

AVT
Internet-Draft
Updates: [3550](#) (if approved)
Intended status: Standards Track
Expires: December 19, 2010

A. Begen
Cisco
C. Perkins
University of Glasgow
D. Wing
Cisco
June 17, 2010

**Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names
(CNAMEs)
draft-ietf-avt-rtp-cnames-00**

Abstract

The RTP Control Protocol (RTCP) Canonical Name (CNAME) is a persistent transport-level identifier for an RTP endpoint. While the Synchronization Source (SSRC) identifier of an RTP endpoint may change if a collision is detected, or when the RTP application is restarted, the CNAME is meant to stay unchanged, so that RTP endpoints can be uniquely identified and associated with their RTP media streams. For proper functionality, CNAMEs should be unique within the participants of an RTP session. However, the existing guidelines for choosing the RTCP CNAME provided in the RTP standard are insufficient to achieve this uniqueness. This memo updates these guidelines to allow endpoints to choose unique CNAMEs.

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 19, 2010.

Copyright Notice

Copyright (c) 2010 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.

Table of Contents

1.	Introduction	3
2.	Requirements Notation	3
3.	Deficiencies with Earlier RTCP CNAME Guidelines	3
4.	Choosing an RTCP CNAME	4
4.1.	Persistent vs. Per-Session CNAMEs	4
4.2.	Guidelines	5
5.	Security Considerations	6
6.	IANA Considerations	6
7.	Acknowledgments	6
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	7
	Authors' Addresses	7

1. Introduction

In [Section 6.5.1 of \[RFC3550\]](#), there are a number of recommendations for choosing a unique RTCP CNAME for an RTP endpoint. However, in practice, some of these methods are not guaranteed to produce a unique CNAME. This memo updates guidelines for choosing CNAMEs, superseding those presented in [Section 6.5.1 of \[RFC3550\]](#).

2. Requirements Notation

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 [\[RFC2119\]](#).

3. Deficiencies with Earlier RTCP CNAME Guidelines

The recommendation in [\[RFC3550\]](#) is to generate the CNAME of the form "user@host" for multiuser systems, or "host" if the username is not available. The "host" part is specified to be the fully qualified domain name (FQDN) of the host from which the real-time data originates. However, FQDNs are not necessarily unique, and can sometimes be common across several endpoints in large service provider networks. Thus, the use of FQDN as the CNAME is strongly discouraged.

IPv4 addresses are also suggested for use in CNAMEs in [\[RFC3550\]](#), where the "host" part of the RTCP CNAME is the numeric representation of the IP address of the interface from which the RTP data originates. As noted in [\[RFC3550\]](#), the use of private network address space [\[RFC1918\]](#) can result in hosts having network addresses that are not globally unique. However, this shared use of the same IP address can also occur with public IP addresses if multiple hosts are assigned the same public IP address and connected to a Network Address Translation (NAT) device [\[RFC3022\]](#). When multiple hosts share the same IP address, whether private or public, using the IP address as the CNAME leads to CNAMEs that are not necessarily unique.

[\[RFC3550\]](#) also notes that if hosts with private addresses and no direct IP connectivity to the public Internet have their RTP packets forwarded to the public Internet through an RTP-level translator, they may end up having non-unique CNAMEs. [\[RFC3550\]](#) suggests that such applications provide a configuration option to allow the user to choose a unique CNAME, and puts the burden on the translator to translate CNAMEs from private addresses to public addresses if necessary to keep private addresses from being exposed. Experience has shown that this does not work well in practice.

4. Choosing an RTCP CNAME

It is difficult, and in some cases impossible, for a host to determine if there is a NAT between itself and its RTP peer. Furthermore, even some public IPv4 addresses can be shared by multiple hosts in the Internet. Thus, using the numeric representation of the IPv4 address as the "host" part of the RTCP CNAME is NOT RECOMMENDED.

4.1. Persistent vs. Per-Session CNAMEs

The RTCP CNAME can either be persistent across different RTP sessions for an RTP endpoint, or it can be unique per session, meaning that an RTP endpoint chooses a different CNAME for each RTP session.

Persistent CNAMEs: To provide a binding across multiple media tools used by one participant in a set of related RTP sessions, the CNAME SHOULD be fixed for that participant. An RTP endpoint that is emitting multiple related streams that require synchronization at the other endpoint(s) SHOULD use a persistent CNAME. A persistent CNAME is also useful to facilitate third-party monitoring, allowing network management tools to correlate the ongoing quality of service across multiple RTP sessions for fault diagnosis and to understand long-term network performance statistics.

Note: A persistent CNAME will not provide a unique identifier for each source if an application permits a user to generate multiple sources from one host. Such an application would have to rely on the SSRC to further identify the source, or the profile for that application would have to specify additional syntax for the CNAME identifier.

Note: If each RTP application creates its CNAME independently, the resulting CNAMEs may not be identical as would be required to provide a binding across multiple media tools belonging to one participant in a set of related RTP sessions. If cross-media binding is required, it may be necessary for the CNAME of each tool to be externally configured with the same value by a coordination tool.

Per-Session CNAMEs: The advantage of this approach is that the CNAME is unique for each RTP session. This prevents the CNAME from being used for traffic analysis. In other words, the RTP endpoints cannot be identified based on their CNAMEs. This provides privacy, but inhibits the use of RTCP as a tool for long-term network management and monitoring.

4.2. Guidelines

RTP endpoints SHOULD practice one of the following guidelines in choosing RTCP CNAME:

- o Given that IPv6 addresses are naturally unique, an endpoint MAY use one of its IPv6 address(es) as the "host" part of its CNAME regardless of whether that IPv6 interface is being used for RTP communication or not. If the RTP endpoint is associated with an IPv6 privacy address [[RFC4941](#)] or a unique local IPv6 unicast address [[RFC4193](#)], that address MAY be used as well. The IPv6 address is converted to its textual representation [[I-D.ietf-6man-text-addr-representation](#)], resulting in a printable string representation as short as 24 bits and as long as 304 bits. Using IPv6 addresses as the "host" part of a CNAME was originally suggested in [[RFC3550](#)].
- o An endpoint that does not know its fully qualified domain name, and is configured with a private IP address on the interface it is using for RTP communication, MAY use the numeric representation of the layer-2 (MAC) address of that interface as the "host" part of its CNAME. For IEEE 802 MAC addresses, such as Ethernet, the standard colon-separated hexadecimal format is to be used, e.g., "00:23:32:af:9b:aa" resulting in a 136-bit printable string representation.
- o An endpoint MAY use its Universally Unique Identifier (UUID) [[RFC4122](#)] to generate the "host" part of its CNAME. The string representation described in [Section 3 of \[RFC4122\]](#) SHOULD be used without "urn:uuid:", which results in a 288-bit printable string representation.
- o To generate a per-session CNAME, an endpoint MAY perform SHA1-HMAC [[RFC2104](#)] on the concatenated values of the RTP endpoint's initial SSRC, the source and destination IP addresses and ports, and a randomly-generated value [[RFC4086](#)], and then truncate the 160-bit output to 96 bits and finally convert the 96 bits to ASCII using Base64 encoding [[RFC4648](#)]. This results in a 128-bit printable string representation. Note that the CNAME MUST NOT change if an SSRC collision occurs, hence only the initial SSRC value chosen by the endpoint is used.

Each of the techniques is equally effective in generating unique CNAMEs, and an RTP application MAY choose any of these techniques to use.

5. Security Considerations

The security considerations of [\[RFC3550\]](#) apply to this document as well.

In some environments, notably telephony, a fixed CNAME value allows separate RTP sessions to be correlated and eliminates the obfuscation provided by IPv6 privacy addresses [\[RFC4941\]](#) or IPv4 NAPT [\[RFC3022\]](#). Secure RTP (SRTP) [\[RFC3711\]](#) can help prevent such correlation by encrypting Secure RTCP (SRTCP) but it should be noted that SRTP only mandates SRTCP integrity protection (not encryption). Thus, RTP applications used in such environments should consider encrypting their SRTCP or generate a per-session CNAME as discussed in [Section 4.1](#).

6. IANA Considerations

There are no IANA considerations in this document.

7. Acknowledgments

Thanks to Marc Petit-Huguenin who suggested to use UUIDs in generating CNAMEs.

8. References

8.1. Normative 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.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 4941](#), September 2007.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique Identifier (UUID) URN Namespace", [RFC 4122](#), July 2005.

- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997.
- [RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", [BCP 106](#), [RFC 4086](#), June 2005.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", [RFC 4648](#), October 2006.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", [RFC 3711](#), March 2004.
- [I-D.ietf-6man-text-addr-representation]
Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation",
[draft-ietf-6man-text-addr-representation-07](#) (work in progress), February 2010.

8.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), February 1996.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), January 2001.

Authors' Addresses

Ali Begen
Cisco
181 Bay Street
Toronto, ON M5J 2T3
CANADA

Email: abegen@cisco.com

Colin Perkins
University of Glasgow
Department of Computing Science
Glasgow, G12 8QQ
UK

Email: csp@cspcrkins.org

Dan Wing
Cisco Systems, Inc.
170 West Tasman Dr.
San Jose, CA 95134
USA

Email: dwing@cisco.com

