Network Working Group
Request for Comments: 4855
Obsoletes: 3555
Category: Standards Track

S. Casner
Packet Design

February 2007

Media Type Registration of RTP Payload Formats

Status of This Memo

Copyright Notice

Abstract

This document specifies the procedure to register RTP payload formats as audio, video, or other media subtype names. This is useful in a text-based format description or control protocol to identify the type of an RTP transmission.

Table of Contents

1. Introduction

RFC 4288 [1] defines media type specification and registration procedures that use the Internet Assigned Numbers Authority (IANA) as a central registry. That document covers general requirements independent of particular application environments and transport modes. This document defines the specific requirements for registration of media types for use with the Real-time Transport Protocol (RTP), RFC 3550 [2], to identify RTP payload formats.

1.1. Terminology

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 [3] and indicate requirement levels for implementations compliant with this specification.

2. Procedure For Registering Media Types for RTP Payload Types

Registering an RTP payload type as a media type follows the same procedures as described in RFC 4288 [1] and uses the registration template shown in Section 10 of that RFC. To specify how the particular payload format is transported over RTP, some additional information is required in the following sections of that template:

Required parameters:

If the payload format does not have a fixed RTP timestamp clock rate, then a "rate" parameter is required to specify the RTP timestamp clock rate. A particular payload format may have additional required parameters.

Optional parameters:

Most audio payload formats can have an optional "channels" parameter to specify the number of audio channels included in the transmission. The default channel order is as specified in RFC 3551 [4]. Any payload format, but most likely audio formats, may also include the optional parameters "ptime" to specify the recommended length of time in milliseconds represented by the media in a packet, and/or "maxptime" to specify the maximum amount of media that can be encapsulated in each packet, expressed as time in milliseconds. The "ptime" and "maxptime" parameters are defined in the Session Description Protocol (SDP) [5].

A particular payload format may have additional optional parameters. As allowed in Section 4.3 of [1], new parameters MAY be added to RTP media types that have been previously

defined, but the new parameters MUST NOT change existing functionality and it MUST be possible for existing implementations to ignore the additional parameters without impairing operation.

Encoding considerations:
Most RTP payload formats include binary or framed data as described in Section 4.8 of [1]. The appropriate encoding considerations MUST be noted.

Published specification:
A description of the media encoding and a specification of the payload format must be provided, usually by reference to an RTP payload format specification RFC. That RFC may be separate, or the media type registration may be incorporated into the payload format specification RFC. The payload format specification MUST include the RTP timestamp clock rate (or multiple rates for audio encodings with multiple sampling rates).

A reference to a further description of the data compression format itself should be provided, if available.

Restrictions on usage:
The fact that the media type is defined for transfer via RTP MUST be noted, in particular, if the transfer depends on RTP framing and hence the media type is only defined for transfer via RTP.

Depending on whether or not the type has already been registered for transfer with a non-RTP protocol (e.g., MIME mail or http), several different cases can occur:

a) Not yet registered as a media type

A new registration should be constructed using the media type registration template. The registration may specify transfer via other means in addition to RTP if that is feasible and desired. The appropriate encoding considerations must be specified, and the restrictions on usage must specify whether the type is only defined for transfer via RTP or via other modes as well.

Optional parameters may be defined as needed, and it must be clearly stated to which mode(s) of transfer the parameters apply.

b) Media type exists for a non-RTP protocol

The restrictions on usage of the existing type should be changed, if present, or added, if not, to indicate that the type can also be transferred via RTP.

RTP-specific parameters may be added, and it must be clearly stated that these are only to be used when the media type is transmitted via RTP transport.

c) Update an existing media type for RTP to be used for a non-RTP protocol

The restrictions on usage of the existing type should be changed to indicate that the type can also be transferred via a non-RTP protocol (e.g., SMTP, HTTP).

Non-RTP-specific parameters can be added, and it must be clearly stated that these are only to be used when the media type is transmitted via a non-RTP transport.

2.1. Example Media Type Registration

The following sample registration of a fake media type audio/example provides examples for some of the required text. References to RFC nnnn would be replaced by references to the RFC that contains the payload format specification and the media type registration.

Type name: audio

Subtype name: example

Required parameters:
rate: RTP timestamp clock rate, which is equal to the sampling rate. The typical rate is 8000; other rates may be specified.

Optional parameters:
channels: number of interleaved audio streams, either 1 for mono or 2 for stereo, and defaults to 1 if omitted. Interleaving takes place between on a frame-by-frame basis, with the left channel followed by the right channel.

ptime: recommended length of time in milliseconds represented by the media in a packet (see RFC 4566).

maxptime: maximum amount of media that can be encapsulated in each packet, expressed as time in milliseconds (see RFC 4566).

Encoding considerations:
This media type is framed binary data (see RFC 4288, Section 4.8).

Security considerations: See Section n of RFC nnnn

Interoperability considerations:
Some receivers may only be capable of receiving single-channel audio.

Published specification: RFC nnnn

Applications that use this media type:
Audio and video streaming and conferencing tools.

Additional information: none

Person & email address to contact for further information:
Fred Audio <fred@example.com>

Intended usage: COMMON

Restrictions on usage:
This media type depends on RTP framing, and hence is only defined for transfer via RTP (RFC 3550). Transfer within other framing protocols is not defined at this time.

Author:
Fred Audio

Change controller:
IETF Audio/Video Transport working group delegated from the IESG.

2.2. Restrictions on Sharing a Subtype Name

The same media subtype name MUST NOT be shared for RTP and non-RTP (file-based) transfer methods unless the data format is the same for both methods. The data format is considered to be the same if the file format is equivalent to a concatenated sequence of payloads from RTP packets not including the RTP header or any RTP payload-format header.

The file format MAY include a magic number or other header at the start of the file that is not included when the data is transferred via RTP.

A second requirement for sharing a media subtype name is that the sets of required parameters must be the same for both methods.

For cases where the data format or required parameters cannot be the same for RTP and non-RTP transfer methods, the data formats MUST be registered as separate types. It is RECOMMENDED that the subtype names be related, such as by using a common root plus a suffix. For those cases where a suffix is applied in the subtype name for the RTP transfer method, the suffix "+rtp" is suggested.

3. Mapping to SDP Parameters

The representation of a media type is specified in the syntax of the Content-Type header field in RFC 2045 [6] as follows:

```
type "/" subtype  *(";" parameter)
```

Parameters may be required for a particular type or subtype or they may be optional. For media types that represent RTP payload formats, the parameters "rate", "channels", "ptime", and "maxptime" have general definitions (given above) that may apply across types and subtypes. The format for a parameter is specified in RFC 2045 as

```
attribute "=" value
```

where attribute is the parameter name and the permissible values are specified for each parameter. RFC 2045 specifies that a value MUST be present and that the value MUST be a quoted string if it contains any of the special characters listed in that RFC.

The information carried in the media type string has a specific mapping to fields in the Session Description Protocol (SDP) [5], which is commonly used to describe RTP sessions. The mapping is as follows:

- The media type (e.g., audio) goes in SDP "m=" as the media name.
- The media subtype (payload format) goes in SDP "a=rtpmap" as the encoding name.
- The general (possibly optional) parameters "rate" and "channels" also go in "a=rtpmap" as clock rate and encoding parameters, respectively.
- The general (and optional) parameters "ptime" and "maxptime" go in the SDP "a=ptime" and "a=maxptime" attributes, respectively.

- Any payload-format-specific parameters go in the SDP "a=fmtp" attribute. The set of allowed parameters is defined by the RFC that specifies the payload format and MUST NOT be extended by the media type registration without a corresponding revision of the payload format specification. The format and syntax of these parameters may also be defined by the payload format specification, but it is suggested that the parameters be copied directly from the media type string as a semicolon separated list of parameter=value pairs. For payload formats that specify some other syntax for the fmtp parameters, the registration of that payload format as a media type must specify what the parameters are in MIME format and how to map them to the "a=fmtp" attribute.

An example mapping is as follows:

```
audio/L16; rate=48000; channels=2; ptime=5; emphasis=50-15

m=audio 49170 RTP/AVP 97
a=rtpmap:97 L16/48000/2
a=fmtp:97 emphasis=50-15
a=ptime:5
```

Note that the payload format (encoding) names defined in the RTP Profile [4] are commonly shown in upper case. Media subtype names are commonly shown in lower case. These names are case-insensitive in both places. Similarly, parameter names are case-insensitive both in media type strings and in the default mapping to the SDP a=fmtp attribute.

4. Changes from RFC 3555

This document updates RFC 3555 to conform to the revised media type registration procedures in RFC 4288 [1]. Whereas RFC 3555 required the encoding considerations to specify transfer via RTP, that is now specified under restrictions on usage. This document also specifies the conditions under which new optional parameters may be added to a previously defined RTP media type and adds a new Section 2.2 to clarify the requirements for sharing a media type among RTP and non-RTP transfer methods.

RFC 3555 included media type registrations for the RTP payload formats defined in the RTP Profile for Audio and Video Conferences, RFC 3551 [4]. Those media type registrations have been removed from this document. Some of them have been assembled into a separate companion RFC 4856 [8], leaving out those that have been, or are intended to be, registered in revisions of their own payload format specification RFCs.

Philipp Hoschka is a co-author of RFC 3555; his contributions to the foundation of this document are appreciated.

5. Security Considerations

The media type registration procedure specified in this memo does not impose any security considerations on its own. Also, registrations conforming to this procedure do not themselves impose security risks. However, use of the media type being registered could very well impose security risks:

- Any media type that contains "active content" imposes the risk of malicious side-effects unless execution of that content is adequately constrained.
- Several audio and video encodings are perfect for hiding data using steganography.
- The RTP specification, RFC 3550, provides security considerations for the transport of audio and video data over RTP, including the use of encryption where confidentiality is required.

Therefore, each media type registration is required to state any security considerations that apply to the use of that type. The remainder of this section is copied from RFC 4288 [1], which specifies media type registration procedures in general.

An analysis of security issues MUST be done for all types registered in the standards tree. A similar analysis for media types registered in the vendor or personal trees is encouraged but not required. However, regardless of what security analysis has or has not been done, all descriptions of security issues MUST be as accurate as possible regardless of registration tree. In particular, a statement that there are "no security issues associated with this type" MUST NOT be confused with "the security issues associated with this type have not been assessed".

There is absolutely no requirement that media types registered in any tree be secure or completely free from risks. Nevertheless, all known security risks MUST be identified in the registration of a media type, again regardless of registration tree.

The security considerations section of all registrations is subject to continuing evaluation and modification, and in particular MAY be extended by use of the "comments on media types" mechanism described in RFC 4288, Section 6.

Some of the issues that should be looked at in a security analysis of a media type are:

- Complex media types may include provisions for directives that institute actions on a recipient's files or other resources. In many cases, provision is made for originators to specify arbitrary actions in an unrestricted fashion that may then have devastating effects. See the registration of the application/postscript media type in RFC 2046 [7] for an example of such directives and how they should be described in a media type registration.

- All registrations MUST state whether or not they employ such "active content", and if they do, they MUST state what steps have been taken to protect users of the media type from harm.

- Complex media types may include provisions for directives that institute actions that, while not directly harmful to the recipient, may result in disclosure of information that either facilitates a subsequent attack or else violates a recipient's privacy in some way. Again, the registration of the application/postscript media type illustrates how such directives can be handled.

- A media type that employs compression may provide an opportunity for sending a small amount of data that, when received and evaluated, expands enormously to consume all of the recipient's resources. All media types SHOULD state whether or not they employ compression, and if they do they should discuss what steps need to be taken to avoid such attacks.

- A media type might be targeted for applications that require some sort of security assurance but not provide the necessary security mechanisms themselves. For example, a media type could be defined for storage of confidential medical information that in turn requires an external confidentiality service or is designed for use only within a secure environment.

6. IANA Considerations

The purpose of this document is to specify the requirements and procedures for registering RTP payload formats in the IANA media type registry. No registrations are defined here.

7. References

7.1. Normative References

[1] Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures", BCP 13, RFC 4288, December, 2005.

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

[3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[4] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", RFC 3551, July 2003.

[5] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.

[6] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.

[7] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, November 1996.

7.2. Informative References

[8] Casner, S., "Media Type Registration of Payload Formats in the RTP Profile for Audio and Video Conferences", RFC 4856, February 2007.

Author's Address

Stephen L. Casner
Packet Design
3400 Hillview Avenue, Building 3
Palo Alto, CA 94304
United States

Phone: +1 650 739-1843
EMail: casner@acm.org

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.