

Network Working Group
Internet-Draft
Expires: January 1, 2006

P. Levis
M. Boucadair
France Telecom
June 30, 2005

The Meta-QoS-Class concept
draft-levis-meta-qos-class-00.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 1, 2006.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document describes a framework for inter-domain Quality of Service (QoS). It makes use of a new concept denoted by Meta-QoS-Class. This new concept guides and simplifies QoS agreements between providers and opens up new perspectives for a global QoS-enabled Internet.

Internet-Draft

The Meta-QoS-Class concept

June 2005

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Assumptions	4
4.	Binding l-QCs as a fundamental inter-domain QoS process . .	5
5.	Provider to provider agreements analysis	5
5.1	About traps and glaciation	5
5.2	Recommendations for provider agreements	6
5.3	Edge to edge guarantees	7
5.4	The need for MQC	7
6.	The Meta QoS Class concept	8
6.1	Definition	8
6.2	Template	8
6.3	Binding process	8
6.4	Properties	9
6.5	Common understanding	10
7.	New perspectives for a gobal QoS-enabled Internet	10
8.	IANA Considerations	11
9.	Security Considerations	11
10.	Acknowledgements	12
11.	References	12
11.1	Normative References	12
11.2	Informative References	12
	Authors' Addresses	13
	Intellectual Property and Copyright Statements	14

1. Introduction

Inter-domain Quality of Service (QoS) is seen as one of the future challenges for the Internet community [[RFC3869](#)]. At a first stage, most of the effort has been put into intra-domain specific problems. A number of issues still appear to need further elaboration [[RFC2990](#)], before it is conceivable to have an operational deployment of QoS services including inter-domain aspects.

As pointed in [[WIP](#)], concepts should always be created in relation to specific problems. This memo focuses on inter-domain QoS and proposes some mechanisms to ease extending intra-domain QoS capabilities. It introduces a new concept denoted by Meta-QoS-Class (MQC). This concept is based on a very simple assumption: for two adjacent domains, each specifically engineered to convey traffic for a given application, the quality of the transportation across the two networks as a whole is likely to be satisfactory for the application, even if the two networks are managed by two different entities. Basically this is how the Plain Old Telephone Service (POTS) works. Finally, this memo explains how MQC could be used to promote a global QoS-enabled Internet. Note that this document doesn't specify any protocols or systems.

2. Terminology

Service Provider

An entity that provides Internet connectivity. Sometimes simply referred to as provider or SP. This document assumes that a provider owns and administers an IP network called a domain.

Domain

A network infrastructure composed of one or several Autonomous Systems.

Local-QoS-Class (l-QC)

A QoS transfer capability across a single domain, characterized by a set of QoS performance parameters denoted by (D, J, L). From a Diffserv [\[RFC2475\]](#) perspective, an l-QC is an occurrence of a Per Domain Behavior (PDB) [\[RFC3086\]](#). It is then signaled by one or several Differentiated Services Code Point (DSCP) [\[RFC2474\]](#).

(D, J, L)

D: one-way transit delay [\[RFC2679\]](#), J: one-way transit delay variation or jitter [\[RFC3393\]](#), and L: packet loss rate [\[RFC2680\]](#).

Inter-domain QoS

Refers to the level of network QoS guarantees for communications that span several domains.

Service scope

Network boundaries demarcating the guarantees of a service.

[3.](#) Assumptions

This memo considers the Internet subset composed of all domains with Diffserv-like capabilities. These capabilities can differ from one provider to another by the number of deployed l-QCs, by their respective QoS characteristics, and also by the way they have been implemented and engineered. This memo does not put any specific requirements on the intra-domain traffic engineering policies and the way they are enforced.

When crossing a domain, a traffic requesting a particular QoS treatment experiences conditions constrained by the values of the (D, J, L) tuple corresponding to the l-QC applied by the provider.

A provider negotiates QoS agreements only with its BGP peers [\[RFC1771\]](#). Potentially, all SPs that are directly accessible without the need to cross a third party domain (immediate neighbors). This is referenced as the cascaded model, see (Figure 1). The opposite approach is the centralized model, see (Figure 2), where agreements

can be reached directly with any remote providers.




Figure 1: Cascaded model

There is a great deal of complexity and scalability issues related to the centralized approach, which represents a radical shift from current Internet business model. Therefore, the only realistic way forward seems to be the cascaded approach. This is the approach adopted in the remainder of this document.




Figure 2: Centralized model

[4.](#) Binding l-QCs as a fundamental inter-domain QoS process

Inter-domain QoS can be approached from many points of view. For providers it is a matter of extending their service scopes, beyond the boundaries of their own domains. They want to benefit from the Internet QoS infrastructure formed by all QoS-enabled domains. On a low level (with regard to protocol layering) extending the service scopes means extending the l-QCs. Packets leaving a domain that applied a given l-QC (signaled by a given DSCP if Diffserv is used), should experience similar treatment when crossing external domains up

to their final destination.

By definition, two l-QCs from two neighboring domains are bound together once the two providers have agreed to transfer traffic from one l-QC to the other. Binding l-QC appears as a fundamental process, it should ensure the consistency of inter-domain QoS treatments.

A provider needs to find the best match between its deployed l-QCs and the neighbor's l-QCs. One potential difficulty is that providers are, in general, very reluctant to communicate on how their networks are engineered. They consider this kind of information as private and confidential. L-Qc binding should operate with minimal exchange of information.


[5.](#) Provider to provider agreements analysis

[5.1](#) About traps and glaciation

In order to identify relevant criteria to bind adjacent l-QCs, this section analyzes provider to provider agreements based on chains of domains.

Provider SP_n offers IP connectivity services to its customers that are part of its BGP neighbors. An IP connectivity service is a set of (Destination, D, J, L) where Destination is a group of IP

addresses reachable via SP_n, and (D, J, L) is the QoS performance to get from SP_n to Destination. SP_n guarantees the level of QoS for the crossing of the whole chain of providers (SP_n, SP_{n-1}, SP_{n-2}, ..., SP₁) from its own domain to the domain where Destination is located, see (Figure 3). That means SP_n implicitly guarantees the level of QoS for the crossing of distant domains like SP_{n-2}. In the same way, SP_{n-2} is likely to be part of numerous SP chains and to see its level of QoS guaranteed by many providers it has maybe even never heard of.






Figure 3: Provider to provider agreement

Any modification in a given agreement is likely to have an impact on numerous external agreements that make use of it. A provider sees the degree of freedom to renegotiate, or terminate, one of its own agreements, being restricted by the number of external (to its domain) agreements that depend on it. This is referenced as the SP chain trap issue. Within the scope of global Internet services, each provider would find itself being part of a large number of SP chains. This solution is not appropriate for a worldwide QoS coverage as it would lead to glaciation phenomena, ending up with a completely petrified QoS infrastructure, where nobody could renegotiate any agreement.

If a QoS-enabled Internet is deemed desirable, with QoS services available potentially to and from any destination, as we are used to with the current Internet, any solution must resolve this problem and find alternate schemes for provider to provider agreements.

[5.2](#) Recommendations for provider agreements

These recommendations are based on two assumptions. First, to avoid forming SP chain traps, provider to provider agreements should not cover several domains. Second, since it is very hard to know about agreements more than one domain hop, and that these agreements can change, it is almost impossible to have an accurate visibility of their evolution.

Therefore, a provider should take the decision to bind one of its l-QCs to one of its neighbor l-QCs based exclusively on:

- What it knows about its own l-QCs;
- What it knows about its neighbor l-QCs.

Agreements are then based on guarantees covering a single domain.

For any n , SP_n should guarantee SP_{n+1} only the crossing performance of SP_n .

5.3 Edge to edge guarantees

It is very important to note that the proposition to limit guarantees to only one domain hop applies exclusively to provider to provider agreements. It does not in any way preclude edge to edge guarantees for a user communication.

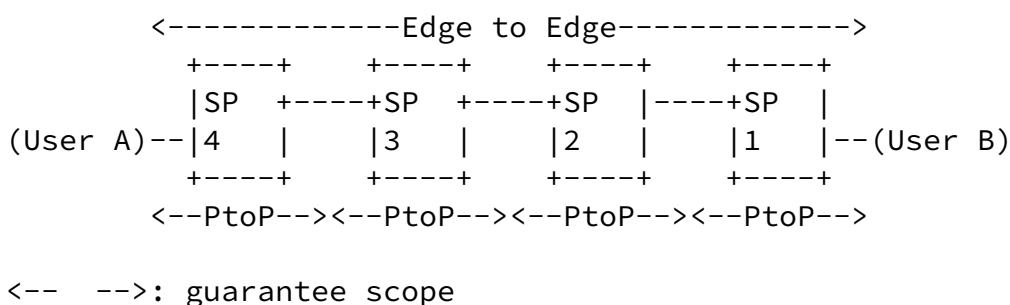


Figure 4: User communication guarantees

Any QoS inter-domain solution, either based on MQC or on a completely different approach, is valid as long as each provider claiming to offer some QoS performance actually delivers the expected level of guarantee. In Figure 4, edge to edge guarantee between users A et B is ensured by concatenation of local provider to provider (PtoP) agreements.

5.4 The need for MQC

For its binding process, a provider will not use any information related to what is happening more than one domain away. It will try to find the best match between its l-QCs and its neighbor l-QCs. That is to say, it will bind its l-QC with the neighbor l-QC that has the closest performance. However, at the scale of a communication, if there is systematically a slight difference between each upstream and downstream l-QC, there can be a significant difference between the first and the last l-QC. There must be a means to ensure the consistency and the coherency of a QoS treatment when traversing a given path.

A solution is to have a classification tool that defines two l-QCs as

being able to be bound together if, and only if, they are classified in the same category. Each category is called a Meta-QoS-Class (MQC). Two l-QCs can be bound together if, and only if, they belong to the same MQC.

[6.](#) The Meta QoS Class concept

[6.1](#) Definition

An MQC can be loosely defined as a label associated with a set of applications that bear similar network QoS requirements. It can be put on any l-QC suited to convey packets from this type of applications. It is a means to certify that this l-QC is suitable for the traffic of this application.

[6.2](#) Template

The following attributes should be documented in any specification of an MQC.

- o A list of services (e.g. VoIP) the MQC is particularly suited for.
- o Boundaries for each QoS performance attribute (D, J, L), whenever required. Several levels can be specified depending on the size of the network provider (regional, national, etc.).
- o Constraints on traffic (e.g. only TCP-friendly).
- o Constraints on the ratio: network resources for the class to overall traffic using this class (e.g. less resource than peak traffic).

[6.3](#) Binding process

A provider goes through several steps to extend its internal l-QCs. First, it classifies its own l-QCs based on MQCs. Second, it learns about available MQCs advertised by its neighbors. To advertise an MQC, a provider must have at least one compliant l-QC and should be ready to reach agreements to let neighbor traffic benefits from it. Third, it contracts an agreement with its neighbor to send some traffic that will be handled accordingly to the agreed MQCs. The latter stage is the binding process.

Note that when a provider contracts an agreement with a neighbor, it may well not know to what downstream l-QCs its own l-QCs are going to

be bound. It only knows that when it sends a packet requesting a given MQC treatment (for example, owing to an agreed DSCP marking) the packet will be handled in the downstream domain by an l-QC compliant with the treatment associated with this MQC.

[6.4](#) Properties

An MQC typically bears properties relevant to the crossing of one and only one domain. However this notion can be extended, in a straightforward manner, to the crossing of several domains, as long as the set of consecutive domains is considered as a single virtual domain.

MQC should not be confused with PDB concept defined in Diffserv architecture. The two notions share the common characteristic of specifying some QoS performance values. But the two concepts differ in their purposes. The objective for the definition of a PDB is to help implementation of l-QCs within a single administrative network. The objective for an MQC is to help agreement negotiation between providers, and therefore the process of binding two adjacent l-QCs.

The MQC concept is very flexible with regard to new unanticipated applications. According to the end-to-end principle [[E2E](#)], a new unanticipated application should have little impact on existing l-QCs, because the l-QCs should have been designed and engineered, to the extent possible, to gracefully allow any new application to benefit from the existing QoS infrastructure. However, this issue does not concern the MQCs as such, because an MQC is an abstract concept that has no physical existence. It is only the problem of l-QCs design and engineering. Therefore, a new unanticipated application could simply drive a new MQC and a new classification process for the l-QCs.

A hierarchy of MQCs can be defined for a given type of service (e.g. VoIP with different qualities: VoIP for residential and VoIP for enterprise). A given l-QC can be suitable for several MQCs (even outside the same hierarchy). Several l-QCs in a given domain can be classified as belonging to the same MQC.

MQC is a concept. MQC does not prohibit the use of any particular mechanism or protocol at the data, control, or management plane. For example, Diffserv, Intserv, traffic shaping, traffic engineering, admission control, and so forth, are completely legitimate. MQC simply drives and federates the way QoS inter-domain relationships are built.

[6.5](#) Common understanding

The need for standardization is strong as far as inter-domain QoS is concerned [[RFC2990](#)]. Each provider must have the same understanding of what a given MQC is about. A global agreement on a set of MQC standards is needed. The number of classes to be defined must remain very small to avoid an overwhelming complexity. There must be also a means to certify that the l-QC classification made by a provider conforms to the MQC standards. So the standardization effort should go along with some investigations on conformance testing requirements.

[7](#). New perspectives for a global QoS-enabled Internet

The MQC concept opens up new perspectives for a QoS enabled Internet that keeps, as much as possible, the openness characteristics of the existing best-effort Internet. This is certainly not the only domain of application to explore, but it is sufficiently important to deserve some special considerations.

The resulting QoS Internet can be viewed as a set of parallel Internets or MQC planes. Each plane consists of all the l-QCs bound accordingly to the same MQC. When an l-QC maps to several MQCs, it potentially belongs to several planes. Users can select the MQC plane that is the closest to their needs, as long as there is a path available for the destination.

The best effort service can be considered as relevant to the so-called best-effort MQC. It is of primary importance to maintain a best-effort route available when no QoS route is known. Best-effort delivery must survive QoS and must remain the main Internet transport service.

When a provider contracts with another provider based on the use of MQCs, it simply adds a logical link to the corresponding MQC plane, basically like what current traditional inter-domain agreement means for the existing Internet. As soon as a new domain joins an MQC plane, it can reach all domains and networks within this plane.

Figure 5 a) depicts the physical layout of a fraction of the Internet, comprising four domains with full-mesh connectivity.




Figure 5: MQC planes

Figure 5 b) depicts how these four domains are involved in a given MQC plane. SP1, SP2 and SP4 have at least one compliant l-QC (SP3 maybe has or not) for this MQC. A bi-directional agreement exists between SP1 and SP2, SP1 and SP4, SP2 and SP4.

Route path selection within a selected MQC plane can be envisaged as the traditional routing system used by the Internet routers: do your best to find one path, which is as good as possible. Thus, we can rely on a BGP-like protocol, for instance the proposal detailed in [\[q-BGP\]](#), for the inter-domain routing process. The resilience feature of the IP routing system is preserved: if a QoS path breaks somewhere, the routing protocol will make it possible to dynamically compute another QoS path if available, in the proper MQC plane.

8. IANA Considerations

There are no IANA considerations.

Note to RFC Editor: this section may be removed on publication as an RFC.

9. Security Considerations

This document describes a framework and not protocols or systems. Potential risks and attacks will directly depend on the implementation technology. Solutions to implement this proposal must detail security issues in the relevant protocol documentation.

Particular attention should be paid on giving access to MQC resources only to authorized traffic. Unauthorized access can lead to denial of service when the network resources get overused.

Levis & Boucadair	Expires January 1, 2006	[Page 11]
-------------------	-------------------------	-----------

Internet-Draft	The Meta-QoS-Class concept	June 2005
----------------	----------------------------	-----------

10. Acknowledgements

Part of this work is funded by the European Commission, within the context of the MESCAL (Management of End-to-End Quality of Service Across the Internet At Large) project (<http://www.mescal.org>). The authors would like to thank all the other partners for the fruitful discussions.

11. References

11.1 Normative References

- [RFC1771] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", [RFC 1771](#), March 1995.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", [RFC 2679](#), September 1999.
- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", [RFC 2680](#), September 1999.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", [RFC 3393](#), November 2002.

11.2 Informative References

- [E2E] Saltzer, J H., Reed, D P., and D D. Clark, "End-To-End Arguments in System Design", ACM Transactions in Computer Systems, Vol 2, Number 4, pp 277-288, November 1984.
- [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.
- [RFC2990] Huston, G., "Next Steps for the IP QoS Architecture", [RFC 2990](#), November 2000.
- [RFC3086] Nichols, K. and B. Carpenter, "Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification", [RFC 3086](#), April 2001.
- [RFC3869] Atkinson, R., Floyd, S., and Internet Architecture Board,

Levis & Boucadair	Expires January 1, 2006	[Page 12]
-------------------	-------------------------	-----------

Internet-Draft	The Meta-QoS-Class concept	June 2005
----------------	----------------------------	-----------

"IAB Concerns and Recommendations Regarding Internet Research and Evolution", [RFC 3869](#), August 2004.

- [WIP] Deleuze, G. and F. Guattari, "What Is Philosophy?", Columbia University Press ISBN: 0231079893, April 1996.
- [q-BGP] Boucadair, M., "QoS-Enhanced Border Gateway Protocol", [draft-boucadair-qos-bgp-spec-00.txt](#), Work in Progress, June 2005.

Authors' Addresses

Pierre Levis
France Telecom
42 rue des Coutures
BP 6243
Caen Cedex 4 14066

France

Email: pierre.levis@francetelecom.com

Mohamed Boucadair
France Telecom
42 rue des Coutures
BP 6243
Caen Cedex 4 14066
France

Email: mohamed.boucadair@francetelecom.com

Levis & Boucadair	Expires January 1, 2006	[Page 13]
-------------------	-------------------------	-----------

Internet-Draft	The Meta-QoS-Class concept	June 2005
----------------	----------------------------	-----------

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Copyright Statement

Copyright (C) The Internet Society (2005). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

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