# The Meta-QoS-Class concept: a step towards global QoS inter-domain services

Pierre Levis[1], Mohamed Boucadair[1], Pierrick Morand[1], Panos Trimintzios[2]

[1]France Telecom R&D, Caen France
E-mail: {pierre.levis, mohamed.boucadair, pierrick.morand}@francetelecom.com
[2]C.C.S.R. University of Surrey, Guildford United Kingdom
E-mail: p.trimintzios@eim.surrey.ac.uk

*Abstract:* **IP QoS (Quality of Service) is commonly presented as a radical shift in the IP paradigm and therefore, in the way we are used to accessing Internet services. Indeed, several architectures have been designed in order to introduce QoS that violate the fundamental features of the best-effort Internet. In this paper we demonstrate that having QoS services that maintain these features, especially the ability to easily connect any pair of users worldwide, naturally leads to the definition of the new concept of Meta-QoS-Class.**

## 1. INTRODUCTION

A true IP-based QoS delivery solution should prevent QoS techniques and architectures from impairing the spirit in which the Internet has been devised [1]. It should preserve the facility to (1) spread Internet access, (2) welcome new applications and (3) communicate from any location to any location. This position paper investigates in this direction and logically leads to the definition of the new concept of Meta-QoS-Class.

The remainder of this paper is organized as follows. Section 2 analyses the problem of end-to-end QoS based on agreements between Service Providers. Section 3 develops the concept of Meta-QoS-Class. Section 4 explains the use of Meta-QoS-Classes to build a QoS-enabled Internet.

This work has been conducted within the scope of the IST MESCAL[1] (Management of End to end quality of Service in the internet At Large) project.

## 2. FROM A BASIC INTER-DOMAIN QOS PROBLEM TO META-QOS-CLASS

### 2.1 Problem statement

We adopt a pragmatic view to tackle the problem of Internet QoS delivery. The solution will consist of Service Providers' QoS capabilities concatenated over the end-to-end path. These meshed QoS capabilities constitute the QoS infrastructure.

We start by closely examining the requirements, the opportunities and the consequences, for a given Service Provider (SP) to integrate this QoS infrastructure. We focus mainly on SP-to-SP agreements rather than on SP-to-customer agreements.

Let's consider a given Service Provider that offers QoS-based services to its customers. The scope of these services is limited to its network domain boundaries. On the other hand, this Service Provider is aware that many other Service Providers, scattered over the Internet, also offer QoS-based services to their customers. This Service Provider is expected to want to benefit from the QoS infrastructure in order to expand its QoS-based service offerings to destinations outside its own administrative domain.

### 2.2 Reaching QoS agreements with neighbors

There are at least two main approaches for this QoS service expansion. In the first approach, the Service Provider negotiates agreements only with its immediate neighboring Service Providers (that is to say, the ones that are directly accessible without the need to cross a third party SP). We call it the *cascaded* approach. In the second approach, the Service Provider negotiates directly with an appropriate number of downstream providers, one or more than one domain hop away. We call it the *centralized* approach.

There is a great deal of complexity and scalability issues related to the centralized approach, which represents a radical shift from current Internet best practice. Therefore, we believe that the only realistic way forward is the cascaded approach. This is the approach we adopt in the rest of this paper.

### 2.3 Binding l-QCs

A SP domain implements QoS capabilities in order to provide QoS-based services. We use the term local-QoS-Class or l-QC to denote a basic QoS transfer capability within a SP domain. A l-QC is characterized by a set of attribute-value pairs, where the attributes express various packet transfer performance parameters such as (D, J, L): one-way transit delay (D), one-way transit variation delay (also known as jitter) (J) and loss rate (L). The provisioning of an l-QC

[1] http://www.mescal.org

solely relies upon engineering policies deployed within the domain. Typically, a combination of the elementary IP DiffServ QoS capabilities with traffic-engineering functions should ensure the l-QC performances. A l-QC is one occurrence of a Per-Domain Behavior or PDB [2].

On a physical level, the QoS service extension to another domain signifies the l-QC extension outside the scope of a single domain. In particular, this means that packets from a flow originated in a domain, with a given DSCP (DiffServ Code Point) indicating a given l-QC, should experience a similar treatment when crossing the set of domains on the path towards its destination.

Two l-QCs from two neighboring SP domains are bound together when the two Service Providers have agreed to transfer traffic from one l-QC on the upstream domain to another l-QC on the downstream domain [3]. Then, if we assume that a Service Provider knows l-QCs capabilities advertised by its service peers, the basic technical question that this provider has to face is: on what basis shall I bind my l-QC to my neighbor SP l-QCs? Given one of my own l-QCs, which is the best match? Based on what criteria?

### 2.4 Limiting the scope of SP-to-SP agreements

In this section we look into the problems when SP-to-SP agreements guarantee QoS over a chain of downstream domains.

Let's assume that SPn knows from its neighbor SPn-1 a set of (Destination, D, J, L) where Destination is a group of IP addresses reachable via SPn-1, and (D, J, L) is the QoS performance to get from SPn-1 to Destination. SPn uses this information to bind its own l-QCs with SPn-1 l-QCs. SPn knows the QoS performance of its own l-QCs and therefore, deduces the QoS performance it could guarantee to its customers in order to join Destination. If this is a viable service and business opportunity, it will buy the (Destination, D, J, L) that best fit its operational objectives.

End-to-end QoS performance is guaranteed in a recurrent manner: SP1 guarantees QoS performance for its own domain crossing; while for a given n, SPn guarantees SPn+1 QoS performance for the crossing of the whole chain of SPs (SPn, SPn-1, …, SP1).

In this model, when a Service Provider contracts an agreement with a neighbor SP, a large number of other SP-to-SP and SP-to-customer agreements are likely to rely on that single agreement if it happens to be part of the chain of Service Providers. Any modification in that agreement is likely to have an impact on the numerous depending external agreements. The problem that arises is that you're not free to reconsider your own agreements, because other Service Providers, that may be you haven't even heard of, include the guarantee of your own agreements in their own agreements.

We call *SP chain trap* the fact that the degree of freedom to renegotiate, or terminate, one of your own agreements is restricted by the number of external (to your domain) agreements that depend on your own agreements. Within the scope of global Internet services, each Service Provider would find itself involved in a large number of SP chains.

This solution is not appropriate for global QoS coverage, it would lead to what we call *lake-freezing phenomenon*, ending up with a completely petrified QoS infrastructure, where nobody could renegotiate any agreement. We deem this lack of flexibility unacceptable for any Service Provider.

We do think that if a QoS-enabled Internet is 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 other schemes for SP-to-SP agreements. For this purpose, we introduce in the next section the concept of Meta-QoS-Class.

### 2.5 The need for Meta-QoS-Classes

A Service Provider knows very little about agreements more than one domain hop away. These agreements can change and it is almost impossible to have an accurate visibility of their evolution.

Furthermore, a Service Provider cannot guarantee anything but its own l-QCs in order to avoid being trapped in SP chains. Therefore, a provider should take the decision to bind one of its l-QCs to one of its neighbor SP l-QCs based solely on:

- What it knows about its own l-QCs
- What it knows about its neighbor SP l-QCs

A Service Provider should not use any information related to what is happening more than one domain hop away. It should try to find the best match between its l-QCs and its neighbor SP l-QCs. That is to say, it should bind one of its l-QC with the neighbor l-QC that has the closest performance (idea of extending l-QC). Agreements are then based on guarantees covering a single SP domain.

*For any n, SPn-1 guarantees SPn nothing but the crossing performance of SPn-1.*

We are confronted, at this point, with a problem of QoS path consistency. If there is systematically a slight difference between the upstream l-QC and the downstream l-QC we can wind up with a significant slip between the first and the last l-QC. Therefore, we must have a means to ensure the consistency and the coherency of a QoS SP's domain path. The idea 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. We call Meta-QoS-Class (MQC) each category of this l-QC taxonomy.

From this viewpoint: *two l-QCs can be bound together if, and only if, they correspond to the same Meta-QoS-Class.*

## 3. THE META-QOS-CLASS CONCEPT

### 3.1 Meta-QoS-Class based on application needs

The philosophy behind MQCs relies on a global common understanding of QoS application needs. Wherever end-users are connected, they more or less use the same kind of applications in quite similar business contexts. They also experience the same QoS difficulties and are likely to express very similar QoS requirements to their respective providers. Globally confronted with the same customers requirements, providers are likely to define and deploy similar l-QCs, each of them being particularly designed to support applications of the same kind of QoS constraints. There are no particular objective reasons to consider that a Service Provider located in Japan would design a “Voice over IP” compliant l-QC with short delay, low loss and small jitter while another Service Provider located in the US would have an opposite view. Applications impose constraints on the network, independently of where the service is offered on the Internet; see [4] for a survey on application needs.

Therefore, even if we strongly believe there is no Internet God, we consider that:

*There is a Customer God and he invented the Meta-QoS-Class concept.*

It should be understood that a MQC is actually an abstract concept. It is not a real l-QC provisioned in a real network.

### 3.2 Meta-QoS-Class definition

A MQCs is defined with the following attributes:

- A list of services (e.g. VoIP) the MQC is particularly suited for.
- Boundaries for the QoS performance attributes (D, J, L), if required.
- Constraints on type of traffic to put onto the MQC (e.g. only TCP-friendly).
- Constraints on the ratio: (resource for the class) to (traffic for the class).

Attributes could depend on the SP's domain diameter, for example a longer delay could be allowed for large domains. Performance attributes can be weighed in order to prioritize the ones the service is more sensitive to.

### 3.3 Compliance of l-QCs to a Meta-QoS-Class

A Service Provider goes through several steps to expand its internal l-QCs.

First, it classifies its own l-QCs based on MQCs. Second, it learns about available MQCs advertised by its neighbor. To advertise a MQC, a Service Provider must have at least one compliant l-QC and should be ready to reach agreements to let neighbor SP 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. A l-QC can be bound only with a neighbor l-QC that is classified as belonging to the same MQC.

Note that when a Service 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, thanks to an agreed DSCP marking) the packet will be handled in the downstream SP domain by an l-QC compliant with the requested MQC.

### 3.4 What's in and out of a Meta-QoS-Class

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

The MQC concept is very flexible with regard to new unanticipated applications. According to the end-to-end principle [5], a new unanticipated application should have little impact on existing l-QCs, because the l-QCs should have been designed, to the extent possible, to gracefully allow any new application to benefit from the existing QoS infrastructure they form. However, this issue does not concern the MQCs per se, because a MQC is an abstract object that has no physical existence. It is solely 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). A given l-QC can be suitable for several MQCs (even outside the same hierarchy). In this case, several DSCPs are likely to be associated with a same l-QC in order to differentiate between traffic classes. Several l-QCs in a given SP domain can be classified as belonging to the same MQC.

The DiffServ concept of PDB [2] should not be confused with the MQC concept. The two concepts share the common characteristic of specifying some QoS performance values. The two concepts differ in their purposes. The objective for the definition of a PDB is to help implementation of QoS capabilities within a network. A MQC does not describe the way to implement a l-QC or PDB. The objective for a MQC is to help agreement negotiation between Service Providers.

### 3.5 Meta-QoS-Class interest summary

In summary the interest of the MQC concept is threefold. It:

- Provides guidance for l-QC binding.
- Allows relevant l-QC binding with no knowledge of the following distant provider agreements.

- Enforces coherency in a QoS path without any knowledge of the complete domain path.

## 4. THE FUNDAMENTAL USE CASE: THE QOS INTERNET AS A SET OF META-QOS-CLASS PLANES

### 4.1 MQC planes

We describe here the fundamental use case, for a QoS-enabled Internet, based on the MQC concept. Our purpose is to build a QoS-enabled Internet that keeps, as much as possible, the openness characteristics of the existing best-effort Internet, and more precisely conforms to the requirements expressed earlier in this paper.

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

Figure 1 depicts the physical layout of a fraction of Internet, comprising four domains from four different SPs, with full-mesh connections.

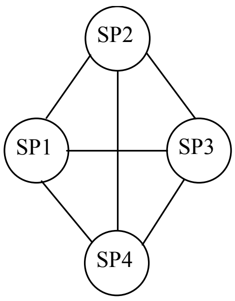


*Figure 1 - A physical configuration*

Figure 2 depicts how these four SPs are involved in two different MQC planes.

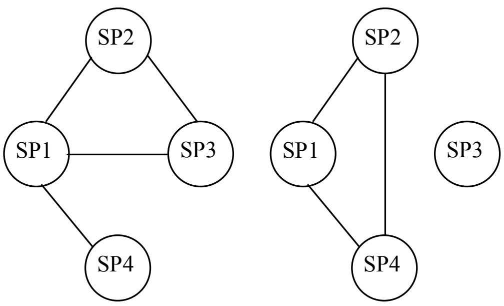


*Figure 2 - Two Meta-QoS-Class planes*

Considering the left-hand plane, we can deduce: (1) each SP has at least one compliant l-QC for the given first MQC, and (2) a bi-directional agreement to exchange traffic for this class exists between SP1 and SP2, SP1 and SP3, SP1 and SP4, SP2 and SP3.

Considering the right-hand plane, we can deduce: (1) SP1, SP2 and SP4 have at least one compliant l-QC (SP3 maybe has, maybe has not) for the given second MQC, and (2) a bi-directional agreement exists between SP1 and SP2, SP1 and SP4, SP2 and SP4.

We assume that in each MQC plane, because we want to stay close to the Internet paradigm, all paths are equal. Therefore, the problem of path selection amounts to: do your best to find one path, as good as possible, for the selected MQC plane. This sounds like the traditional routing system used by the Internet routers. Thus, we can rely on a BGP-like (Border Gateway Protocol) inter-domain routing protocol for the path selection process. We can call this protocol q-BGP. By destination, q-BGP selects and advertises one path for each MQC plane. From an abstract view, each MQC plane runs its own BGP protocol. When it comes to implementation, there can be only one q-BGP session between two SP domains, shared by all MQC planes, thanks to the use of the QoS_NLRI attribute [6].

When, for a given MQC plane, there is no path available to a destination, the only way for a datagram to reach this destination is to use another MQC plane. The only MQC plane available for all destinations is the best-effort MQC plane (also known as the current Internet).

For global access services our solution stands only if MQC-based binding is largely accepted and become a current practice. Otherwise, a MQC plane will have too many holes. Noteworthy, this limitation is due to the nature of the service itself, and not to the use of MQCs. Insofar as we target global services we are bound to provide QoS in as many SP domains as possible. However, any MQC-enabled part of the Internet that forms a connected graph can be used for QoS communications, and incrementally extended. Therefore, incremental deployment does lead, to a certain extent, to incremental benefits. For example, in Figure 2 right-hand plane, as soon as SP3 connects to the MQC plane it will be able to benefit from SP1, SP2 and SP4 QoS capabilities.

We can now elaborate a bit more on what it means for a Service Provider to contract an agreement with another Service Provider based on the use of the MQC concept. It simply means adding a link to the corresponding MQC plane, basically just what current traditional inter-domain agreement means for the existing Internet. As soon as a SP domain joins a MQC plane, it can reach all domains and networks within the plane.

This set of domains and networks is prone to evolve dynamically along with the appearance of new inter-domain agreements and the revocation of old inter-domain agreements. However, for a given SP-to-SP agreement, in a given MQC plane, any evolution elsewhere in this plane has no direct impact on this agreement. We are not, therefore, prone to the *lake-freezing phenomenon* and we can easily change our inter-domain agreements so far as our neighbor Service Providers agree.

We fully benefit from the resiliency feature of the IP routing system: if a QoS path breaks somewhere, the q-BGP protocol will make it possible to compute another QoS path dynamically in the proper MQC plane.

Each Service Provider must have the same understanding of what a given MQC is about. A global agreement, on a set of standards, is needed. This agreement could be typically reached in an international standardization body. The number of MQCs defined, and consequently the number of MQC planes, must remain very small to avoid an overwhelming complexity. The need for some sort of standardization is rather evident as far as inter-domain QoS is concerned [7]. There must be also a means to certify that the l-QC classification made by a Service Provider conforms to the MQC standards. So the MQCs standardization effort should go along with some investigations on conformance testing requirements.

### 4.2 Levels of guarantee

Any QoS inter-domain solution, either based on MQC or on a completely different approach, is valid as long as each Service Provider claiming some QoS performances actually delivers the expected level of guarantee. In our solution this is ensured by concatenation of local binding agreements, without any broad agreement covering the whole QoS path.

It's very important to notice, that here, we are only speaking of SP-to-SP agreements. Having SP-to-SP agreements limited to only one domain span, should not preclude having SP-to-customer agreements guaranteed edge-to-edge, from first domain ingress point to last domain egress point.

There is often confusion about QoS as an underlying technology and QoS as a service offering [8].

If we target harder administrative guarantees, for well-delimited services like VPN, we can, for example, use MQC information exchange by q-BGP to find a QoS path that fits the demand, and then reach an agreement for all SP actors of the selected QoS path, possibly enforcing the path by a MPLS TE tunnel [9].

More generally, MQC is a concept. MQC doesn't prohibit the use of any particular QoS mechanism or protocol whether at the data plane, control plane or management plane. For example, DiffServ, Traffic Shaping, Traffic Engineering, Admission Control, Bandwidth Broker, Billing, and so forth, are completely legitimate. MQC simply drives and federates the way QoS inter-domain relationships are built.

## 5. CONCLUSION

In this is paper we have introduced the new concept of Meta-QoS-Class. It significantly helps Service Providers to negotiate agreements. It avoids what we have called *SP chain traps* leading to *lake-freezing phenomenon*. It Provides guidance for l-QC binding. It allows relevant l-QC binding with no knowledge of the following distant provider agreements. It enforces coherency in a QoS path without any knowledge of the complete domain path.

Meta-QoS-Class concept opens up an innovative way to achieve global QoS Internet connectivity that maintains the main features of the Internet. It could open a new path in the inter-domain QoS research area and enable new QoS models to be introduced.

## 6. AKNOWLEDGMENT

This work was undertaken in the Information Society Technology (IST) MESCAL project, which is partially funded by the European Commission. We would like to thank the other MESCAL partners for their inputs and discussions.

## REFERENCES

[1] D.D. Clark, "The Design Philosophy of the DARPA Internet Protocols", Proc. ACM SIGCOM 88.

[2] K. Nichols, B. Carpenter, "Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification", RFC3086, Apr. 2001.

[3] P. Levis, A. Asgari, P.Trimintzios "Considerations on inter-domain QoS and Traffic Engineering issues through a utopian approach", to appear in proceedings of SAPIR Workshop, Fortaleza Brazil, Aug. 2004, ©Springer-Verlag.

[4] D. Miras, "A Survey on Network QoS Needs of Advanced Internet Applications", Internet2 QoS Working Group working document, Dec. 2002.

[5] J.H. Saltzer, D.P. Reed, D.D. Clark, "End to End Arguments in System Design", in ACM Transactions on Computer Systems*,* 1984.

[6] G. Cristallo, C. Jacquenet, "An Approach to Inter-domain Traffic Engineering", in Proceedings of XVIII World Telecommunications Congress (WTC2002), Sep. 2002.

[7] M. Eder, H. Chaskar, S. Nag, "Considerations from the Service Management Research Group (SMRG) on Quality of Service (QoS) in the IP Network", RFC3387, Sep. 2002.

[8] B. Teitelbaum, S. Shalunov, "What QoS Research Hasn't Understood About Risk", RIPQoS Workshop SIGCOMM2003, Aug. 2003.

[9] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC2702, Sep. 1999.