Description: When true, this hint indicates that the site allows clients to reuse persistent connections keyed by IP address, rather than by hostname. Note that all sites that are sharing the connection MUST declare this hint for it to be used, and if a transport-layer certificate is in use (e.g., for TLS [[RFC5246](#)]), it MUST be valid for all sites.
- o Value Type: true | false
- o Contact: mnot@mnot.net
- o Specification: [this document]

In other words, if both `www.example.com` and `foo.example.org` resolve to the address `192.0.2.5`, and indicate this hint, then clients can send a request to `www.example.com` and then a request to `foo.example.org` on the same TCP connection to that address.

If any of the sites grouped together for the purposes of pconn-ip declare a max-conns hint, the max-conns value for that address is considered to be the maximum of the declared max-conn hints present.

### [5.3.](#) max-pipeline-depth

- o Browser Hint Name: max-pipeline-depth
- o Description: When present, this hint indicates the maximum number of pipelined requests per connection that the site would like clients to use.
- o Value Type: number
- o Contact: mnot@mnot.net

### [5.4.](#) small-hdrs

- o Browser Hint Name: small-hdrs
- o Description: When true, this hint indicates that clients can omit the Accept and Accept-Charset request headers when communicating with the site, and that they can use a shortened version of the User-Agent header.
- o Value Type: true | false
- o Contact: mnot@mnot.net

### [5.5.](#) no-referer

- o Browser Hint Name: no-referer
- o Description: When true, this hint indicates that clients can omit the Referer HTTP request header when sending requests to the site.

- o Value Type: true | false
- o Contact: mnot@mnot.net

### [5.6.](#) send-ref

- o Browser Hint Name: send-ref
- o Description: When true, this hint indicates that clients can omit the Referer HTTP request header when sending requests to the site, if they instead send a Ref HTTP request header (see [Section 6](#)).
- o Value Type: true | false
- o Contact: mnot@mnot.net

### [5.7.](#) chunk-req-bodies

- o Browser Hint Name: chunk-req-bodies

- o Description: When true, this hint indicates that the server can successfully process a request with a chunk-encoded body; i.e., that it should not return a 411 Length Required status. Note that clients may still require a 411 response status, even in when this hint is present. When false, the server may or may not require a Content-Length on requests with bodies.
- o Value Type: true | false
- o Contact: mnot@mnot.net

#### [5.8.](#) omit-cookies

- o Browser Hint Name: omit-cookies
- o Description: This hint contains a list of cookie [\[RFC6265\]](#) cookie-names which can be omitted in requests sent to this site.
- o Value Type: Array of strings
- o Contact: mnot@mnot.net

### [6.](#) The Ref HTTP Request Header

TBD: relative URI referer header

### [7.](#) Security Considerations

TBD

### [8.](#) IANA Considerations

#### [8.1.](#) The 'browser-hints' Well-Known URI

This document defines the "browser-hints" Well-Known URI [\[RFC5785\]](#).

- o URI suffix: browser-hints
- o Change controller: mnot@mnot.net
- o Specification document(s): [this document]
- o Related information:

## [8.2.](#) The HTTP Browser Hints Registry

This document establishes the HTTP Browser Hints Registry.

New hints are registered First Come First Served (see [[RFC5226](#)]), by sending e-mail to <<mailto:iana@iana.org>> (or using other mechanisms, as established by IANA).

Registration requests MUST use the following template:

- o Browser Hint Name: [name of hint]
- o Description: [description of hint]
- o Value Type: [JSON value type]
- o Contact: [e-mail address(es)]
- o Specification: [optional; reference or URI to more info]

New hints MUST be optional; they cannot place requirements upon implementations.

Likewise, new hints MUST be relevant to browser use cases; other non-browsing hints and metadata would make the hints response undesirably large. However, note that non-browser clients MAY use them.

Finally, new hints MUST NOT make communication non-conformant with HTTP itself; i.e., this is not a mechanism for changing the HTTP protocol in incompatible ways. For example, if a hint indicates that browsers can compress request headers using GZIP, intermediaries that are interposed are likely to fail.

The initial contents of the registry are defined in [Section 5](#).

## [9.](#) References

### [9.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.



Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.

[RFC4627] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", [RFC 4627](#), July 2006.

## [9.2.](#) Informative References

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), August 2008.

[RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", [RFC 5785](#), April 2010.

[RFC6265] Barth, A., "HTTP State Management Mechanism", [RFC 6265](#), April 2011.

## [Appendix A.](#) Acknowledgements

Thanks to Mike Belshe, Poul-Henning Kamp, Anirban Kundu, Patrick McManus, and Steve Souders for their suggestions and feedback.

The author takes all responsibility for errors and omissions.

## Author's Address

Mark Nottingham

Email: [mnot@mnot.net](mailto:mnot@mnot.net)

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