TRAM
Internet-Draft
Intended status: Standards Track
Expires: January 5, 2015

M. Thomson
Mozilla
B. Aboba
Microsoft
A. Johnston
Avaya
O. Moskalenko
public project

rfc5766-turn-server
July 4, 2014

A Bandwidth Attribute for TURN
draft-thomson-tram-turn-bandwidth-01

Abstract

An attribute is defined for Session Traversal Utilities for NAT (STUN) that allows for declarations of bandwidth limits on the negotiated flow. The application of this attribute is the negotiation of bandwidth between a Traversal Using Relays around NAT (TURN) client and a TURN server.

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 January 5, 2015.

Table of Contents

1. Introduction

This document defines a BANDWIDTH attribute that can be used to request and allocate bandwidth at a Traversal Using Relays around NAT (TURN) relay [RFC5766].

The operator of a TURN server will likely wish to provide fairness between relayed sessions. A TURN server might also wish to limit the use of service to audio-only sessions, or low bandwidth video and audio sessions. In addition, the server may apply rate-limiting policy depending on the credential used for authentication, or the origin of the client. Without the BANDWIDTH attribute, there is no way for a client to indicate the expected bandwidth utilization, or for the server to indicate the maximum bandwidth utilization allowed before rate limiting could be applied.

This attribute is used for indicating a bandwidth limit that is set in policy. The sender is not advised or required to utilize bandwidth up to this limit; limits are usually set well in excess of application needs. Senders also limit their use of bandwidth in reaction to path congestion and "circuit breakers".

Note that the BANDWIDTH attribute was originally in the TURN draft up to version draft-ietf-behave-turn-07 where it was removed as "the requirements for this feature were not clear and it was felt the feature could be easily added later." This draft proposes adding this attribute back into TURN. A related error code 507 "Insufficient Bandwidth Capacity" was also defined in the TURN Internet-Draft, but is not proposed in this draft. This attribute has also been proposed to be used by ICE to provide communication consent [I-D.thomson-mmusic-rtcweb-bw-consent]. No use cases have been identified where bandwidth information is useful for a STUN server which is responding to STUN binding requests.

There have been discussions about what other media-related information could be usefully exchanged between a TURN client and a TURN server. One proposal was for the actual media type (voice, video, data) to be exchanged. Other proposals include more granularity over the bandwidth, including max, min, average, etc. While these could be added, the authors do not feel the use cases for these data have been sufficiently developed yet. Also, this information is known in signaling through the SDP attributes and parameters. In a particular implementation, it could be possible for a signaling-aware entity to share this information with a TURN server in order to apply policy for the media relay.

2. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, RFC 2119 [RFC2119] and indicate requirement levels for compliant implementations.

The terms client, server, and peer are those used for TURN, as defined in [RFC5766].

3. The BANDWIDTH Attribute

The BANDWIDTH attribute (identifier TBD) identifies the rate of packet transmission in kilobits per second that is permitted for a given transport flow. The BANDWIDTH attribute is a comprehension-optional attribute (see Section 15 from [RFC5389]). Figure 1 shows the format of this attribute.

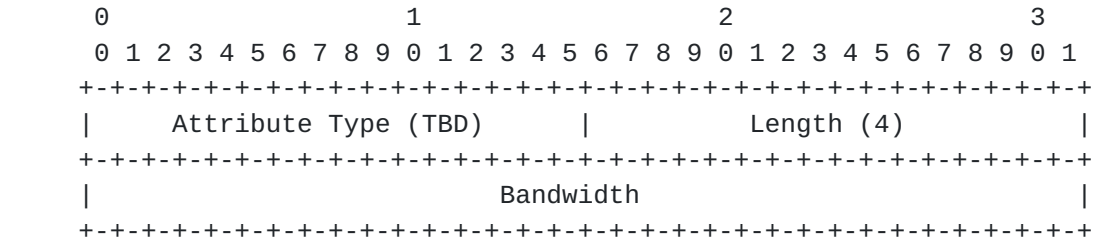

Figure 1: Bandwidth Attribute Format

The value of this attribute is an unsigned integer that represents the maximum bandwidth for the flow in kilobits per second (1 kilobit = 1024 bits). This is the original format of the Bandwidth attribute. This format could include a maximum and average bandwidth, as the BANDWIDTH-USAGE attribute proposed in [I-D.martinsen-tram-discuss].

4. Applications

This section discusses the application of the BANDWIDTH attribute for STUN, TURN, and ICE.

4.1. STUN Usage

Since the bandwidth of a communications session has no bearing on a STUN server that simply responds to binding requests, this attribute MUST NOT be used for client-server STUN requests or responses.

4.2. TURN Usage

This attribute can be useful for communication between a TURN client and a TURN server.

The BANDWIDTH attribute indicates a limit to available bandwidth for TURN [RFC5766] allocation. The bandwidth limit is symmetric; the value covers the bandwidth of data sent from a peer toward the TURN server and the bandwidth of data sent from client to the TURN server.

A BANDWIDTH attribute MAY be present in an Allocate request. This attribute indicates that the given bandwidth is requested. A BANDWIDTH attribute MAY be present in an Allocate response. This attribute in a response indicates the limit that will be applied by the TURN server. The value a TURN server provides could be influenced by the value that a TURN client requests at the discretion of server policy. A client could use this bandwidth limitation of the TURN server in choosing media types or in choosing codecs for a media session.

4.3. ICE Usage

While [I-D.thomson-mmusic-rtcweb-bw-consent] proposed the use of the BANDWIDTH attribute to provide bandwidth consent for ICE, this draft does not do so. This attribute MUST NOT be used with ICE.

5. Bandwidth Measurement Considerations

Allocation messages (Binding and Allocate) sent to and from the TURN server are exempt from any bandwidth measurement accounting.

In calculating bandwidth, the entire IP packet - including the header - is measured. This is identical to the measurement performed by the Real-time Transport Protocol (RTP) [RFC3550]. At a TURN server, bandwidth measurement is performed on the packets arriving at or leaving from the TURN server, prior to the encapsulation that occurs between TURN server and TURN client.

Determining the rate requires that the bits be allocated to specific intervals of time. How bits are allocated MAY vary between implementations.

Measurement of bandwidth is imperfect and inconsistent. Packet jitter can result in fluctuations in received packet rate so that a receiver might see an instantaneous bandwidth that is different to what the sender might have transmitted. Jitter can cause the observed bandwidth of incoming packets to temporarily increase above

the permitted rate. At a minimum, implementations SHOULD allow for short periods of excessive bandwidth to allow for these temporary increases.

5.1. Rate Enforcement

Enforcement of limits by the TURN server SHOULD provide an allowance for application usages that temporarily exceed the limit. For example, assessing observed bandwidth usage as an average over 10 seconds ensures that real-time video does not clip unnecessarily; shorter durations could result in the enforcement affecting valuable intra-frames.

6. Security Considerations

For STUN requests or responses that are not sent using TLS or DTLS transport, the bandwidth information contained in the BANDWIDTH attribute will be available to an eavesdropper who could use it to learn about the nature of a session to be established. For example, they might be able to deduce from the bandwidth requested that the session is likely to be audio only, or audio and video. However, an on-path attacker can likely learn this same information from either the signaling channel or by inspecting the RTP packet headers, which are in the clear for SRTP, or simply by measuring the media bandwidth used.

If a STUN request or response is transported using TCP or UDP, the BANDWIDTH attribute will have integrity protection from the MESSAGE-INTEGRITY attribute if the request is authenticated using the STUN short-term or long-term authentication method. Unauthenticated TCP or UDP requests will not have integrity protection and could be modified by a MitM attacker. The use of DTLS transport [I-D.ietf-tram-stun-dtls] provides integrity protection for the BANDWIDTH attribute regardless of the STUN authentication method used.

7. IANA Considerations

The STUN BANDWIDTH attribute uses the TBD value in the comprehension-optional range. This attribute is registered in the "STUN Attribute" Registry following the procedures of Section 18.2 of [RFC5389].

8. Implementation Status

Note to RFC Editor: Please remove this entire section prior to

publication, including the reference to RFC 6982.

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC6982]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC6982], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

A multiple realms capable advanced open source TURN server (named 'Coturn') has been created by Oleg Moskalenko and is freely licensed under the New BSD license. This reference implementation and proof-of-concept provides a clone (a spin-off) of the rfc5766-turn-server project adding STUN BANDWIDTH attribute support, among other TRAM Working Group STUN and TURN extensions.

'Coturn' is backward-compatible with rfc5766-turn-server project but the code is more complex and it uses a different (also more complex) database structure. It is the intent to add all IETF TRAM TURN server related capabilities to this project as they mature. 'Coturn' is publicly available and can be found at:
https://code.google.com/p/coturn/

9. References

9.1. Normative References

[I-D.ietf-tram-stun-dtls]
Petit-Huguenin, M. and G. Salgueiro, "Datagram Transport Layer Security (DTLS) as Transport for Session Traversal Utilities for NAT (STUN)", draft-ietf-tram-stun-dtls-05 (work in progress), June 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.

[RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.

[RFC5766] Mahy, R., Matthews, P., and J. Rosenberg, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)", RFC 5766, April 2010.

9.2. Informative References

[I-D.martinsen-tram-discuss]
Martinsen, P. and H. Wildfeuer, "Differentiated prIorities and Status Code-points Using Stun Signalling (DISCUSS)", draft-martinsen-tram-discuss-00 (work in progress), February 2014.

[I-D.thomson-mmusic-rtcweb-bw-consent]
Thomson, M. and B. Aboba, "Bandwidth Constraints for Session Traversal Utilities for NAT (STUN)", draft-thomson-mmusic-rtcweb-bw-consent-00 (work in progress), October 2012.

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.

[RFC6982] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", RFC 6982, July 2013.

Authors' Addresses

Martin Thomson
Mozilla
331 E Evelyn Street
Mountain View, CA 94041
USA

Phone: +1 650-353-1925
Email: martin.thomson@gmail.com

Bernard Aboba
Microsoft
One Microsoft Way
Redmond, WA 98052
USA

Email: bernard_aboba@outlook.com

Alan Johnston
Avaya
St. Louis, MO
USA

Email: alan.b.johnston@gmail.com

Oleg Moskalenko
public project rfc5766-turn-server
Walnut Creek, CA
USA

Email: mom040267@gmail.com
URI: https://code.google.com/p/rfc5766-turn-server/