Network Working Group
Internet-Draft
Intended status: Experimental
Expires: December 12, 2010

H. Alvestrand
j

June 10, 2010

TOC

A Datagram Transport for the WEBM profile draft-alvestrand-webm-datagram-00

Abstract

This document describes a combination and profiling of existing IETF protocols to provide a datagram service that is suitable as a generic transport substrate for the WEBM family of real-time audio/video applications.

Requirements Language

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 **RFC 2119** [RFC2119].

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 December 12, 2010.

Copyright Notice

Table of Contents

---

1. Introduction

TOC

When transporting audio / video data between participants on the current Internet, there are a number of obstacles to be faced.

Among them are NAT boxes, firewalls, connection interruptions, the availability of multiple paths between participants, and capacity issues.

This memo describes a combination of existing protocols that can be used to achieve a seamless datagram transport service across this very heterogenous environment.

---

2. Terminology

TOC

This draft uses a couple of commonly used terms in quite specific ways. The reader is advised to study these definitions carefully.

(TODO: Agree on terminology to use)

Session
: An association with two endpoints, between which datagrams flow.

Datagram
: A sequence of octets, of a given length. In this specification, a datagram does not carry addressing information.

Channel
: One means of transporting a datagram over a session. A session may have multiple channels at any time.

Endpoint
: One end of a session. This document does not distinguish between an initiator and a responder endpoint.

Control channel
: A means of communication between the endpoints of a session that does not require a transport to be active. Typically, authentication, authorization and negotiation is carried out over the control channel. The specification of the control channel is out of scope for this specification.

---

3. Service model

TOC

The basic model presented is a datagram model. On top of this one can layer various services, such as pseudoTCP (REF), RTP (REF) or any other higher layer protocol that is capable of running across a datagram service.

The addressing model departs from the traditional Internet model in that end point addresses are not used for endpoint identification, only for channel establisment; instead, an initial packet exchange, using ICE (REF), is used to bind a channel to a prenegotiated session.

The datagram service is not completely transparent; in particular, it is not possible to carry a datagram where the two highest bits of the first octet are zero and octet 5 to 8 contain the value 0x2112A442, since these datagrams are reserved for use of the STUN protocol (RFC 5389 section 6).

4. Channel types

TOC

4.1. UDP channel

TOC

An UDP channel is negotiated using ICE. Each datagram is simply carried as the content of an UDP packet.

4.2. TCP channel

TOC

A TCP channel consists of a TCP connection, over which are sent datagrams packaged according to (REF). The binding of a TCP channel is done by executing an ICE negotiation over the first few packets passed across the TCP channel.

4.3. TLS channel

TOC

A TLS channel consists of a standard TLS negotiation, followed by passing datagrams over the TLS record layer; the length fields of (REF) are not used. A TLS channel is bound to its session by <insert process description>.

4.4. DTLS channel

TOC

A DTLS channel is created by executing a DTLS connection negotiation, followed by datagram exchange, where the datagrams are protected by DTLS mechanisms. The DTLS channel is bound to its session by <insert process>.

4.5. Channels with relay

TOC

(this section intentionally left blank for now)

5. Channel setup, teardown and usage

TOC

The service model envisioned here is that all datagrams arriving on a session are considered equally valid. The session gives no guarantees against duplication, loss or reordering; such concerns are left to the higher protocol layers.

The expected normal usage is that two endpoints will exchange addressing information that can be used for a series of potential channels, that the endpoints will probe for working channels using ICE (RFC 5245), and use the "best" candidate, while using the STUN probing facilities to keep some number of "second best" candidates alive if the "best" candidate stops working.

A data-sending endpoint may unilaterally decide to start or stop using an established channel at any time. No negotiation is necessary.

A receiving endpoint will learn that a channel has been removed by not seeing any more STUN keepalive messages on that channel within <timeout>.

A session is considered closed when all channels that have been successfully established have timed out.

6. IANA Considerations

TOC

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

TOC

As with all layered protocols, it is a matter for the application to decide which level security should be provided at. For instance, an RTP session protected using SRTP <ref> can be considered to not need any further safeguards against interception, modification or replay, so can be passed "in the clear" across any channel type here. For data without such protection, adequate measures need to be taken; in particular, it is trivially easy for someone with the ability to snoop and insert packets to insert fake packets into an established UDP channel.

The main defense against denial-of-service attacks is the fact that the ICE mechanisms were designed for low cost refusal of unauthorized connections.

8. Acknowledgements

TOC

9. Normative References

TOC

[RFC2119] **Bradner, S.**, "**Key words for use in RFCs to Indicate Requirement Levels**," BCP 14, RFC 2119, March 1997 (**TXT**, **HTML**, **XML**).

Author's Address

TOC

Harald Tveit Alvestrand
j