

Network Working Group
Request for Comments: 3140
Obsoletes: 2836
Category: Standards Track

D. Black
S. Brim
B. Carpenter
F. Le Faucheur
June 2001

Per Hop Behavior Identification Codes

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This document defines a 16 bit encoding mechanism for the identification of differentiated services Per Hop Behaviors in protocol messages. It replaces RFC 2836.

Table of Contents

1. Introduction.....	2
1.1. Usage Scenarios.....	2
2. Encoding.....	3
3. Signalling the Class Selector Codepoints.....	4
4. IANA Considerations.....	5
5. Security Considerations.....	5
Changes from RFC 2836.....	5
Acknowledgements.....	6
References.....	6
Authors' Addresses.....	6
Intellectual Property.....	7
Full Copyright Statement.....	8

1. Introduction

Differentiated Services [RFC 2474, RFC 2475] introduces the notion of Per Hop Behaviors (PHBs) that define how traffic belonging to a particular behavior aggregate is treated at an individual network node. In IP packet headers, PHBs are not indicated as such; instead Differentiated Services Codepoint (DSCP) values are used. There are only 64 possible DSCP values, but there is no such limit on the number of PHBs. In a given network domain, there is a locally defined mapping between DSCP values and PHBs. Standardized PHBs recommend a DSCP mapping, but network operators may choose alternative mappings.

In some cases it is necessary or desirable to identify a particular PHB in a protocol message, such as a message negotiating bandwidth management or path selection, especially when such messages pass between management domains. Examples where work is in progress include communication between bandwidth brokers, and MPLS support of diffserv.

In certain cases, what needs to be identified is not an individual PHB, but a set of PHBs. One example is a set of PHBs that must follow the same physical path to prevent re-ordering. An instance of this is the set of three PHBs belonging to a single Assured Forwarding class, such as the PHBs AF11, AF12 and AF13 [RFC 2597].

This document defines a binary encoding to uniquely identify PHBs and/or sets of PHBs in protocol messages. This encoding MUST be used when such identification is required.

This document replaces RFC 2836, which omitted considerations for the Class Selector codepoints.

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].

1.1. Usage Scenarios

Diffserv services are expected to be supported over various underlying technologies which we broadly refer to as "link layers" for the purpose of this discussion. For the transport of IP packets, some of these link layers make use of connections or logical connections where the forwarding behavior supported by each link layer device is a property of the connection. In particular, within the link layer domain, each link layer node will schedule traffic depending on which connection the traffic is transported in. Examples of such "link layers" include ATM and MPLS.

For efficient support of diffserv over these link layers, one model is for different Behavior Aggregates (BAs) (or sets of Behavior Aggregates) to be transported over different connections so that they are granted different (and appropriate) forwarding behaviors inside the link layer cloud. When those connections are dynamically established for the transport of diffserv traffic, it is very useful to communicate at connection establishment time what forwarding behavior(s) is (are) to be granted to each connection by the link layer device so that the BAs transported experience consistent forwarding behavior inside the link layer cloud. This can be achieved by including in the connection establishment signaling messages the encoding of the corresponding PHB, or set of PHBs, as defined in this document. Details on proposed usage of PHB encodings by some MPLS label distribution protocols (RSVP and LDP) for support of Diff-Serv over MPLS, can be found in [MPLS-DS].

In another approach, the ATM Forum has a requirement to indicate desired IP QOS treatments in ATM signaling, so that ATM switches can be just as supportive of the desired service as are IP forwarders. To do so the Forum is defining a new VC call setup information element is which will carry PHB identification codes (although will be generalized to do more if needed).

2. Encoding


PHBs and sets of PHBs are encoded in an unsigned 16 bit binary field.

The 16 bit field is arranged as follows:


Case 1: PHBs defined by standards action, as per [RFC 2474].

The encoding for a single PHB is the recommended DSCP value for that PHB, left-justified in the 16 bit field, with bits 6 through 15 set to zero. Note that the recommended DSCP value MUST be used, even if the network in question has chosen a different mapping.

The encoding for a set of PHBs is the numerically smallest of the set of encodings for the various PHBs in the set, with bit 14 set to 1. (Thus for the AF1x PHBs, the encoding is that of the AF11 PHB, with bit 14 set to 1.)



Case 2: PHBs not defined by standards action, i.e., experimental or local use PHBs as allowed by [RFC 2474]. In this case an arbitrary 12 bit PHB identification code, assigned by the IANA, is placed left-justified in the 16 bit field. Bit 15 is set to 1, and bit 14 is zero for a single PHB or 1 for a set of PHBs. Bits 12 and 13 are zero.



Bits 12 and 13 are reserved either for expansion of the PHB identification code, or for other use, at some point in the future.

In both cases, when a single PHBID is used to identify a set of PHBs (i.e., bit 14 is set to 1), that set of PHBs MUST constitute a PHB Scheduling Class (i.e., use of PHBs from the set MUST NOT cause intra-microflow traffic reordering when different PHBs from the set are applied to traffic in the same microflow). The set of AF1x PHBs [RFC 2597] is an example of a PHB Scheduling Class. Sets of PHBs that do not constitute a PHB Scheduling Class can be identified by using more than one PHBID.

3. Signalling the Class Selector Codepoints

[RFC 2474] defines the eight DS codepoint values of the form 'xxx000' (where x may be '0' or '1') as the Class Selector Codepoints. Codepoint 000000 is the recommended DSCP value for the Default PHB, and hence the Case 1 PHBID constructed from that codepoint is used to signal the Default PHB (see Section 2 above).

For convenience and consistent operation with networks that employ IP Precedence [RFC 1812], the Case 1 format PHBIDs constructed from the other seven Class Selector Codepoints may also be used to signal PHBs. In each case, the PHB signaled by such a PHBID is the PHB to which the embedded class selector codepoint (or IP Precedence value that corresponds to it in non-diffserv domains) is mapped in the recipient's network. Note that different networks will employ different mappings; see Section 4 of [RFC 2474] for further discussion.

Any specified use of PHBIDs SHOULD allow the use of the eight Case 1 PHBIDs constructed from the Class Selector Codepoints.

4. IANA Considerations

IANA is requested to create a new assignment registry for "Per-Hop Behavior Identification Codes", initially allowing values in the range 0 to 4095 decimal.

Assignment of values in this field require:

- the identity of the assignee
- a brief description of the new PHB, with enough detail to distinguish it from existing standardized and non-standardized PHBs. In the case of a set of PHBs, this description should cover all PHBs in the set.
- a reference to a stable document describing the PHB in detail.

During the first year of existence of this registry, IANA is requested to refer all requests to the IETF diffserv WG for review. Subsequently, requests should be reviewed by the IETF Transport Area Directors or by an expert that they designate.

If the number of assignments begins to approach 4096, the Transport Area Directors should be alerted.

5. Security Considerations

This encoding in itself raises no security issues. However, users of this encoding should consider that modifying a PHB identification code may constitute theft or denial of service, so protocols using this encoding must be adequately protected.

Just signalling a PHBID SHOULD NOT be sufficient to grant the sender access to a PHB that it would otherwise not be able to use. In cases where this is an issue, receivers SHOULD treat received PHBIDs as requests for service, and use local policy to determine whether to grant or deny such requests.

Changes from RFC 2836

[RFC 2836] did not consider the Class Selector code points, which are covered by section 3 of the present document. A clarification has been added at the end of section 2 for the case of PHB Scheduling Classes. The second paragraph of section 5 has been added.

Acknowledgements

Useful comments were made by members of the IETF Diffserv working group.

References

- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC 2474] Nichols, K., Blake, S., Baker, F. and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC 2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [RFC 2597] Heinanen, J., Baker, F., Weiss, W. and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, June 1999.
- [RFC 2836] Brim, S., Carpenter, B. and F. Le Faucheur, "Per Hop Behavior Identification Codes", RFC 2836, May 2000.
- [MPLS-DS] Le Faucheur, F., et al., "MPLS Support of Differentiated Services", Work in Progress.

Authors' Addresses

David L. Black
EMC Corporation
42 South St.
Hopkinton, MA 01748

EMail: black_david@emc.com

Scott W. Brim
146 Honness Lane
Ithaca, NY 14850
USA

EMail: sbrim@cisco.com

Brian E. Carpenter
IBM
c/o iCAIR
Suite 150
1890 Maple Avenue
Evanston, IL 60201
USA

E-Mail: brian@icair.org

Francois Le Faucheur
Cisco Systems
Petra B - Les Lucioles
291, rue Albert Caquot
06560 Valbonne
France

E-Mail: flefauch@cisco.com

Intellectual Property

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

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