Network WG
Internet-Draft
Intended status: Informational
Expires: January 9, 2012

James Polk
Cisco
July 9, 2012

The Problem Statement for the Standard Configuration of DiffServ Service Classes
draft-polk-tsvwg-diffserv-stds-problem-statement-00.txt

Abstract

This document describes the problem statement on two recently proposed expansions to DiffServ. The first of these expansions proposes updating the informational RFC 4594 document to standards track status, while making the necessary changes to make it current; for example, creating more granular traffic treatments, some with new Per Hop Behaviors (PHB). The second proposal defines 6 new DiffServ Codepoints necessary from these new PHBs in the proposal within the first draft.

Table of Contents

1. Introduction 2
2. Brief Overview of RFC 4594 and RFC 5127 3
2.1 Brief Overview of RFC 4594 3
2.2 Brief Overview of RFC 5127 4
3. Brief Discussion of the RFC 4594 Update Draft 5
4. Conclusion and What's Next 7
5. Acknowledgements 7
6. IANA Considerations 7
7. Security Considerations 8
8. References 8
8.1 Normative References 8
8.2 Informative References 8
Author's Address 9

1. Introduction

Differentiated Services (DiffServ) [RFC2474] creates an IP header marking or indicator with which intermediate nodes (i.e., routers and switches) can make policy decisions. These 6-bit values are called Differentiated Services Codepoint Point (DSCP) values. DSCP values are used to differentiate packet treatment within an intermediate node, not across a network, as the conditions affecting that marking are different within each node. This is called Per Hop Behavior (PHB). In other words, even though a packet has the same DSCP from source to destination, it can and often does experience different treatment depending on the conditions of the nodes it traverses on its journey.

The DiffServ architecture allows for DSCP values within a packet to be changed, or remarked, any number of times. In other words, a packet can have its DSCP remarked at every layer-3 hop throughout the life of that packet. This practice actually occurs infrequently, but it is allowed.

At issue is a combination of the number of networks or endpoints that are choosing to use DiffServ markings, and the number of administrative domains (called "networks" in this document) a packet traverses with different policies for how packet flows of a similar type (e.g., a voice flow, or an email flow, etc.) are to be marked.

The community presently has RFC 4594 [RFC4594], which is an informational guideline on how networks can or should mark certain packet flows with differing traffic characteristics using DiffServ. There are several reasons why this informational RFC lacks the necessary clarity and strength to reach widespread adoption:

- confusion between RFC 4594 and RFC 5127 [RFC5127], the latter of which is for aggregating many 6-bit DSCP values into a 3-bit (8

value) field used specifically by service provider (SP) networks.

- some believe both RFCs are for SPs, while others ignore RFC 5127 and use RFC 4594 as if it were standards track or BCP.

- some believe RFC 5127 is for SPs only, and want RFC 4594 to reduce the number of DSCPs within its guidelines to recommend using only 3 or 4 DSCPs. This seems to stem from a manageability and operational perspective.

- some know RFC 4594 is informational and do not follow its guidelines specifically because it is informational.

- some use DSCP values that are not defined within RFC 4594, making mapping between different networks using similar or identical application flows difficult.

- some believe enterprise networks should not use either RFC except at the edge of their networks, where they directly connect to SP networks.

- some argue that the services classes guidance per class is too broad and are therefore not sure in which service class a particular application is to reside.

This document is not intended to reach RFC status. Rather, it is to stimulate discussion on both RFC 4594 and 5127 to lessen existing confusion within the community. It should be noted that RFC 4594 has an offered update within TSVWG [ID-4594-UPDATE]. This draft has created some heated discussions within that WG before and during the Paris IETF meeting.

First, we'll discuss briefly RFCs 4594 and 5127 in Section 2. Then we will discussion what the update to RFC 4594 proposes differently and what we expect to happen to RFC 5127 in Section 3.

2. Brief Overview of RFC 4594 and RFC 5127

2.1 Brief Overview of RFC 4594

Essentially, RFC 4594 is a guideline for how to choose which DSCP to use based on the traffic characteristics an application flow needs to experience within a network for optimal performance. RFC 4594 specifically points to several existing standards-track DiffServ RFCs to augment the text in each of those RFCs, without violating any of the rules within each of those documents. RFC 4594:

- painstakingly lays out definitions and guidelines for each service class.

o clearly indicates each service class's tolerance to delay, jitter

and packet loss.

- details the conditioning treatments at the Differentiated Services (DS) edge.

- categorizes traffic characteristics into 12 service classes utilizing one or more DSCPs:

Network Control	Broadcast Video
Telephony	Low-Latency Data
Signaling	OAM
Multimedia Conferencing	High-throughput Data
Realtime Interactive	Standard
Multimedia Streaming	Low-priority Data

2.2 Brief Overview of RFC 5127

At its barest, RFC 5127 recommends that, of the many service classes described within RFC 4594, each having different traffic characteristics, similar service classes be grouped or aggregated into 3, 4, or 5 markings for SP traversal. This limitation of the number of individual service classes is partly to reduce the number of separate distinctions traversing over their network because SPs have difficulty managing what is deemed 'too many' different classes. Another part for this reduction is customer expectations of meeting contractual Service Level Agreements (SLAs).

To this end, and perhaps because of it, MPLS was designed with only 8 values of priority differentiation, i.e., the 3 EXP bits. To be fair, LAN based IEEE has only a 3-bit priority field as well within its specifications, known as the Priority Code Point (PCP), as part of the 802.1Q header spec. IEEE 802.1e, which defines QoS over Wi-Fi, also only defines 8 levels (called User Priority or UP codes).

The result is to have the IETF within RFC 5127 recommend the following (which is Figure 2 within that RFC):

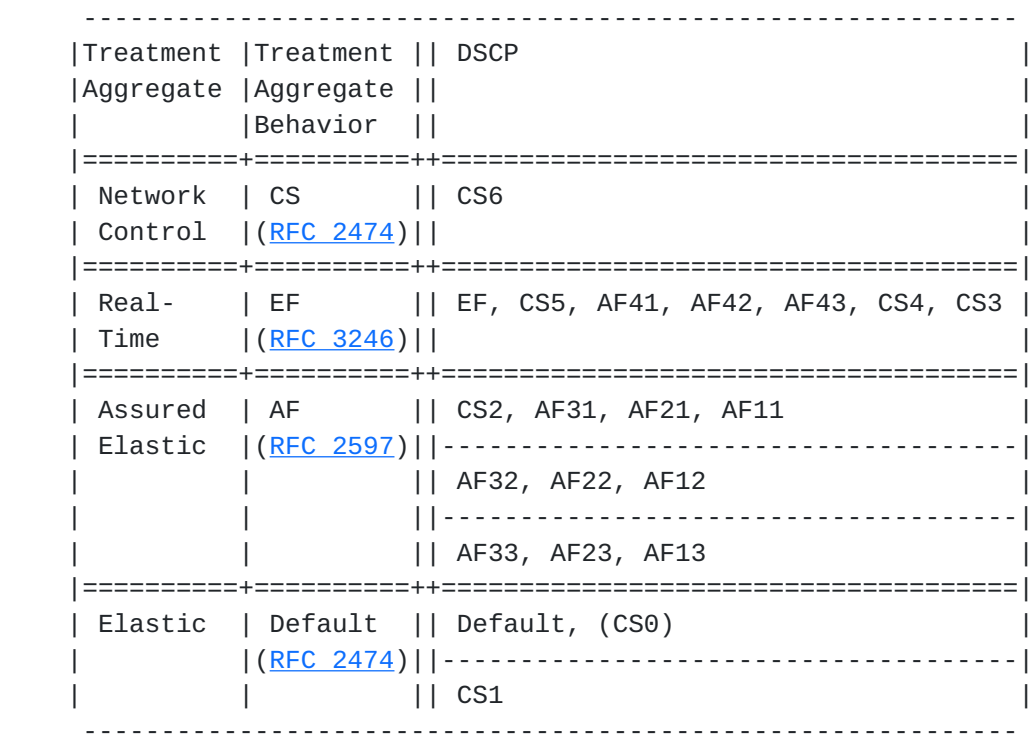

Figure 1: Treatment Aggregate Behavior

RFC 5127 goes on to recommend the marking and treatments on either side of the provider edge remain the same. In other words, the DSCP values remain the same and are used to determine which queue to place the packets into within the aggregates, where the packets are treated the same within that tunnel until the egress provider edge.

Many within enterprise networks do not pay attention to what RFC 5127 says because they are sufficiently removed from dealing with the constraints of very few DSCP values or the need to aggregate DSCP values into groups.

3. Brief Discussion of the RFC 4594 Update Draft

The RFC 4594 update draft [ID-4594-UPDATE] proposes to update what has occurred since RFC 4594 was written (i.e., 2006), in which more granular service classes can be differentiated by application requirements. For example, Figure 2 within RFC 4594 identifies "Telephony" as having 'Fixed-size small packets'. That is not true for today's video flow, therefore it needs to be modified. The update draft currently breaks out audio and video separately to reflect this different, as well as the ability to treat each traffic type differently within a network. Another example is gaming and TCP. The two were believed by most, and it is still believed by many that gaming requires a UDP delivery due to the requirements for

timely delivery of packets and that retransmissions would cause delays and bad things to happen to gaming applications. This was proved false within [ID-TCMTF], in which the author of that document

had a presentation showing TCP was used and viable.

[RFC5865] created a new Expedited Forwarding (EF) DSCP value called VOICE-ADMIT, the second time an application is identified within the DiffServ realm. The first was the service class Broadcast Video, which is poorly used within RFC 4594 because other types of flows can be 'broadcast' other than video, such as audio. From this, [ID-4594-UPDATE] moved in two directions:

- it called out two service classes (audio and video), even though audio and video packets are not the only types of packets within each traffic characteristic.

- it removed "Video" from the Broadcast service class name.

From the resistance to this proposal within [ID-4594-UPDATE], perhaps other service class label names should be used.

The draft also recognizes the differences in video traffic, even though it is always carried over RTP [RFC3550]. Aside from silence suppression, video traffic varies far more than audio traffic. For example, video is

- far more variable in bandwidth utilization within the same flow.

- far more variable in packet size.

- at different business priorities in some networks based on a configuration. For example, desktop video often is of less important than Telepresence video on the same network. Lacking congestion, the two are treated the same. When congestion exists, one is given priority over the other.

Consequently any service class that contains video needs to account for larger packet size variation than audio, which was equally true in 2006, but not contained in RFC 4594.

Further, with the publication of RFC 5865, the concept of 'capacity admitted' traffic flows have been defined within DiffServ, and are being expanded with the proposal within this new draft [ID-NEW-DSCPS]. There are differing opinions as to whether the realtime Treatment Aggregate in Figure 1 above should also contain these capacity admitted flows, or if 'capacity admitted' traffic flows should have their own Treatment Aggregate containing all realtime capacity admitted traffic. Mixing capacity admitted traffic with unbounded realtime traffic seems to be trouble from a predictability point of view within routers believing they individually understand exactly how much traffic will be traversing each interface and at what rate.

All this said, there is a valid argument to constrain or prevent any DSCP value from being assigned to a single application, mostly due

to the limitation of the overall number of DSCP values available for use. [ID-4594-UPDATE] provides at least several applications per service class (or DSCP); a fact many have overlooked to date.

[ID-4594-UPDATE] is not only about or because of realtime traffic. It is also an overall update to the ideas and guidelines within RFC 4594, with the intent to make that document a standards track document for interoperability purposes.

4. Conclusion and What's Next

Without attempting to fundamentally change the guidelines within RFC 5127, this effort should not be as controversial as it has been, if we understand that those networks that need more granular traffic treatments can be configured with more granularity while not violating the needs of other networks that do not wish to be made aware of the increased treatment differences.

Everyone involved in this discussion needs to have a clear understanding of the difference points of view within the RFC 4594 effort (i.e., the RFC and the update draft) as well as within RFC 5127. One focuses on defining each service class and the other focuses on determining which of the existing service classes go into which aggregate, if present.

We hope to form a BoF on this subject that will explicitly *not* form a working group or produce any documents, or even drafts, but will gather the community from several (if not all) areas, and not just within the transport area. That is the purpose of this draft: to stimulate discussion towards the goal of discussion within the community on DiffServ. If the community does not believe a BoF is necessary, the work will proceed, or not, in TSVWG. Knowing how many within the community have attended TSVWG in each meeting for the last 9 or so years, it is felt that a much wider audience is necessary, given how much impact [ID-4594-UPDATE] can potentially have.

5. Acknowledgements

The author would like to thank Gorry Fairhurst and David Black for their positive discussions towards the formation of a BoF in Vancouver IETF. The author would also like to thank Paul Jones for doing a valuable proof read to catch points I didn't make clear, as well as identify simple nits I should have caught the nth time I reread this.

6. IANA Considerations

There are no IANA considerations as a result of this document.

7. Security Considerations

There are no security considerations within this document because it will not be progressed beyond this individual contributor stage, and all the specifying will be done in other drafts that will wholly contain all the security considerations for this goal/idea.

8. References

8.1 Normative References

There are no normative references within this document.

8.2. Informative References

[RFC2474] K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers ", RFC 2474, December 1998

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.

[RFC4594] J. Babiarz, K. Chan, F Baker, "Configuration Guidelines for Diffserv Service Classes", RFC 4594, August 2006

[RFC5127] Chan, K., Babiarz, J., and F. Baker, "Aggregation of DiffServ Service Classes", RFC 5127, February 2008.

[RFC5865] F. Baker, J. Polk, M. Dolly, "A Differentiated Services Code Point (DSCP) for Capacity-Admitted Traffic", RFC 5865, May 2010

[ID-4594-UPDATE] J. Polk, "Standard Configuration of DiffServ Service Classes", "work in progress", March 2012

[ID-NEW-DSCPS] J. Polk, "New Differentiated Services Code Point Assignments for Rich Media Traffic", "work in progress", March 2012

[ID-TCMTF] J. Saldana, D. Wing, J. Fernandez Navajas, Muthu A M. Perumal, J. Ruiz Mas, "Tunneling Compressed Multiplexed Traffic Flows (TCMTF)", "work in progress", March 2012

Authors' Address

James Polk
3913 **Treemont Circle**
Colleyville, Texas 76034

Phone: +1.817.271.3552
Email: jmpolk@cisco.com