

Network WG
Internet-Draft
Intended status: PS
Expires: January 15, 2014

James Polk
Toerless Eckert
Cisco Systems
July 15, 2013

Differentiated Services Delay-and-Loss vs. Loss-Rate-Adaptive
Service Classes
draft-polk-tsvwg-delay-vs-loss-ds-service-classes-00.txt

Abstract

This document discusses how RTCWeb/RMCAT applications could best leverage existing and new DiffServ assignments and why we think it is necessary to differentiate the assignments of delay-and-loss-based vs. just-loss-based rate-adaptive video applications.

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 15, 2014.

Copyright Notice

Copyright (c) 2013 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.

1. Introduction

The current guidelines for DSCP assignments of traffic in the IETF is the informational [RFC 4594](#) [[RFC4594](#)]. The TSV Working Group is currently reviewing how to update these assignments with one proposal being [[ID-4594bis](#)].

This document discusses how RTCWeb/RMCAT applications could best leverage [[RFC4594](#)] & [[ID-4594bis](#)] assignments and why we think it is necessary to differentiate the assignments of video for these type of applications from non-rate adaptive video flows (i.e., MPEG2).

The current text of [[ID-4594bis](#)] has audio flows within an interactive communication assigned to EF. The current text of [[RFC4594](#)] has interactive video flows would be assigned to AF4x.

The definition of AF4x directly matches the requirements for delay and loss-based rate-adaptive, self-friendly video traffic as targeted by RTCWeb applications and RMCAT congestion mechanisms. Splitting audio and video this way into separate traffic classes would specifically provide the benefit of guaranteeing no congestion/loss for the audio part while allowing congestion in video to cause rate adaptation where there is contention in the network. Unfortunately, the reality of deployments of AF4x in network in the last decade or longer does not match this original target of [[RFC4594](#)], instead relying only on loss in the network to indicate congestion.

2. 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 [[RFC2119](#)] when they appear in ALL CAPS. These words may also appear in this document in lower case as plain English words, absent their normative meanings.

3. Rate-Adaptive Traffic is Inelastic (and often badly provisioned)

Because delay-based rate-adaptive video for conferencing has for the most part been a fairly recent development, the majority of video traffic deployed in many networks into AF41 or AF4x in general has been inelastic, highly loss-sensitive video traffic at fixed bitrates. This type of video deployment is accompanied by some form of either

- o a vendor specific "Locations CAC" (LCAC) mechanism,
- o an on-path RSVP based CAC or

- o a much more commonly what is perceived to be sufficient overprovisioning.

This is because LCAC is often considered to be too provisioning intense and RSVP CAC is not supported on a critical mass of application/devices to be deployable in most networks.

In addition to Video, AF4x often also carries in practice, the (non-rate adaptive) voice part of collaborative sessions, primarily because creating lip-sync between EF audio and AF4x video traffic when they experiences differing jitter and latency in the network was seen as a complexity that especially lower end hardware based video endpoints wanted to avoid implementing. Additional reasons for this:

- o Endpoints are too lazy to mark, so network devices mark and cannot tell what flow is video vs. audio;
- o Implementers simply did not give the option to mark voice or video differently;
- o People want the same per-hop behavior for audio and video. After all, getting one type of media there faster is of little value since it just has to be put in a buffer and queued, waiting for the other flow.

Alas, what might have been sufficient overprovisioning often turns out to be "pray and suffer" in the face of often badly controllable addition of new applications, devices, users or increased in utilization. For example, the busy hour in large enterprises often shrinks and becomes even more busy when expanding operations across

multiple time zones. All these problems are the primary reason to propose the move to rate adaptive video, as seen in many recent products and IETF RTCweb/RMCAT work, but with the often long-term investments in a lot of this equipment in enterprises and cost of migrating deployment designs, the issues with this current use of AF4x is not going to go away for a long time. The recent interest in more application or network based circuit-breaker methods [\[ID-AVT-CB\]](#) for this type of inelastic traffic is also a good proof for this pain point (including the demands within the IETF to support them on any inelastic solution (i.e., RTCWeb & RMCAT efforts)).

Because of these fairly widespread deployments of AF4x and their issues, these flows make for a fairly bad PHB to put RMCAT/RTCweb traffic into. Rate adaptive traffic will always come up short against non-rate adaptive, incongestible traffic when being put into the same queue. This is especially true when provisioning for AF4x is leveraging the above "pray and suffer + circuit breaker" approach.

[4.](#) Jitter from Non-Rate Adaptive Video Traffic is Bad

Even if the existing non-rate-adaptive traffic was correctly provisioned such that there is no congestion, and even if all that traffic is known in a particular deployment to only use AF41, it would still not be a good idea to put rate-adaptive RTCweb/RMCAT traffic into AF42/AF43.

The reason is that according to the definition of these traffic classes, all AF4x traffic should be in a single queue because there must be no reordering between AF4x packets. The permissible differentiation between AF41/AF42/AF43 therefore are differential drop profiles via WRED or other methods. This means that even if there is never ongoing congestion caused by badly behaving legacy non-rate adaptive traffic, the new rate-adaptive video traffic would still suffer from the jitter introduced by that old video traffic because it has to share a queue. And it would suffer proportional to the amount of traffic in the queue, aka: it would suffer most during the busy hours.

This jitter from unknown/legacy/bad video traffic is especially troublesome because delay variation schemes considered for

rate-control in RMCAT will be conscious of that jitter and likely work less well when exposed to it.

It is hard enough - even unavoidable - for RMCAT to figure out how deal with competing TCP traffic on a single BE "Internet Queue". It is a total waste of effort and easily avoidable for RMCAT having to figure out how to deal with the jitter and loss introduced by legacy non-rate-adaptive video traffic in controlled environments where traffic can be separated by DSCP.

[5.](#) Suggested Behavior in the Network

[5.1](#) CS4 and CS4-Discardable for rate-adaptive video

This document is therefore asking to assign a Service Class to delay and loss-based rate-adaptive, self-friendly conversational video traffic that is separate from AF4x.

We think that CS4 has been a traffic class that so far has seen little use and most likely is the easiest traffic class to use for this traffic.

In addition, it would be very helpful for future improvements in delay and loss-based rate-adaptive video traffic in the network if there is a second traffic class, e.g., "CS4-Discardable", for lower-priority packets within video flows that can easily be dropped.

Given how rate-adaptive video will not always be able to avoid all

packet drops, video codecs can improve the visual quality in conjunction with the network by creating discardable P-frames (e.g., P-frames not referenced by other P-frames) and having the network preferentially drop those frames. This can easily achieved with two DSCP assigned values, where one has a higher drop priority than the other. This, of course, is exactly the logic also designated for AF41/AF42/AF43, in other words we are simply asking for a new PHB be assigned to the existing AF41 (suggest: CS4) and a new AF42 (suggested: CS4-Discardable).

[5.2](#) AF4x for Non-Rate-Adaptive Video

To represent the reality of deployments in the IETF guidelines, AF41 should be re-designated for non-rate adaptive conversational video with explicit admission control (off-path or non-path). AF42 should be re-designated for non-rate-adaptive conversational video without explicit admission control (e.g., relying solely on circuit breakers).

This proposal, in effect, is suggesting the text in [RFC 4594](#) regarding AF4X, as written, be ported/moved over to replace the text in that same RFC that was for CS4 PHB. New text will need to be written in the RFC4594bis document addressing AF4X taking over the more legacy traffic behaviors from non-adaptive (based on loss) video traffic.

[6.](#) Additional Recommendations for Other PHBs Affected

Specifying that delay and loss-based rate-adaptive video use CS4 (and 'CS4+' or 'CS4-Discardable' or 'CS4-D') means the current [RFC 4594](#) assignment of CS4 to the Realtime-Interactive (RTI) service class needs modification. Rather than create new definitions for CS4, currently occupied with the Realtime Interactive (RTI) PHB traffic according to [RFC 4594](#), our proposal is to move the RTI service class definitions to CS5 (as is the case in [[ID-4594bis](#)] - where the RTI will have a minimal effect on the existing installed base of implementations because it has few.

The CS5 PHB, which is assigned for H.323 and SIP signaling, would be then moved to another DSCP. Perhaps near Low-Latency Data (CS2) according to [[ID-4594bis](#)].

[7.](#) Acknowledgements

The authors would like to thank Paul Jones, Charles Eckel, Charles Ganzhorn, Fred Baker, Mo Zanaty, Michael Ramalho for their active participation in this effort.

[8.](#) IANA Considerations

If the WG agrees with this effort, then there will need to be a reassignment of AF4X and CS4, as well as the new assignment of an

additional CS4+ or CS4-discardable PHB.

9. Security Considerations

TBW

10. References

10.1 Normative References

- [RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997
- [RFC4594] J. Babiarz, K. Chan, F Baker, "Configuration Guidelines for Diffserv Service Classes", [RFC 4594](#), August 2006
- [ID-4594bis] J. Polk, "Standard Configuration of DiffServ Service Classes", "work in progress", March 2012
- [ID-AVT-CB] C. Perkins, V. Singh, "Multimedia Congestion Control: Circuit Breakers for Unicast RTP Sessions", work in progress, July 2013

10.2 Informative References

none at this time

Author's Addresses

James Polk
Cisco Systems, Inc.
8017 Hallmark Dr.
North Richland Hills, TX 76182
USA

Phone: +1 817 271 3552
Email: jmpolk@cisco.com

Toerless Eckert
Cisco Systems, Inc.
San Jose
US

Email: eckert@cisco.com