

TSVWG
Internet-Draft
Intended status: Informational
Expires: August 18, 2014

R. Geib, Ed.
Deutsche Telekom
February 14, 2014

DiffServ interconnection classes and practice
draft-geib-tsvwg-diffserv-intercon-05

Abstract

This document proposes a limited and well defined set of QoS PHBs and PHB groups to be applied at (inter)connections of two separately administered and operated networks. Many network providers operate Aggregated DiffServ classes. This draft contains DiffServ aggregation friendly interconnection concepts.

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 August 18, 2014.

Copyright Notice

Copyright (c) 2014 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.

Internet-Draft

Abbreviated Title

February 2014

Table of Contents

1.	Introduction	3
1.1.	Related work	5
2.	Terminology	5
3.	Aggregating PHBs of a class by a DSCP Precedence Prefix . . .	6
4.	An Interconnection class and codepoint scheme	6
4.1.	Treatment of Network Control traffic at carrier interconnection interfaces	9
5.	DiffServ Intercon relation to other QoS standards	10
5.1.	MPLS, Ethernet and DSCP Precedence Prefixes for aggregated classes	11
5.2.	Proposed GSMA IR.34 to DiffServ Intercon mapping	11
5.3.	Proposed MEF 23.1 to DiffServ Intercon mapping	12
6.	Contributors	14
7.	Acknowledgements	14
8.	IANA Considerations	14
9.	Security Considerations	14
10.	References	14
10.1.	Normative References	14
10.2.	Informative References	15
Appendix A.	Change log	16
	Author's Address	16

Internet-Draft

Abbreviated Title

February 2014

1. Introduction

DiffServ has been deployed in many networks. As described by [section 2.3.4.2 of RFC 2475](#), remarking of packets at domain boundaries is a DiffServ feature [[RFC2475](#)]. This draft proposes a set of standard QoS classes and code points at interconnection points to which and from which locally used classes and code points should be mapped.

IP precedence has been deprecated. MPLS and Ethernet support 3 bit code point fields to differentiate service quality (see MPLS TC / Traffic Class [[RFC5462](#)] and PCP, Priority Code Point [[IEEE802.1Q](#)]). The concept of classifying DiffServ traffic classes by the bits 0-2 of a DSCP has been part of Diffserv from start on. This is also reflected by the DiffServ codepoint definitions of AF and EF. It is common practice today also to copy these three DSCP bits into MPLS TC or Ethernet P-Bits. PHBs based on DSCP bit 0-2 classification may be applied in core network sections rather than then DSCP based PHBs. Network providers make use of this feature for their own IP QoS concepts. This draft suggests to expand it to interconnections between operators of different domains in a simple manner while each operator may maintain the own class and codepoint scheme within the own domain.

The scope of this draft is limited to 4 specified interconnection classes having four different 3 bit code points in DSCP bits 0-2. Using more than the 4 proposed IP precedences at interconnection could result in non-revertible IP Precedence or DSCP rewrites and avoid sustaining end-to-end QoS classes, if a receiving provider operates more than 4 MPLS TCs. Assume a provider operating 4 QoS classes available at interconnection and MPLS within his backbone. Further assume this carrier to support MPLS based ECN marking and assume this carrier to operate a newtork control class with an own MPLS TC. Two codepoints are left for future use. If 5 or more PHBs each with different DSCP bits 0-2 are offerd at an interconnection point and no more than a single MPLS label needs to be pushed, two (or more) PHBs will carry the same DSCP bits 0-2 after re-marking to

the provider internal QoS scheme. Due to MPLS pen ultimate hop popping, DSCPs must be re-written in this case. That may work if bits 3-5 of the DSCP can be varied without introducing ambiguities. Should this traffic later pass another QoS interconnection point further downstream, the original sending domain may not be able to ensure proper class mapping for the PHBs merged into a single class by the receiving domain.

At first glance, the interconnection codepoint scheme looks like an additional effort. But there are some obvious benefits: each party sending or receiving traffic has to specify the mapping from or to the interconnection class and code point scheme only once. Without

it, this is to be negotiated per interconnection party individually. Further, end-to-end QoS in terms of traffic being classified for the same class in all passed domains is more likely to result if an interconnection code point scheme is used. It is not necessarily resulting from individual per provider mapping negotiations, as can be seen from the example given above.

The standards and deployments known to the author of this draft are limited to 4 DiffServ classes at interconnection points (or less). The example given in [RFC 5127](#) on aggregation of DiffServ service classes picks 4 Treatment aggregates (note that this document prefers class instead of treatment aggregate). Reasons to favour working with 4 DiffServ interconnection classes:

- o There should be a coding reserve for interconnection classes. This leaves space for future standards, for private bilateral agreements and for provider internal classes.
- o The fields available for carrying QoS information (e.g., DiffServ PHB) in MPLS and Ethernet are only 3 bits in size, and are intended for more than just QoS purposes (see e.g. [\[RFC5129\]](#)).
- o Migrations from one code point scheme to another may require spare QoS code points.

IP Precedence has been deprecated when DiffServ was standardised. It is common practice today however to copy the DSCPs Bits 0-2 (called DSCP Precedence Prefix in the following) into MPLS TC or Ethernet P-Bits. This is also reflected by the DiffServ codepoint definitions

of AF and EF. Class based PHBs may be applied in core network sections rather than then DSCP based PHBs.

The set of available router and traffic management tools to configure and operate DiffServ classes is limited. This should be reflected by class definitions. These may in the end be more related to transport properties (e.g., whether the traffic in a class is congestion controlled or not) than to application requirements.

[RFC5127](#) provides recommendations on domain internal aggregation of DiffServ traffic and offers a deployment example [[RFC5127](#)]. This draft differs from the [RFC5127](#) aggregation deployment example in the following points:

- o the basic concept of this draft is to maintain classes, while expecting DSCP remarking at provider edges.
- o This draft follows [RFC4594](#) in the proposed marking of provider Network Control traffic and expands [RFC4594](#) on treatment of CS6

marked traffic at interconnection points (see [section 5.2](#)).

The proposed Interconnection class and code point scheme tries to reflect and consolidate related DiffServ and QoS standardisation efforts outside of the IETF, namely MEF [MEF 23.1], GSMA [[IR.34](#)] and ITU [[Y.1566](#)]. GSMAs IR.34 specifies an inter provider VPN, but it is nevertheless specifying a kind of DiffServ aware IP based carrier interconnection.

[1.1](#). Related work

In addition to the standardisation activities which triggered this work, other authors published RFCs or drafts which may benefit from an interconnection class- and codepoint scheme. [RFC 5160](#) suggests Meta-QoS- Classes to enable deployment of standardised end to end QoS classes [[RFC5160](#)]. The authors agree that the proposed interconnection class- and codepoint scheme as well as the idea of standardised end to end classes would complement their own work. Work on signaling Class of Service at interconnection interfaces by BGP [[I-D.knoll-idr-cos-interconnect](#)], [[ID.idr-sla](#)] is beyond the scope of this draft. Should the basic transport and class properties be standardised as proposed here, signaled access to QoS classes may

be of interest. The current BGP drafts focus on exchanging SLA and traffic conditioning parameters. They seem to assume that common interpretation of the PHB properties identified by DSCPs has been established prior to exchanging further details by BGP signaling.

[2.](#) Terminology

This draft re-uses existing terminology.

DSCP Precedence Prefix Bits 0-2 of the DSCP ("x" in this generic DSCP: xxxddd) are called the DSCP Precedence Prefix. [Section 4.2 of \[RFC2474\]](#) discusses the role of these bits in enabling use of DiffServ with network equipment that is not fully DiffServ-compliant; this term provides a formal for these bits that is preferable to referring to them as "the former IP Precedence field".

DSCP Precedence Class This is a set of one or more PHBs that utilize the same DSCP Precedence Prefix on an interconnection between two networks.

[3.](#) Aggregating PHBs of a class by a DSCP Precedence Prefix

Configuration and operation of MPLS networks is simplified, if a DSCP Precedence Class can be aggregated into a single PSC by classifying them on their DSCP Precedence Prefix. The DSCP Precedence Prefix of an interconnection DSCP Precedence Class may be rewritten at the ingress edge router and then simply be copied into the MPLS TC field of one or more labels to be pushed.

To allow for simple carrier interconnection agreements, carriers sending traffic belonging to the same class but marked by DSCPs with differing DSCP Precedence Prefixes should apply the interconnection marking and code point scheme specified below, if they interconnect to a carrier applying DSCP Precedence Prefix based traffic aggregation. An example where this may be applicable is the

Interactive Class of GSMA IR.34 [[IR.34](#)]). Another option is to negotiate a customised interconnection agreement of course.

Classification by DSCP Precedence Prefix is useful for links aggregating DiffServ traffic. DSCP Precedence Prefix based classification is not recommended as a general mode of operation. Edge systems, QoS policy enforcement nodes, service areas and hosts benefit from fine grained DSCP based classification and should continue to do so.

[RFC 2474](#) specifies the Class Selector Codepoints [[RFC2474](#)]. These offer a similar concept, but they are strictly limited to xxx000 DSCPs. The Class Selector Code points don't offer aggregation, they just simplify classification. This draft intends to aggregate several PHBs of a single class by a DSCP Precedence Prefix, which a different concept than that of the Class Selector Code points.

[4.](#) An Interconnection class and codepoint scheme

Interconnecting parties face the problem of matching classes to be interconnected and then to agree on code point mapping. As stated by the DiffServ Architecture [[RFC2475](#)], remarking is a standard behaviour at interconnection interfaces. This draft proposes a standard interconnection set of 4 DSCP precedence classes with well defined DSCP and DSCP Precedence Prefix values. A sending party remarks DSCPs from internal schemes to the Interconnection code points. The receiving party remarks DSCP Precedence Prefixes and / or DSCPs to her internal scheme. Thus the interconnection code point scheme fully complies with the DiffServ architecture.

This draft picks up the DiffServ interconnection class definitions proposed by ITU-T Y.1566 [[Y.1566](#)]. In addition to the classes

defined there, this draft proposes a complete set of PHBs and DSCPs. Like in the example given by [RFC 5127](#) for domain internal class aggregation, Y.1566 specifies four PHB scheduling classes (for provider interconnection however). Their properties slightly from those of the [RFC5127](#) example:

Class Priority: PHB EF, DSCP 101 110. The figures of merit describing the PHB to be in the range of low single digit

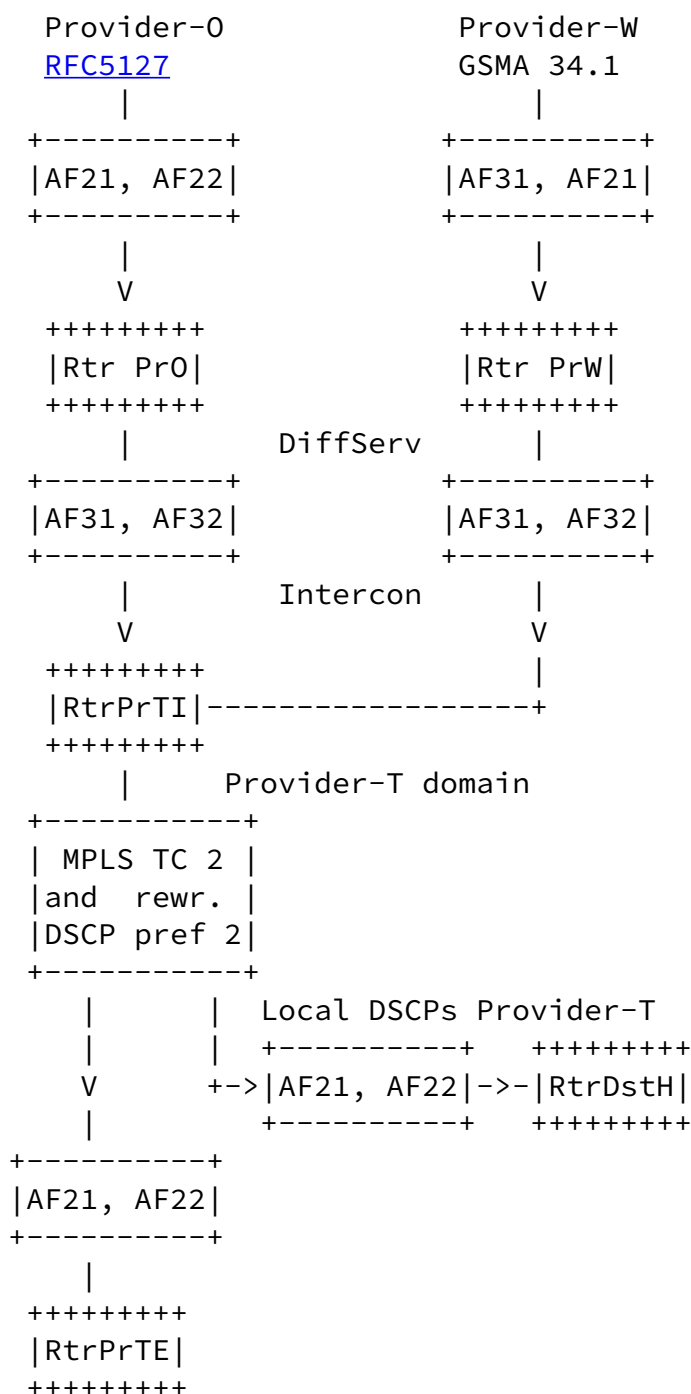
milliseconds. See [[RFC3246](#)]. This class corresponds to [RFC 5127](#)'s real time class, but it is limited to traffic for which node configuration "ensures that the service rate of EF packets on a given output interface exceeds their arrival rate at that interface over long and short time intervals" (see [RFC 3246](#)).


Bulk inelastic: PHB AF41, DSCP 100 010 (the other AF4 PHB group PHB's and DSCPs should be reserved for future extension of this DSCP scheduling class). Optimised for low loss, low delay, low jitter at high bandwidth. Traffic load in this class must be controlled, e.g. by application servers. One example could be flow admission control. There may be infrequent retransmissions requested by the application layer to mitigate low levels of packet losses. Discard of packets through active queue management should be avoided in this class. Congestion in this class may result in bursty packet loss. If used to carry multimedia traffic, it is recommended to carry audio and video traffic in a single PHB (note that video conferencing may require separate PHBs for audio and video traffic, which could also be realised by utilising two AF 4 PHBs). All of these properties influence the buffer design. This class is designed to transport those parts of [RFC 5127](#)'s Real Time class, which consume considerable QoS bandwidth at the interconnection interface.

Assured: The complete PHB group AF3, DSCPs 011 010, 011 100 and 011 110. This class may be optimised to transport traffic without bandwidth requirements. It aims on very low loss at high bandwidths. Retransmissions after losses characterise the class and influence the buffer design. Active queue management with probabilistic dropping may be deployed. The [RFC 5127](#) example calls this class Assured Elastic.

Default: Default PHB, CS0 with DSCP 000 000. This class may be optimised to transport traffic without bandwidth requirements. Retransmissions after losses characterise the class and influence the buffer design. Active queue management with probabilistic dropping may be deployed. The [RFC 5127](#) example calls this class Elastic.

peer with provider T. They have agreed upon a QoS interconnection. Traffic of provider 0 terminates within provider Ts network, while the GSMA IR.34 traffic transits through the network of provider T to provider F. Assume all providers to run their own internal codepoint schemes for a class with properties of the DiffServ Intercon Assured class. See section for a description of GSMA IR.34.





DiffServ Intercon example

Figure 1

It is easily visible that all providers only need to deploy internal DSCP to DiffServ Intercon DSCP mappings to exchange traffic in the desired classes.

[RFC5127](#) specifies a separate PHB scheduling class for network control traffic. This class may be present at interconnection interfaces too, but depending on the agreement between providers, it may also be classified for another interconnection class. See [section 4.2](#) for a detailed discussion.

The proposed class and code point scheme is designed for point to point IP layer interconnections. Other types of interconnections are out of scope of this document. The basic class and code point scheme is applicable on Ethernet layer too.

[4.1](#). Treatment of Network Control traffic at carrier interconnection interfaces

As specified by [RFC4594, section 3.2](#), Network Control (NC) traffic marked by CS6 is to be expected at interconnection interfaces. This document does not change NC specifications of [RFC4594](#). The latter specification is detailed on domain internal NC traffic and on traffic exchanged between peering points. Further, it recommends not to forward CS6 marked traffic originating from user-controlled end

points by the NC class of a provider domain.

Internet-Draft

Abbreviated Title

February 2014

As a minor clarification to [RFC4594](#), "peering" shouldn't be interpreted in a commercial sense. The NC PHB is applicable also in the case of a purchased network service based on a transit agreement with an upstream provider. [RFC4594](#) recommendations on NC traffic are applicable for IP carrier interconnections in general.

Some CS6 traffic exchanged accross carrier interconnections will terminate at the domain ingress node (e.g., if BGP is running between the two routers on opposite ends of the interconnection link).

An IP carrier may limit access to the NC PHB for traffic which is recognised as network control traffic relevant to the own domain. Interconnecting carriers should specify treatment of CS6 marked traffic received at a carrier interconnection which is to be forwarded beyond the ingress node. An SLA covering the following cases is recommended, if a carrier wishes to send CS6 marked traffic accross an interconnection link which isn't terminating at the interconnected ingress node:

- o classification of traffic which is network control traffic for both domains. This traffic should be classified and marked for the NC PHB.
- o classification of traffic which is network control traffic for the sending domain only. This traffic should be classified for a PHB offering similar properties as the NC class (e.g. AF31 as specified by this document). As an example GSMA IR.34 proposes an Interactive class / AF31 to carry SIP and DIAMETER traffic. While this is service control traffic of high importance to the interconnected Mobile Network Operators, it is certainly no Network Control traffic for a fixed network providing transit. The example may not be perfect. It was picked nevertheless because it refers to an existing standard.
- o any other CS6 marked traffic should be remarked or dropped.

[5.](#) DiffServ Intercon relation to other QoS standards

This sections provides suggestions how to aggregate traffic by DSCP Precedence Prefexies to Ethernet and MPLS. Other Standardisation Organsiations deal with QoS aware provider interconnection. As IP is the state of the art realisation of provider interconnections, these standards bodies specify DiffServ aware interconnections. Some of these bodies are industry alliances chartered also to promote interconnection of wireless or Ethernet technology including the exchange of IP datagrams at interconnection points. Examples are the Metro Ethernet Forum (MEF) or the GSM Alliance (GSMA). The ITU was

mentioned earlier. ITU works across and between responsibilities of different Standardisation Organisations and liaising with them, if their responsibilities are touched, is traditional part of that process.

[5.1.](#) MPLS, Ethernet and DSCP Precedence Prefixes for aggregated classes

The interconnection class and code point scheme respects properties and limits of a 3 bit PHB coding space in different ways:

- o it allows to classify four interconnection classes based on the DSCP Precedence Prefix.
- o it supports a single PHB group (AF3), whose DSCPs may be aggregated into a sinle MPLS TC (or Ethernet priority) based on their joint DSCP Precedence Prefix. This kind of aggregation will work for the AF4 PHB group, if the PHBs AF42 and AF43 need to be supported in addition to AF41.
- o Applying only 4 aggregated classes at interconnection allows for bilateral extensions, if desired. Should two carriers agree to map AF32 and AF33 to an additional individual MPLS TC and offer an Ordered Aggregate across the interconnecting domain, this proposal at leaves some MPLS TC coding space for such an extension (although this draft doesn't recommend interconnections of that type).

The above statement is no requirement to depricate any DSCP to MPLS TC or Ethernet P-Bit mapping functionality. In the opposite, by limiting the interconnection scheme to 6 PHBs, each PHB may be mapped to an individual Traffic Class or Priority also within MPLS or Ethernet (if desired).

5.2. Proposed GSMA IR.34 to DiffServ Intercon mapping

GSMA IR.34 provides guidelines on how to set up and interconnect Internet Protocol (IP) Packet eXchange (IPX) Networks [IR.34]. An IPX network is an inter-Service Provider IP backbone which comprises the interconnected networks of IPX Providers. IPX is a telecommunications interconnection model for the exchange of IP based traffic between customers of separate mobile and fixed operators as well as other types of service provider (such as ISP), via IP based network-to-network interface. Note that IPX is not a public interconnection model however, it is designed as a private IP Backbone of the interconnected parties. Two IPX partners may interconnect using transit offered by an Inter-Service Provider IP Backbone. This requires an IP QoS aware interconnection as described by this draft between IPX provider and Inter-Service Provider IP

Geib

Expires August 18, 2014

[Page 11]

Internet-Draft

Abbreviated Title

February 2014

Backbone.

GSMA IR.34 specifies 4 aggregated classes and 7 PHBs. Mapping to DiffServ Intercon is smooth apart from the GSMA aggregated class Interactive, which consists of 4 PHBs. The table below lists a suggested mapping to DiffServ Intercon.

GSMA IR.34		DiffServ-Intercon	
Agg. Class	PHB	Agg. Class	PHB
Conversational	EF	Priority	EF
Streaming	AF41	Bulk inelastic	AF41
Interactive	AF31	Assured	AF31
(ordered by	AF32		
priority,			AF32
AF3 highest)	AF21		
	AF11		AF33

Background	CS0	Default	CS0	
+-----+	+-----+	+-----+	+-----+	+-----+

Suggested mapping of GSMA IR.34 classes and PHBs to DiffServ Intercon

Figure 2

The motivation resulting in the design of the IR.34 Intercative class are unknown to the author of this draft. It is out of scope of this draft to decide how 4 PHBs can be mapped to a to single aggregated class. The suggested mapping is pragmatic and tries to come as close as possible to other design criteria pursued by GSMA IR.34.

[5.3.](#) Proposed MEF 23.1 to DiffServ Intercon mapping

MEF 23.1 is an implemenation agreement facilitating Ethernet service interoperability and consistency between Operators and the use of a common CoS Label set for Subscribers [[MEF23.1](#)]. It pursues the same aims and method on Ethernet layer as this draft does on IP layer (i.e. providing an interconnection class and codepoint scheme). MEF 23.1 addresses external network to network interfaces typically interconnecting metro ethernet providers. This may typically involve

Geib

Expires August 18, 2014

[Page 12]

Internet-Draft

Abbreviated Title

February 2014

Ethernet Network Sections associated with typical Operator domains that interconnect at external network to network interfaces. MEF 23.1 specifies three aggregated CoS classes. In addition, the usage of a subset of Ethernet PCP and IP DSCP values is specifiedthus facilitating synergies between Ethernet and IP services and networks. The main purpose is specifying operator virtual (Ethernet) connections. As an IP QoS model is added, this draft proposes an IP class and codepoint mapping.

MEF 23.1 QoS mapping requires some thought as 3 classes are supported of which two are Ordered Aggregates. MEF 23.1s [section 8.5.1](#) example on IP DSCP mapping may suggest supporting classification based on the DSCP Precedence Prefix. MEF 23.1 [section 8.5.2.1](#) allows the conclusion that MEF class M is best mapped to this drafts Bulk inelastic class. The suggested mapping MEF to DiffServ Intercon is limited to the the MEF color blind mode (3 classes, no sub-classes):

MEF 23.1		DiffServ-Intercon	
Agg. Class	PHB	Agg. Class	PHB
High	EF	Priority	EF
Medium	AF3	Bulk inelastic	AF41
Low	CS1	Default	CS0

Suggested mapping of MEF 23.1 color blind mode classes and PHBs to DiffServ Intercon

Figure 3

The MEF color aware mode supports Ordered Aggregates in the MEF classes M and L. The MEF L specification classifies AF1 and Best Effort for the same Ordered Aggregate. A Better than Best Effort service produced in the same queue as best effort traffic can be realized, but hasn't been standardized by IETF. This document doesn't suggest any mapping. Diffserv Intercon also doesn't suggest an Ordered Aggregate in the Bulk Inelastic class. Later versions of this document may do so and specify a solution. Both issues are left for bilateral negotiation.

6. Contributors

David Black provided many helpful comments and inputs to this work.

7. Acknowledgements

Al Morton and Sebastien Jobert provided feedback on many aspects during private discussions. Brian Carpenter, Mohamed Boucadair and Thomas Knoll helped adding awareness of related work.

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

This document does not introduce new features, it describes how to use existing ones. The security section of [RFC 2475](#) [[RFC2475](#)] and [RFC 4594](#) [[RFC4594](#)] apply.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2474] 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.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", [RFC 2475](#), December 1998.
- [RFC2597] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, "Assured Forwarding PHB Group", [RFC 2597](#), June 1999.
- [RFC3246] Davie, B., Charny, A., Bennet, J., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", [RFC 3246](#), March 2002.

- [RFC3260] Grossman, D., "New Terminology and Clarifications for Diffserv", [RFC 3260](#), April 2002.
- [RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen,

P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", [RFC 3270](#), May 2002.

[RFC5129] Davie, B., Briscoe, B., and J. Tay, "Explicit Congestion Marking in MPLS", [RFC 5129](#), January 2008.

[RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", [RFC 5462](#), February 2009.

[min_ref] authSurName, authInitials., "Minimal Reference", 2006.

[10.2](#). Informative References

[I-D.knoll-idr-cos-interconnect]
Knoll, T., "BGP Class of Service Interconnection",
[draft-knoll-idr-cos-interconnect-11](#) (work in progress),
November 2013.

[ID.idr-sla]
IETF, "Inter-domain SLA Exchange", IETF, [http://
datatracker.ietf.org/doc/draft-ietf-idr-sla-exchange/](http://datatracker.ietf.org/doc/draft-ietf-idr-sla-exchange/),
2013.

[IEEE802.1Q]
IEEE, "IEEE Standard for Local and Metropolitan Area
Networks - Virtual Bridged Local Area Networks", 2005.

[IR.34] GSMA Association, "IR.34 Inter-Service Provider IP
Backbone Guidelines Version 7.0", GSMA, GSMA IR.34 [http://
www.gsma.com/newsroom/wp-content/uploads/2012/03/
ir.34.pdf](http://www.gsma.com/newsroom/wp-content/uploads/2012/03/ir.34.pdf), 2012.

[MEF23.1] MEF, "Implementation Agreement MEF 23.1 Carrier Ethernet
Class of Service Phase 2", MEF, MEF23.1 [http://
metroethernetforum.org/PDF_Documents/
technical-specifications/MEF_23.1.pdf](http://metroethernetforum.org/PDF_Documents/technical-specifications/MEF_23.1.pdf), 2012.

[RFC4594] Babiarz, J., Chan, K., and 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.

[RFC5160] Levis, P. and M. Boucadair, "Considerations of Provider-to-Provider Agreements for Internet-Scale Quality of Service (QoS)", [RFC 5160](#), March 2008.

[Y.1566] ITU-T, "Quality of service mapping and interconnection between Ethernet, IP and multiprotocol label switching networks", ITU, <http://www.itu.int/rec/T-REC-Y.1566-201207-I/en>, 2012.

[Appendix A](#). Change log

00 to 01 Added terminology and references. Added details and information to interconnection class and codepoint scheme. Editorial changes.

01 to 02 Added some references regarding related work. Clarified class definitions. Further editorial improvements.

02 to 03 Consistent terminology. Discussion of Network Management PHB at interconnection interfaces. Editorial review.

03 to 04 Again improved terminology. Better wording of Network Control PHB at interconnection interfaces.

Author's Address

Ruediger Geib (editor)
Deutsche Telekom
Heinrich Hertz Str. 3-7
Darmstadt, 64295
Germany

Phone: +49 6151 5812747
Email: Ruediger.Geib@telekom.de

