Workgroup: tcpm
Internet-Draft:
draft-boucadair-tcpm-rst-diagnostic-payload-00
Published: 30 March 2022
Intended Status: Standards Track
Expires: 1 October 2022
Authors: M. Boucadair
Orange

TCP RST Diagnostic Paylaod

Abstract

This document specifies a diagnostic payload format to be returned in TCP RST segments. Such payloads are used to share with the endpoints the reasons for which a TCP connection has been reset. This is meant to ease diagnostic and troubleshooting.

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 1 October 2022.

Table of Contents

1. Introduction

A TCP connection [I-D.ietf-tcpm-rfc793bis] can be reset by a peer for various reasons, e.g., a received data does not correspond to an active connection. Also, a TCP connection can be reset by an on-path service function (e.g., CGN [RFC6888], NAT64 [RFC6146], firewall) for various reasons. Typically, a NAT function can generate an RST segment to notify the peers upon the expiry of the lifetime of the corresponding mapping entry or because an RST segment was received from a peer (Section 2.2 of [RFC7857]). A TCP connection can also be closed by a user or an application at any time. However, the peer that receives an RST segment does not have any hint about the reason that led to terminating the connection. Likewise, the application that relies upon such a TCP connection may not easily identify the reason for a connection closure. Troubleshooting such events at the terminal side that receives the RST segment may not be trivial.

This document fills this void by specifying a diagnostic payload that is returned in an RST segment. Returning such data is consistent with the provision in Section 3.5.3 of [I-D.ietf-tcpm-rfc793bis] for RST segments.

This document does not change the conditions under which an RST segment is generated (Section 3.5.2 of [I-D.ietf-tcpm-rfc793bis]).

The generic procedure for processing an RST segment is specified in Section 3.5.3 of [I-D.ietf-tcpm-rfc793bis]. Only the deviation from that procedure to identify and validate an enclosed diagnostic payload is provided in Section 3.

A peer that receives a diagnostic payload may pass that information to the local application in addition to the information (MUST-12) described in Section 3.6 of [I-D.ietf-tcpm-rfc793bis]. That information may also be logged locally, unless a local policy specifies otherwise.

The first version of the specification is meant to discuss the format and the overall approach to ease maintaining the list of codes while allowing for adding new codes as needed in the future. As such, this first version of the specification does not include a comprehensive list of error codes. These codes will be completed in future versions (Table 1).

2. Terminology

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 makes use of the terms defined in Section 4 of [I-D.ietf-tcpm-rfc793bis].

3. RST Diagnostic Payload

In order to unambiguously identify an RST diagnostic payload that is compliant with the present specification, the payload MUST use the I-JSON message format [RFC7493]. The following parameters are defined:

rc: Stands for "Reason Code". This parameter takes a value from the "TCP Failure Causes" registry (Section 4.1). This parameter is omitted if none of the values maintained by IANA can be used to report a reset failure cause.

rd: Stands for "Reason Description". It includes a brief description of the reason code. This parameter SHOULD NOT be included if a code that covers this error case is already registered in Section 4.1. This parameter is useful only for codes that are not yet registered or application-specific codes.

At least one of the above parameters MUST be included in an RST diagnostic payload that is compliant with the present specification.

Figure 1 depicts an example of an RST diagnostic payload that is generated to inform the peer that the connection is reset because an ACK was received while the connection is still in the LISTEN state.

```
{
 "rc": 2
}
```

Figure 1: An RST Diagnostic Payload with Reason Code

Figure 2 shows an example of an RST diagnostic payload that includes a free description to report a case that is not covered yet by the table maintained by IANA (Section 4.1).

```
{
 "rd": "brief human-readable description"
}
```

Figure 2: An RST Diagnostic Payload with Reason Description

An RST diagnostic payload may be included by the peer that resets the connection or by an on-path service function. For example, the following payload can be returned by a NAT when a mapping entry expires (Figure 3).

```
{
 "rc": 8
}
```

Figure 3: An RST Diagnostic Payload to Report Connection Timeout

4. IANA Considerations

4.1. New Registry for TCP Failure Causes

This document requests IANA to create a new subregistry entitled "TCP Failure Causes" under the "Transmission Control Protocol (TCP) Parameters" registry [IANA-TCP].

The registry is initially populated with the following values:

Value	Description	Specification (if available)
1	Data lost. New data is received after CLOSE is called	Sections 3.6.1 and 3.10.7.1 of [I-D.ietf-tcpm-rfc793bis]
2	Still in LISTEN. Received ACK while the connection still in the LISTEN state	Section 3.10.7.2 of [I-D.ietf-tcpm-rfc793bis]
3	Malformed Message	N/A
4	Not Authorized	N/A
5	Resource Exceeded	N/A
6	Network Failure	N/A
7	Connection Reset received from the peer	N/A
8	Destination Unreachable	N/A

Value	Description	Specification (if available)
9	Connection Timeout	RFCXXX
10	description XXX	URL

Table 1: Initial TCP Failure Causes

The assignment policy for this registry is "Expert Review" (Section 4.5 of [RFC8126]). The designated experts may approve registration once they checked that the new requested code is not covered by an existing code and if the provided reasoning to register the new code is acceptable. A registration request may supply a pointer to a specification where that code is defined. However, a registration may be accepted even if no permanent and readily available public specification is available.

5. Security Considerations

[I-D.ietf-tcpm-rfc793bis] discusses TCP-related security considerations. RST-specific attacks and their mitigation are discussed in Section 3.10.7.3 of [I-D.ietf-tcpm-rfc793bis].

In addition to these considerations, it is RECOMMENDED to control the size of acceptable diagnostic payload and keep it as brief as possible. Also, it is RECOMMENDED to avoid leaking privacy-related information as part of the diagnostic payload (e.g., including a description such as "user X resets explicitly the connection").

6. Acknowledgements

TBC.

7. References

7.1. Normative References

[I-D.ietf-tcpm-rfc793bis] Eddy, W. M., "Transmission Control Protocol (TCP) Specification", Work in Progress, Internet-Draft, draft-ietf-tcpm-rfc793bis-28, 7 March 2022, <https://www.ietf.org/archive/id/draft-ietf-tcpm-rfc793bis-28.txt>.

[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>.

[RFC7493] Bray, T., Ed., "The I-JSON Message Format", RFC 7493, DOI 10.17487/RFC7493, March 2015, <https://www.rfc-editor.org/info/rfc7493>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[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

[IANA-TCP] IANA YANG, "Transmission Control Protocol (TCP) Parameters", <https://www.iana.org/assignments/tcp-parameters/tcp- parameters.xhtml>.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <https://www.rfc-editor.org/info/rfc6146>.

[RFC6888] Perreault, S., Ed., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, DOI 10.17487/RFC6888, April 2013, <https://www.rfc-editor.org/info/rfc6888>.

[RFC7857] Penno, R., Perreault, S., Boucadair, M., Ed., Sivakumar, S., and K. Naito, "Updates to Network Address Translation (NAT) Behavioral Requirements", BCP 127, RFC 7857, DOI 10.17487/RFC7857, April 2016, <https://www.rfc-editor.org/info/rfc7857>.

Author's Address

Mohamed Boucadair
Orange
35000 Rennes
France

Email: mohamed.boucadair@orange.com