

HTTP
Internet-Draft
Intended status: Standards Track
Expires: May 20, 2020

P. O'Hanlon
J. Guessing
British Broadcasting Corporation
November 17, 2019

The Transport-Info Response Header
draft-ohanlon-transport-info-header-00

Abstract

The Transport-Info header provides a mechanism to inform a client of the last-mile server's view of the network transport related information such as current delivery rate and round-trip time. This information has a wide range of uses such as client monitoring and diagnostics, or allowing a client to adapt to current network conditions.

Note to Readers

RFC Editor: please remove this section before publication

Source code and issues for this draft can be found at
<https://github.com/bbc/draft-ohanlon-transport-info-header> [1].

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 May 20, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

Transport-Info Header

November 2019

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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
1.1.	Motivation	3
1.2.	Notational Conventions	3
2.	The Transport-Info Response Header	3
2.1.	Utilisation of Transport-Info header metrics	5
3.	Server side behaviour	6
4.	Client side proxy considerations	6
5.	IANA Considerations	6
6.	Security Considerations	7
7.	References	7
7.1.	Normative References	7
7.2.	Informative References	8
7.3.	URIs	9
	Authors' Addresses	9

[1.](#) Introduction

The Transport-Info header provides for relaying of transport protocol related information from a last-mile edge server entity to a client with the aim of informing the client of the server's view on the transport state. The state of a connection is dependent upon information based upon packet exchanges during the transport processes. Firstly, there is information that is common to both client and server, such as the calculated round-trip time (RTT), although it may be measured using different packets at each end. Secondly, there is state information that exists only at each endpoint, such as the size of the congestion, and receive windows. Thus certain transport state information is only available at the server which can be useful to the client, for example, to calculate the current transport rate. This information may then be used to better inform a client of the state of the network path and make

appropriate adaptations.

The information can also be utilised by a client to provide for application level client oriented metric logging to back-end systems

for monitoring and analysis purposes. Such data could be utilised in a manner not unlike that proposed in [\[RFC4898\]](#).

This approach is directly applicable to TCP but also can be utilised with other related transport protocols, such as QUIC [\[I-D.ietf-quic-transport\]](#).

[1.1.](#) Motivation

This work is motivated in part by the fact that even in modern web browsers web applications are not currently able to obtain such low level information about their connections. Additionally some information is only available at the server, such as the size of the server congestion window. As a result clients often resort to application level measurements, to infer such things as the current delivery rate, but these are not always indicative of the performance of the transport layer.

There exist W3C specifications such as the Network Information API [\[network-info-api\]](#), which contains an attribute named "downlink" that purports to provide measurement of effective bandwidth "across recently active connections". However, in practice the downlink measurement appears to be a very rough estimate which is of little use for informing an application of dynamic network conditions. Furthermore, it currently has limited browser support.

[1.2.](#) Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [\[RFC2119\]](#) [\[RFC8174\]](#) when, and only when, they appear in all capitals, as shown here.

This document uses the Augmented Backus-Naur Form (ABNF) notation of [\[RFC5234\]](#) with the list rule extension defined in [\[RFC7230\]](#),

[Appendix B](#). It includes by reference the DIGIT rule from [[RFC5234](#)] and the OWS and field-name rules from [[RFC7230](#)].

2. The Transport-Info Response Header

The Transport-Info header uses the proposed Structured Header draft [[I-D.ietf-httpbis-header-structure](#)]

Transport-Info = sh-list

Each member of the parameterised list represents an entry that contains a set of metrics reported by the edge server.

The list members identify the server that inserted the value, and MUST have a type of either sh-string or sh-token. Depending on the deployment, this might be a product or service name (e.g., ExampleEdge or "Example CDN"), a hostname ("edge-1.example.com"), and IP address, or a generated string.

Each member of the list can also have a number of parameters that contain metrics. While all but one of these parameters are OPTIONAL, edge servers are encouraged to provide as much information as possible.

- o Exactly one parameter whose name is "ts", and whose value is an sh-float indicating the measurement timestamp in seconds since UNIX epoch.
- o Optionally one parameter whose name is "alpn", and whose value is an sh-string representing the ALPN protocol identifier [[alpn-ids](#)].
- o Optionally one parameter whose name is "cc_algo", and whose value is sh-string, conveying the name of congestion control algorithm used by the server for this connection.
- o Optionally one parameter whose name is "cwnd", and whose value is a sh-integer, conveying the size of the server's congestion window [[RFC5681](#)] in packets.
- o Optionally one parameter whose name is "rcv_space", and whose value is a sh-integer, conveying the size of the receiver's window in bytes.
- o Optionally one parameter whose name is "dstport", and whose value is a sh-integer, conveying the server's destination port of this connection for correlation of measurements between requests.
- o Optionally one parameter whose name is "mss", and whose value is a

- sh-integer, conveying the size of the server's Maximum Segment Size in bytes.
- o Optionally one parameter whose name is "rtt", and whose value is an sh-float, in milliseconds, indicating the server's estimate of the Round-Trip Time from its transport layer.
 - o Optionally one parameter whose name is "rttvar", and whose value is an sh-float, in milliseconds, indicating the server's estimate of the variation of the Round-Trip Time [[RFC6298](#)] from its transport layer.
 - o Optionally one parameter whose name is "send_rate", and whose value is a sh-float, in kilobits per second, conveying the server's calculation of the sending rate for this connection.

Here is an example of a header with a single set of metrics:

```
Transport-Info = ExampleEdge; ts=1567176968.69; alpn="h2"; cwnd=24;
                 rtt=250; mss=1460; rttvar=10; dstport=12345
```

Whilst it is understood that such metrics may only provide an instantaneous view on the transport state, the Transport-Info header is designed to allow for delivery of multiple timestamped entries in a single header.

Here is an example of header with multiple entries, utilising the structured header inner-list type:

```
Transport-Info = "edge-1.example.com"; ts=1567176968.69; alpn="h2"; cwnd=24;
                 rtt=250; mss=1452; rttvar=10; dstport=123451,
                 "edge-1.example.com"; ts=1567176969.97; alpn="h2"; cwnd=23;
                 rtt=255; mss=1452; rttvar=12; dstport=123451
```

If the end points support HTTP/2, and later, another technique to increase temporal coverage for an ongoing session is for the client to issue additional HEAD or OPTION * requests for a resource at the same origin. This works with HTTP/2 and later as all requests to the same origin utilise one TCP or QUIC connection. Whilst the HTTP priorities can affect the allocation of capacity between streams, one use-case for the Transport-Info header is for information regarding sustained flows, such as media delivery, which tend consist of a known limited number of flows to the same origin so the priorities

would not affect the calculations.

2.1. Utilisation of Transport-Info header metrics

In the case of TCP, calculation of the transport transmission rate is possible using the cwnd and rtt, and knowledge of the mss. The equation being as follows:

$$\text{send_rate} = 8 * \text{send_window} / \text{rtt}$$

$$\text{Where send_window} = \min(\text{cwnd} * \text{mss}, \text{rcv_space})$$

If the mss is not available then it is possible to perform the calculation using an estimate of the mss, or a common value such as 1460 for IPv4. It understood there can be some variation for different network and tunnelled paths (e.g. 1452 for IPv4 PPPoE) as can be seen in recent studies [[exploring-mtu](#)], although the large proportion of mss values fall within a range 1220-1460. The send_window is preferably calculated using a minimum of the cwnd and rcv_space, but if the rcv_space is not available it may be approximated by just using the cwnd.

This equation maybe applied for other related window based transport protocols (e.g. QUIC [[I-D.ietf-quick-transport](#)]) with similar information, although it may need some modification.

The calculation of the send rate maybe performed at the server, or may be left to the client to calculate as and when required.

3. Server side behaviour

With most web server deployments an origin server sits behind some form of CDN, with varying levels of fan-out to a point where an edge server is connected on the last mile to clients. The Transport-Info header SHOULD only be inserted into an HTTP stream by the last hop edge server that is connected to clients so that it conveys information pertinent to the client's direct transport path. Also the Transport-Info MUST not be cached.

RFC Editor: please remove this section before publication

The provision of the Transport-Info header is possible using a number of existing server systems that already provide support for such metrics, which currently utilise operating system support for tcp_info data structure which are available on Linux and BSD systems.

In terms of current implementations there is in-built support in Nginx/Openresty using its variables "var.tcpinfo_rtt" etc. Apache Traffic Server provides support using the TCPInfo plugin. Varnish provides access to "tcp_info" using their "vmod_tcp" module. Node.js has libraries such as "nodejs_tcpinfo" which provide support. Whilst most of the implementations do not provide access to the TCP MSS it is available via the underlying kernel tcp_info data structure so it would be fairly straightforward to provide access to such information.

[4.](#) Client side proxy considerations

In the case where a proxy services client requests, this proxy would be configured according to local policy as to whether it passes through, modifies or drops the Transport-Info header. This decision can depend on a number of factors, including the utility of the header given local network configuration, and also whether the header might reveal unwanted information to end clients, since the Transport-Info header would relate to the connection between the edge CDN node and the proxy.

[5.](#) IANA Considerations

This specification registers the following entry in the Permanent Message Header Field Names registry established by [\[RFC3864\]](#):

- o Header field name: Transport-Info

- o Applicable protocol: http

- o Status: standard

- o Author/Change Controller: IETF

- o Specification document(s): [this document]

o Related information:

6. Security Considerations

The content of the Transport-Info is largely available through other techniques such as packet capture so it should not lead to security issues. Certain metrics, such as the cwnd, may be considered as less visible but since they are part of the transport layer they can be inferred. Any metrics that may be considered private should not be sent in the header, or sent only over an encrypted connection.

In the case where clients are connected via a proxy then organisations may wish to modify or drop the header if they consider that it might reveal unwanted information to end clients.

If the header is delivered over a transport protocol whose content can be modified without detection then parties should be aware that the header could be maliciously modified to alter the metrics values which could result in the client making incorrect adaptations.

7. References

7.1. Normative References

- [I-D.ietf-httpbis-header-structure]
Nottingham, M. and P. Kamp, "Structured Headers for HTTP", [draft-ietf-httpbis-header-structure-14](#) (work in progress), October 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](#), [RFC 3864](#), DOI 10.17487/RFC3864, September 2004, <<https://www.rfc-editor.org/info/rfc3864>>.

- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion

Control", [RFC 5681](#), DOI 10.17487/RFC5681, September 2009, <<https://www.rfc-editor.org/info/rfc5681>>.

[RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", [RFC 6298](#), DOI 10.17487/RFC6298, June 2011, <<https://www.rfc-editor.org/info/rfc6298>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[7.2](#). Informative References

[alpn-ids]

"Application-Layer Protocol Negotiation (ALPN) Protocol ID", IANA , n.d., <<http://www.iana.org/assignments/tls-extensiontype-values>>.

[exploring-mtu]

Custura, A., Fairhurst, G., and I. Learmonth, "Exploring usable Path MTU in the Internet", Network Traffic Measurement and Analysis Conference , April 2018, <<https://doi.org/10.23919/TMA.2018.8506538>>.

[I-D.ietf-quic-transport]

Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", [draft-ietf-quic-transport-23](#) (work in progress), September 2019.

[network-info-api]

Grigorik, I., "Network Information API", W3C , September 2019, <<http://wicg.github.io/netinfo/>>.

[RFC4898] Mathis, M., Heffner, J., and R. Raghunarayan, "TCP Extended Statistics MIB", [RFC 4898](#), DOI 10.17487/RFC4898, May 2007, <<https://www.rfc-editor.org/info/rfc4898>>.

[RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.

[RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", [RFC 7230](#), DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.

[7.3.](#) URIs

[1] <https://github.com/bbc/draft-ohanlon-transport-info-header>

Authors' Addresses

Piers O'Hanlon
British Broadcasting Corporation

Email: piers.ohanlon@bbc.co.uk

James Gruessing
British Broadcasting Corporation

Email: james.gruessing@bbc.co.uk

