Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 1, 2011

H. Alvestrand

June 30, 2010

TOC

Real Time Protocols for Brower-based Applications draft-alvestrand-webm-protocols

Abstract

This document gives an overview of a protocol suite intended for use with real-time applications that can be deployed in browsers - "real time communication on the Web".

It intends to serve as a starting and coordination point to make sure all the parts that are needed to achieve this goal are findable, and that the parts that belong in the Internet protocol suite are fully specified and on the right publication track.

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 January 1, 2011.

Copyright Notice

Table of Contents

---

1. Introduction

TOC

The Internet was, from very early in its lifetime, considered a possible veichle for the deployment of real-time, interactive applications - with the most easily imaginable being audio conversations (aka "Internet telephony") and videoconferencing.

The first attempts to build this were dependent on special networks, special hardware and custom-built software, often at very high prices or at low quality, placing great demands on the infrastructure.

As the available bandwidth has increased, and as processors and other hardware has become ever faster, the barriers to participation have decreased, and it is possible to deliver a satisfactory experience on commonly available computing hardware.

Still, there are a number of barriers to the ability to communicate universally - one of these is that there are, as of yet, no single set of communication protocols that all agree should be made available for communication; another is the sheer lack of universal identification systems (such as is served by telephone numbers or email addresses in other communications systems).

Development of The Universal Solution has proved hard, however, for all the usual reasons. This memo aims to take a more building-block-oriented approach, and try to find consensus on a set of substrate components that we think will be useful in any real-time communications systems.

The last few years have also seen a new platform rise for deployment of services: The browser-embedded application, or "Web application". It turns out that as long as the browser platform has the necessary interfaces, it is possible to deliver almost any kind of service on it.

Traditionally, these interfaces have been delivered by plugins, which had to be downloaded and installed separately from the browser; in the development of HTML5, much promise is seen by the possiblitiy of making those interfaces available in a standardized way within the browser.

Other efforts focus on making standardized APIs and interfaces available, within or alongside the HTML5 effort, for those functions; this memo concentrates on specifying the protocols and subprotocols that are needed to specify the interactions that happen across the network.

2. Functionality groups

TOC

The functionallity groups that are needed can be specified, more or less from the bottom up, as:

- Data transport: TCP, UDP and the means to set up secure connections between entities.
- Data framing: RTP and other data formats that serve as containers.
- Data formats: Codec specifications, format specifications and functionality specifications for the data passed between systems. Audio and video codecs, as well as formats for data and document sharing, belong in this category.
- Connection management: Setting up connections, agreeing on data formats, changing data formats during the duration of a call; SIP and Jabber/XMPP belong in this category.
- Presentation and control: What needs to happen in order to ensure that interactions behave in a non-surprising manner. This can include floor control, screen layout, voice activated image switching and other such functions - where part of the system require the cooperation between parties. Cisco/Tandberg's TIP was one attempt at specifying this functionality.
- Local system support functions: These are things that need not be specified uniformly, because each participant may choose to do these in a way of the participant's choosing, without affecting the bits on the wire in a way that others have to be cognizant of. Examples in this category include echo cancellation (some forms of it), local authentication and authorization mechanisms, OS access control and the ability to do local recording of conversations.

Within each functionality group, it is important to preserve both freedom to innovate and the ability for global communication. Freedom to innovate is helped by doing the specification in terms of interfaces, not implementaiton; any implementation able to communicate according to the interfaces is a valid implementation. Ability to communicate globally is helped both by having core specifications be unencumbered by IPR issues and by having the formats and protocols be fully enough specified to allow for independent implementation.

One can think of the three first groups as forming a "media transport infrastructure", and of the three last groups as forming a "media service". In many contexts, it makes sense to use a common specification for the media transport infrastructure, which can be embedded in browsers and accessed using standard interfaces, and "let a thousand flowers bloom" in the "media service" layer; to achieve interoperable services, however, at least the first five of the six groups need to be specified.

3. Data transport

TOC

Datagram transport is is the subject of a separate draft, "A Datagram Transport for the WEBM profile".

For octet-stream transport, TCP is used. (QUESTION: Do we need a TCP relay specification?)

4. Data framing

TOC

RTP **[RFC3550]**and SRTP **[RFC3711]**. The RTP/SAVP profile, defined as part of SRTP, is supported.

---

5. Data formats

TOC

In audio, the SILC/CELT merged codec**[I-D.valin-codec-prototype]** have been (will be?) made freely available.

In video, the VP8 and Theora codecs are freely available. H.264/AVC and H.264/SVC **[I-D.ietf-avt-rtp-svc]** are widely enough used that it gives a wider range of communcations partners if they are supported.

---

6. Connection management

TOC

XMPP, and its Jabber component, has proved a versatile tool in building interoperable communities. In order to share communication with other systems, SIP may also be supported.

---

7. Presentation and control

TOC

This area is still unstable.

---

8. Local system support functions

TOC

These do not need to be specified, but it is good to speak about available components here.

---

9. IANA Considerations

TOC

This document makes no request of IANA.

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

---

10. Security Considerations

TOC

Security of the web-enabled real time communications comes in two pieces:

- Security of the components: The browsers, and other servers involved. The most target-rich environment here is probably the browser; the aim here should be that the introduction of WEBM components introduces no additional vulnerability.

- Security of the communication channels: It should be easy for a participant to reassure himself of the security of his communication - by verifying the crypto parameters of the links he himself participates in, and to get reassurances from the other parties to the communication that they promise that appropriate measures are taken.
- Part of that last part is securing that the participants are who they say they are (when appropriate).

11. Acknowledgements

TOC

12. Normative References

TOC

[I-D.ietf-avt-rtp-svc]	Wenger, S., Wang, Y., Schierl, T., and A. Eleftheriadis, "**RTP Payload Format for SVC Video**," draft-ietf-avt-rtp-svc-21 (work in progress), April 2010 (**TXT**).
[I-D.valin-codec-prototype]	Valin, J. and K. Vos, "**Prototype for IETF Interactive Audio Codec**," draft-valin-codec-prototype-00 (work in progress), July 2010 (**TXT**).
[RFC2119]	**Bradner, S.**, "**Key words for use in RFCs to Indicate Requirement Levels**," BCP 14, RFC 2119, March 1997 (**TXT**, **HTML**, **XML**).
[RFC3550]	Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "**RTP: A Transport Protocol for Real-Time Applications**," STD 64, RFC 3550, July 2003 (**TXT**, **PS**, **PDF**).
[RFC3711]	Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "**The Secure Real-time Transport Protocol (SRTP)**," RFC 3711, March 2004 (**TXT**).

Author's Address

TOC

Harald T. Alvestrand