

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 22, 2012

A. Petersson
M. Nilsson
Opera Software
April 20, 2012

Forwarded HTTP Extension
draft-ietf-appsawg-http-forwarded-02

Abstract

This document standardizes an HTTP extension header field that allows proxy components to disclose information lost in the proxying process, e.g., the originating IP address of a request or IP number of the proxy on the user-agent facing interface. Given a trusted path of proxying components, this makes it possible to arrange so that each subsequent component will have access to e.g., all IP addresses used in the chain of proxied HTTP requests.

This document also specifies guidelines for a proxy administrator to anonymize the origin of a request.

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 October 22, 2012.

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. 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	3
2.	Notational Conventions	3
3.	Syntax Notations	3
4.	Forwarded	3
5.	Parameters	4
5.1.	Forwarded by	5
5.2.	Forwarded for	5
5.3.	Forwarded host	5
5.4.	Forwarded proto	6
5.5.	Extensions	6
6.	Node identifiers	6
6.1.	IPv4 and IPv6 identifiers	7
6.2.	The "unknown" identifier	7
6.3.	Obfuscated identifier	7
7.	Implementation considerations	8
7.1.	Example	8
8.	Security considerations	9
8.1.	Header validity and integrity	9
8.2.	Information leak	9
9.	IANA considerations	9
10.	References	10
10.1.	Normative references	10
10.2.	Informative references	10
Appendix A.	Change Log (to be removed by RFC Editor before publication)	10
A.1.	Since draft-petersson-forwarded-for-00	11
A.2.	Since draft-petersson-forwarded-for-01	11
A.3.	Since draft-petersson-forwarded-for-02	11
A.4.	Since draft-ietf-appsawg-http-forwarded-00	11
A.5.	Since draft-ietf-appsawg-http-forwarded-01	11
	Authors' Addresses	12

1. Introduction

In today's HTTP landscape, there are a multitude of different applications acting as a proxy for the user agent and effectively anonymizing the requests to look as if they originated from the proxy IP address or in other ways changing the information in the original request. Examples of such applications include caching, content filtering, content compression, crypto offload, and load balancing. As most of the time this destructive behavior is not the primary purpose, or even a desired effect, a way of disclosing the original information at HTTP level instead of depending on the TCP/IP connection remote IP address and transport port number is needed.

In addition to the above mentioned problems, there may also be issues due to the use of NAT. This is further discussed in [[RFC6269](#)].

A common way to disclose this information is by using the de facto standard header fields such as X-Forwarded-For, X-Forwarded-By, and X-Forwarded-Proto. This document intends to standardize syntax and semantics for disclosing such information. The header field also combines all information within one single header field, making it possible to correlate the header fields to each other. With the header field format described in this document, it is possible to know what information belongs together, given that the proxies are trusted. Such conclusions are not possible to make with the X-Forwarded class of header fields. This new header field also extends the de facto standard of, e.g., X-Forwarded-For with features for which real life deployments have shown a need.

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 [[RFC2119](#)].

3. Syntax Notations

This specification uses the Augmented Backus-Naur Form (ABNF) notation of [[RFC5234](#)] with the list rule extension defined in [Section 3.2.5](#) of [[I-D.ietf-httpbis-p1-messaging](#)].

4. Forwarded

The Forwarded HTTP header field is an OPTIONAL header field that, when used, contains a list of parameter-identifier pairs that

disclose information that is altered or lost when a proxy is involved in the path of the request. This applies to forwarding proxies, as well as reverse proxies. Information passed in this header can be, e.g., the source IP address of the request, the IP address of the incoming interface on the proxy, or whether HTTP or HTTPS is used. If the request is passing through several proxies, each proxy MAY add a set of parameters; it MAY also remove earlier added Forwarded-header fields.

The top-level list is represented as a list of HTTP header field-value as defined in Section 3.2 of [[I-D.ietf-httpbis-p1-messaging](#)]. The first element in this list holds information added by the first proxy, followed by information added by any subsequent proxy. Each field-value is a semicolon-separated list, this sub-list consists of parameter-identifier pairs. Parameter-identifier pairs are grouped together by an equals sign. The parameter names are case-insensitive. The header field can be defined in augmented BNF syntax as:

```
Forwarded    = "Forwarded" ":" LWS Forwarded-v
Forwarded-v  = 1#forwarded-element

forwarded-element =
    [ forwarded-pair ] *( ";" [ forwarded-pair ] )

forwarded-pair = token "=" value
value           = token / quoted-string
```

Example:

```
Forwarded: For=192.0.2.43,"for=[2001:db8:cafe::17]:47011"
Forwarded: proto=https;by=198.51.100.60
```

Given that a proxy wishes to add a Forwarded header field to the outgoing request, if the incoming request has no such header field, the proxy simply adds the header with the list of parameters desired. If, on the other hand, the incoming request has such a header field, the proxy adds a comma and the list of parameters. A proxy MAY remove all Forwarded header fields from a request. It MUST, however, ensure that the correct header field is updated in case of multiple Forwarded header fields.

5. Parameters

This document specifies a number of parameters and valid values for each of them.

- o "by" identifies the user-agent facing interface of the proxy.
- o "for" identifies the node making the request to the proxy.
- o "host" is the host request header-field as received by the proxy.
- o "proto" indicates what protocol was used to make the request.

5.1. Forwarded by

The "by" parameter is used to disclose the interface where the request came in to the proxy server. Typically, the value of this parameter is an IP address and optionally a port number, but it can, however, be some other kind of identifier.

The syntax of a "by" value, after potential quoted-string unescaping, MUST conform to the "node" ABNF described in [Section 6](#).

This is primarily added by reverse proxies that wish to forward this information to the backend server.

5.2. Forwarded for

The "for" parameter is used to disclose information about the user agent that initiated the request. Typically the value of this parameter is an IP address, but it MAY also be some other kind of identifier.

The syntax of a "for" value, after potential quoted-string unescaping, MUST conform to the "node" ABNF described in [Section 6](#).

In a chain of proxy servers where this is fully utilized, the first for-parameter will disclose the user agent where the request first was made, followed by any subsequent proxy identifiers. The last proxy in the chain is not part of the list of for-parameters. The last proxy's IP address, and optionally a port number, are, however, readily available as the remote IP address of the TCP/IP connection.

5.3. Forwarded host

The "host" parameter is used to forward the original value of the "Host" header field. This MAY be used for example by the origin server if a reverse proxy is rewriting the "Host" header field to some internal host name.

The syntax for a "host" value, after potential quoted-string unescaping, MUST conform to the Host ABNF described in Section 5.4 of [\[I-D.ietf-httpbis-p1-messaging\]](#).

5.4. Forwarded proto

The "proto" parameter has the value of the used protocol type. The syntax of a "proto" value, after potential quoted-string unescaping, MUST conform to the URI scheme name as defined in [Section 3.1 in \[RFC3986\]](#) and registered to IANA according to [\[RFC4395\]](#). Typical values are "http" or "https".

For example, in an environment where a reverse proxy is also used as a crypto offloader, this allows the origin server to rewrite URLs in a document to match the type of connection as the user agent requested, even though all connections to the origin server are unencrypted HTTP.

5.5. Extensions

Private extensions allow for adding own parameters and values. This may be particularly useful in a reverse proxy environment. If these extensions are to be widely spread it is RECOMMENDED that they are standardized. IANA MUST in such cases be notified, and parameters should be registered according to [\[RFC3864\]](#). Parameters MUST be registered with the "Forwarded." prefix.

New parameters and their values MUST conform the forwarded-pair as defined in ABNF in [Section 4](#).

6. Node identifiers

The node identifiers are the IP address, and optionally port number, of the network node, a predefined token hiding the real identity, but signaling that such a component exists in the network path, or a generated token allowing for tracing and debugging without revealing network internals.

nodename = IPv4address / IPv6address /
 "unknown" / obfnode

All of the identifiers may optionally have the port identifier, for example, allowing the identification of the end point in a NATted environment.

node = nodename [":" node-port]

The node-port can be identified either by its TCP port number or by a generated token obfuscating the real port number.


```
node-port      = port / obfport
port           = 1*5DIGIT
obfport        = 1*(ALPHA / DIGIT)
```

Note that this also allows port numbers to be appended to the the "unknown" identifier. Interpretation of such notation is, however, left to the possessor of a proxy adding such a value to the header field. To distinguish an obfport from a port, we RECOMMEND that an obfport always should contain at least one ALPHA.

Example:

```
192.0.2.43:47011
[2001:db8:cafe::17]:47011
```

6.1. IPv4 and IPv6 identifiers

The ABNF rules for "IPv6address" and "IPv4address" are defined in [\[RFC3986\]](#) The IPv6address SHOULD comply with textual representation recommendations [\[RFC5952\]](#) (e.g., lowercase, zero compression).

Note that the IP address may be one from the internal nets, as defined in [\[RFC1918\]](#) and [\[RFC4193\]](#). Also, note that an IPv6 adress always must be enclosed by square brackets.

6.2. The "unknown" identifier

The "unknown" identifier is used when the identity of the preceding entity is not known. One example would be a proxy server process generating an outgoing request without direct access to the incoming request TCP socket.

6.3. Obfuscated identifier

A generated identifier may be used where there is a wish to keep the internal IP addresses secret, while still allowing the Forwarded header field to be used for tracing and debugging. The identifiers can be randomly generated for each request and do not need to be statically assigned to resources. To distinguish the obfuscated identifier from other identifiers, it MUST have a leading underscore "_". Further, it MUST also consist of only US-ASCII letters and US-ASCII digits.

```
obfnode        = "_" 1*( ALPHA / DIGIT )
```

Example:

```
Forwarded: for=_hidden, for=_SEVKISEK
```


7. Implementation considerations

Note that an HTTP list allows white spaces to occur between the identifiers, and the list may be split over multiple header fields. As an example, the header field

```
Forwarded: for=192.0.2.43,for=[2001:db8:cafe::17],for=unknown
```

is equivalent to the header field

```
Forwarded: for=192.0.2.43, for=[2001:db8:cafe::17], for=unknown
```

which is equivalent to the header fields

```
Forwarded: for=192.0.2.43
Forwarded: for=[2001:db8:cafe::17], for=unknown
```

Note that this header field is relevant on a per request basis and MUST NOT be cached. Also note that this header field MUST NOT be preserved across redirects.

7.1. Example

A request from a client with IP address 192.0.2.43 passes through a proxy with IP address 198.51.100.17, then through another proxy with IP address 203.0.113.60 before reaching a origin server. This could, for example, be an office client behind a corporate malware filter talking to a origin server through a reverse proxy.

- o The HTTP request between the client and the first proxy has no Forwarded header field.
- o The HTTP request between the first and second proxy has a "Forwarded: for=192.0.2.43" header field.
- o The HTTP request between the second proxy and the origin server has a "Forwarded: for=192.0.2.43, for=198.51.100.17;by=203.0.113.60;proto=http;host=example.com" header field.

Note that, at some points in a connection chain, the information might not be correctly updated in the Forwarded header field, either because of lack of support of this HTTP extension or because of a policy decision not to disclose information about this network component.

8. Security considerations

8.1. Header validity and integrity

The Forwarded HTTP header field cannot be relied upon to be correct, as it may be modified, whether mistakenly or for malicious reasons, by every node on the way to the server, including the client making the request.

One approach is to verify the correctness of proxies and whitelist them as trusted. This approach has at least two weaknesses. First, the chain of IP addresses listed before the request came to the proxy cannot be trusted. Second, unless the communication between proxies and the end point is secured, the data can be modified by an attacker with access to the network.

8.2. Information leak

The Forwarded HTTP header field can reveal internal structures of the network setup behind the NAT or proxy setup, which may be undesired. This can be addressed either by preventing the internal nodes from updating the HTTP header field or by having an egress proxy removing entries that reveals internal network information.

9. IANA considerations

This document specifies the HTTP header listed below, which should be added to the permanent HTTP header registry defined in [[RFC3864](#)]. We also suggest that substructures are allowed for individual standards of extension parameters [Section 5.5](#).

Header field: Forwarded

Header field: Forwarded.for

Header field: Forwarded.by

Header field: Forwarded.host

Header field: Forwarded.proto

Applicable protocol: http/https

Status: standard

Author/Change controller:

IETF (iesg@ietf.org)

Internet Engineering Task Force

Specification document(s): this specification ([Section 4](#))

Related information: None

10. References

10.1. Normative references

- [I-D.ietf-httpbis-p1-messaging]
Fielding, R., Lafon, Y., and J. Reschke, "HTTP/1.1, part 1: URIs, Connections, and Message Parsing", [draft-ietf-httpbis-p1-messaging-19](#) (work in progress), March 2012.
- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), February 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](#), [RFC 3864](#), September 2004.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.
- [RFC4395] Hansen, T., Hardie, T., and L. Masinter, "Guidelines and Registration Procedures for New URI Schemes", [BCP 35](#), [RFC 4395](#), February 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", [RFC 5952](#), August 2010.

10.2. Informative references

- [RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", [RFC 6269](#), June 2011.

Appendix A. Change Log (to be removed by RFC Editor before publication)

A.1. Since [draft-petersson-forwarded-for-00](#)

Added IANA considerations.

Expanded scope and add parameterized list.

A.2. Since [draft-petersson-forwarded-for-01](#)

Removed "x-" from private extensions.

Allow for any protocol name.

Rename kv-v to forwarded-element and kv to forwarded-value.

Add informative reference [RFC6269](#).

A.3. Since [draft-petersson-forwarded-for-02](#)

Name change to [draft-ietf-appsawg-http-forwarded-00](#).

Updated proto in list under [section 5](#) Parameters.

Remove "hidden" but mention `_hidden` as an example in 6.3 Obfuscated identifier.

Clarify that IPv6-addresses must be enclosed by square brackets.

Restrict ext-value: do not allow `,` or `;`.

A.4. Since [draft-ietf-appsawg-http-forwarded-00](#)

Write IP address instead of IP number.

Remove BNF for IP addresses.

A.5. Since [draft-ietf-appsawg-http-forwarded-01](#)

Refer to httpbis instead of [RFC2616](#). Thereby also change to [RFC5234](#) ABNF.

Split up ABNF to be more general on top level.

Add some comments on [draft-ietf-httpbis-p2-semantics-19](#)#section-3.1 to "Implementation Considerations"

Removal of ABNF appendix.

Merging of the sections "Private extensions" and "Future extensions".

Authors' Addresses

Andreas Petersson
Opera Software
S:t Larsgatan 12
Linköping SE-582 24

Email: pettson@opera.com

Martin Nilsson
Opera Software
S:t Larsgatan 12
Linköping SE-582 24

