

rtcweb  
Internet-Draft  
Intended status: Standards Track  
Expires: April 15, 2013

M. Thomson  
B. Aboba  
Microsoft  
October 12, 2012

Bandwidth Constraints for Session Traversal Utilities for NAT (STUN)  
draft-thomson-mmusic-rtcweb-bw-consent-00

## 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 to reducing denial of service attacks from internet telephony systems. Other applications include negotiation of bandwidth at packet relays.

## 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 April 15, 2013.

## 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.

Internet-Draft

STUN Bandwidth

October 2012

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Terminology . . . . .	<a href="#">4</a>
<a href="#">3.</a>	The BANDWIDTH Attribute . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Application . . . . .	<a href="#">4</a>
<a href="#">4.1.</a>	ICE Bandwidth Consent . . . . .	<a href="#">5</a>
<a href="#">4.1.1.</a>	STUN Usage . . . . .	<a href="#">5</a>
<a href="#">4.1.2.</a>	Bandwidth Limiting . . . . .	<a href="#">5</a>
<a href="#">4.1.3.</a>	Bandwidth Consent Revocation . . . . .	<a href="#">6</a>
<a href="#">4.2.</a>	Relay Bandwidth Allocation . . . . .	<a href="#">6</a>
<a href="#">5.</a>	Bandwidth Measurement Considerations . . . . .	<a href="#">6</a>
<a href="#">5.1.</a>	Rate Enforcement . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">8</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">8</a>
<a href="#">8.</a>	Acknowledgments . . . . .	<a href="#">8</a>
<a href="#">9.</a>	References . . . . .	<a href="#">8</a>
<a href="#">9.1.</a>	Normative References . . . . .	<a href="#">8</a>
<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">9</a>
	Authors' Addresses . . . . .	<a href="#">9</a>

## 1. Introduction

A key security property that Interactivity Connectivity Establishment (ICE) [[RFC5245](#)] establishes is that the subject of a flow of packets consents to receive packets. This provides a measure of protection against the Voice Hammer attack (see [Section 18.5.1 of \[RFC5245\]](#)) where an attacker with access to signaling induces valid clients to send excessive amounts of data toward a victim.

ICE depends on the use of a Session Traversal Utilities for NAT (STUN) Binding request and response to provide an indication of consent. A host that provides a Binding response demonstrates that they have seen the transaction ID and are therefore either on the path between sender and receiver, or have an agent on the path.

As a result of the connectivity check, the sender determines that either the receiving host consents to receiving packets, or some on-path attacker has faked consent. In the latter case, the on-path attacker is already in a position to generate packets toward the receiving host, so this does not present an advantage to the attacker and we discount the attack.

This basic consent mechanism only establishes that some data is acceptable. It does not establish how much the receiver is prepared to accept. In particular, where media multiplexing is negotiated, consent cannot be provided for individual media; it must apply to all potential streams received on a transport address. Given the wide disparity in potential bandwidth usage by text, audio and video, this enables a volume-based denial of service attack. In this attack a service that is prepared to automatically accept real-time communications can become the target of excessive bandwidth from clients. This only requires that the client visits the attacker's site, there is no user consent required.

This attack is only mitigated by measures such as continuing consent [[I-D.muthu-behave-consent-freshness](#)]. Refusing to provide continuing

consent takes a non-trivial amount of time to effect changes in senders. Thus, continuing consent can only limit the duration of an attack, not prevent it, and only by refusing all media traffic.

This document defines a BANDWIDTH attribute for STUN that can be used to prevent this attack. This attribute can also be used to request and allocate bandwidth at a Traversal Using Relays around NAT (TURN) relay.

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".

## [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 term "sender" and "receiver" refer to peers that are respectively sending or receiving packets on a given transport flow. In a typical peer to peer transport flow, both peers act in both roles. ICE terminology, specifically candidate, flow, and candidate are borrowed from [[RFC5245](#)].

## [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.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								




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).

## 4. Application

Bandwidth limits are applied to both ICE connectivity checking and TURN relay allocation.

### 4.1. ICE Bandwidth Consent

For a negotiated ICE flow, the BANDWIDTH attribute indicates the bandwidth that a receiver consents to receive.

In the absence of a BANDWIDTH attribute, no bandwidth limits apply to a flow.

#### 4.1.1. STUN Usage

A sender that supports this attribute MUST include a BANDWIDTH attribute in the STUN Binding requests it generates for connectivity checks. The primary utility of this inclusion is to indicate support for a parameter.

A receiver that supports this attribute includes the permitted bandwidth in the STUN Binding response it generates. A receiver that requires bandwidth consent MUST ignore Binding requests that do not include the BANDWIDTH attribute. A receiver MAY use the value in the Binding request as guidance on what bandwidth to permit.

#### 4.1.2. Bandwidth Limiting

An ICE-capable sender MUST NOT exceed the allowed bandwidth that has been indicated in connectivity checks for the specific flow. Bandwidth limits for each flow are independent. That is, a bandwidth limit learned from a connectivity check response only applies to packets that are sent from the local candidate to the remote candidate that were used in that connectivity check. An attacker with access to signaling could otherwise insert a candidate that they control. The attacker-controlled candidate could then indicate a larger bandwidth limit that affects other flows.

Independent bandwidth limits for flows implies that where there are multiple valid candidate pairs the overall bandwidth limit is increased for each valid candidate pair. For an attacker, this presents an opportunity to multiply the bandwidth that flows toward a receiver by the number of valid candidate pairs. A receiver SHOULD monitor the incoming bandwidth for a component and limit the aggregate bandwidth to the expected maximum. This depends on application-specific knowledge of how flows are expected to be used, such as knowledge of how the ICE component relates to Session Description Protocol (SDP) [[RFC4566](#)] "m=" lines. If an aggregate limit is exceeded, the receiver SHOULD revoke consent ([Section 4.1.3](#)) on one or more flows.

In ICE, a peer MAY choose to include a BANDWIDTH value of zero prior to nominating a candidate pair. This implies that only STUN Binding

requests are permitted on the flow. The advantage of this is that the bandwidth allowance is not increased by the number of valid candidate pairs. This increases the latency of flow setup because the controlled peer needs to perform a post-nomination connectivity check to check if available bandwidth has increased to a usable value.

#### [4.1.3](#). Bandwidth Consent Revocation

During ongoing consent checks [[I-D.muthu-behave-consent-freshness](#)], the BANDWIDTH attribute allows for a faster revocation of consent by a receiver. Without the BANDWIDTH attribute, the receiver stops responding to Binding requests. The sender only stops sending when enough of these responses are lost. By responding with a BANDWIDTH value of zero, a sender ceases packet transmission on the flow in a

much shorter time.

A sender MUST cease transmission of all packets toward a receiver that indicates a bandwidth of zero. This is important even where rates are averaged over time, where changes in the bandwidth limit might otherwise take some time.

#### 4.2. Relay Bandwidth Allocation

The BANDWIDTH attribute indicates a limit to available inbound bandwidth for TURN [[RFC5766](#)] allocation. Inbound bandwidth is the bandwidth of data sent from a peer toward the TURN server.

A BANDWIDTH attribute - when present in an Allocate request - indicates that the given bandwidth is requested. A BANDWIDTH attribute in an Allocate 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 TURN client can use the indicated bandwidth to limit the value that it sends in ICE connectivity check responses.

Bandwidth that the TURN client sends toward the TURN server is not governed by this attribute. A TURN server is able to terminate an allocation if a TURN client generates excessive outbound bandwidth.

### 5. Bandwidth Measurement Considerations

Connectivity check and allocation messages (Binding and Allocate) are exempt from any bandwidth measurement accounting. This allows consent to be verified without being subject to delays due to

bandwidth throttling. Senders are expected to limit the rate of outgoing connectivity checks using independent mechanisms.

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 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

[[Editor's Note: There are two approaches to this: Either make the limit a hard limit (with a small jitter allowance), which would necessitate fairly high bandwidth limits for normal usage lest there be squeezing or dropping of video i-frames, which can significantly affect a real-time experience. The other approach, taken here, permits a more usage-aware limiting, taking the limit as more of a "guideline", which allows latitude for temporary excess, enabling more realistic limits (and probably some queueing somewhere), with guarantees on compliance with the limit over longer periods.]]

Enforcement of bandwidth limits is a sender responsibility, though a receiver or other middlebox MAY perform enforcement. Senders are able to make temporary allowances for the data that is being transmitted, by averaging bandwidth usage to account. This allows for different bandwidth usage profiles. For example, real-time audio typically uses a nearly constant rate, whereas bandwidth consumption increases significantly for the transmission of intra-frames in real-time video. In contrast, real-time application sharing has highly unpredictable bandwidth consumption.

Enforcement of limits by nodes other than the sender 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 any enforcement affecting valuable intra-frames.

## 6. Security Considerations

ICE negotiation potentially results in multiple candidate pairs for a component becoming valid. If a non-zero bandwidth value is used during the checking phase, the sender might be convinced that there are multiple valid flows. Mitigation measures and considerations for their use are described in [Section 4.1](#).

## 7. IANA Considerations

IANA has allocated a code of (TBD) to the STUN BANDWIDTH attribute. This attribute is registered in the "STUN Attribute" Registry following the procedures of [Section 18.2 of \[RFC5389\]](#). BANDWIDTH is a comprehension-optional STUN attribute.

## 8. Acknowledgments

Humayun Khan, Tim Moore provided input and implementation experience.

## 9. References

### 9.1. Normative References

- [I-D.muthu-behave-consent-freshness]  
Perumal, M., Wing, D., and H. Kaplan, "STUN Usage for Consent Freshness and Session Liveness",  
[draft-muthu-behave-consent-freshness-01](#) (work in progress), July 2012.
- [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

- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.

## Authors' Addresses

Martin Thomson  
Microsoft  
3210 Porter Drive  
Palo Alto, CA 94304  
USA

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

Bernard Aboba  
Microsoft  
One Microsoft Way  
Redmond, WA 98052  
USA

Email: bernard\_aboba@outlook.com

