Internet Engineering Task Force
INTERNET DRAFT
Expires May 1998

Juha Heinanen
Telia Finland, Inc.
November 1997

Use of the IPv4 TOS Octet to Support Differential Services

<draft-heinanen-diff-tos-octet-01.txt>

Status of this Memo

Abstract

This document describes how the TOS octet in the IPv4 header can be used to support differential Internet services. The proposal is based on the existing format of the TOS octet as defined in [RFC791] and [RFC1349].

1. Introduction

In order to support differential services within the Internet, such as those proposed by [Clark] and [Kilkki], the routers must be able to distinguish, how they should treat IP packets in terms of Type of Service (TOS) and Drop Preference (DP). The TOS usually indicates the desired delay characteristics of the packet, whereas the DP identifies how "important" the packet is if packets need to be discarded due to congestion.

In a typical router implementation, each TOS has its own queue, which it serves with a policy that meets the characteristics of the TOS. The DPs, on the other hand, map to queue thresholds so that an

incoming packet with a DP value i may be discarded if the length of the corresponding queue has exceeded a threshold value th(i) of that queue.

The TOS of a packet is set by the application or, if the application is unaware of differential Internet services, by the operating system or a router. In the latter two cases, the TOS may be determined from the TCP/UDP port numbers of the packet.

The DP of a packet can also be set by the application, the OS, or a router, but the criteria based on which the DP is assigned can vary widely. For example, a packet may be assigned a higher DP if it is consider to be "outside" of a given user profile. For more information on policies to set the DP, see [Clark] and [Kilkki].

2. TOS and DP in the IPv4 Header

Ideally, in order to support a wide range of TOSes and DPs, the IP packet header should have several bits available for this purpose. Luckily, the The IPv4 header has an 8 bit TOS octet, that can be used to carry the TOS and DP information.

According to [RFC791], the IPv4 TOS octet is divided into a 3 bit Precedence field and a 3 bit TOS field. The last two bits of the TOS octet are reserved for future use:

```
Bits 0-2:  Precedence.
Bit    3:  0 = Normal Delay,       1 = Low Delay.
Bits   4:  0 = Normal Throughput,  1 = High Throughput.
Bits   5:  0 = Normal Reliability, 1 = High Reliability.
Bit  6-7:  Reserved for Future Use.
```

The following two sections propose how the TOS and DP information needed to support the implementation of differential Internet services can be mapped to the IPv4 TOS octet.

3. Drop Preference (DP)

The semantics of the 3 bit Precedence field as defined in [RFC791] doesn't contradict its use as the DP field:

> "Several networks offer service precedence, which somehow treats high precedence traffic as more important than other traffic (generally by accepting only traffic above a certain precedence at time of high load)."

Further, [RFC791] assigns the following meaning to the eight possible values of Precedence field:

111 - Network Control
110 - Internetwork Control
101 - CRITIC/ECP
100 - Flash Override
011 - Flash
010 - Immediate
001 - Priority
000 - Routine

Consistently with the semantics of DP, the above values can be interpreted to describe the increasing importance of an IPv4 packet from the value 000 to value 111.

It is proposes that differential Internet services use the Precedence field of the IPv4 header to indicate the DP of the packet. In accordance with [RFC791], the DP of the packet can range from value 000, indicating the highest DP, to value 111, indicating the lowest DP. Note that if conformance with [RFC791] if not a requirement, it would be more natural to use the value 000 to indicate the lowest DP and 111 the highest DP, since value 000 is usually the default in IPv4 packets.

At present it is still unclear how many levels of DP will be actually needed. No matter what the conclusion will be, it is highly unlikely that every network would support all three Precedence bits and implement all eight levels of DP. This documemnt therefore proposes that a network is allowed to support one, two, or three DP bits. If it supports only one DP bit, it is the Bit 0, and if it supports two, they are the Bits 0 and 1.

4. Type of Service (TOS)

The three bit TOS field can be used to indicate the Type of Service that an IPv4 packet expects to receive from a network. As already shown above, [RFC791] specifies the following semantics for these 3 bits:

```
Bit    3:  0 = Normal Delay,       1 = Low Delay.
Bits   4:  0 = Normal Throughput,  1 = High Throughput.
Bits   5:  0 = Normal Reliability, 1 = High Reliability.
```

In [RFC1349], the TOS field has been extended by one of the reserved bits of [RFC791] resulting in the following semantics for the Bits 3-6 of the IPv4 TOS octet:

```
1000   --   minimize delay
0100   --   maximize throughput
0010   --   maximize reliability
```

```
0001   --   minimize monetary cost
0000   --   normal service
```

Ideally, the TOS field of IPv4 TOS octet should support a wide variety of delay characteristics. However, it is not known yet, how many different levels of delay can be actually supported by implementations so that each level would have its own distinguised "feel". It is therefore proposed that networks that support differential Internet services initially only use the Bit 3 of the TOS field to indicate the desired delay treatment of a packet. According to both [RFC791] and [RFC1349], by setting the Bit 3 to 0 or 1, the packet requests a normal or low delay, respectively.

This proposal doesn't exclude the possibility to later extend the delay "field" by another bit if it turns out that two levels of delay is not enough and that more levels can be supported by implementations. It is proposed that the Bit 4 of the TOS octet is reserved for the possible extension of the delay "field" until the issue is resolved. If two delay bits will be later defined, a network is still allowed to support only one of them, which is the Bit 3.

The use of the Bits 5-7 of the TOS octet in connection with differential Internet services is left for further study.

5. Summary

This document has proposed how the IPv4 TOS octet can be used to support differential services within the Internet. The proposal does not include any syntactical changes to the IPv4 header and also the semantic changes are kept to the very minimum.

The author hopes that this proposal could help to create a common convention on the usage of the TOS octet for the support of differential services within the Internet. Such services will be or, more precisely, are being implemented anyhow and it is thus very important that the implementations and service offerings will be compatible with each other.

6. Security Considerations

Security is not addressed in the current version of this document.

References

[Clark] Clark, D. and Wroclawski, J., "An Approach to Service Allocation in the Internet". draft-clark-diff-svc-alloc-00.txt, July 1997.

[Kilkki] Kilkki, K., "Simple Integrated Media Access (SIMA)". <draft-kalevi-simple-media-access-01.txt>, June 1997.

[RFC791] Information Sciences Institute, "Internet Protocol". RFC 791, September 1981.

[RFC1349] Almquist, P., "Type of Service in the Internet Protocol Suite". RFC 1349, July 1992.

Acknowledgements

I would like to thank Kalevi Kilkki for his comments on earlier versions of this document. The current version has been influenced by the discussions on the int-serv mailing lists.

Author Information

Juha Heinanen
Telia Finland, Inc.
Myyrmaentie 2
01600 VANTAA
Finland

Phone +358 303 944 808
Email jh@telia.fi