END-TO-END QUALITY IN IP NETWORKS: CAN WE OFFER AND CHARGE IT?

Marco Alfano
CSELT
Via G. Reiss Romoli 274
10148 Torino, ITALY
Ph: +39 011 2286850
Fax: +39 011 2286207
marco.alfano@cselt.it

Magnus Krampell
EURESCOM GmbH
Schloss-Wolfsbrunnenweg 35
D-69118 Heidelberg, GERMANY
Ph: +49 6221 989 381
Fax: +49 6221 989 209
krampell@eurescom.de

Mikhail Smirnov
GMD FOKUS
Kaiserin-Augusta-Allee, 31
Berlin 10589, GERMANY
Ph: +49 30 3463 7113
Fax: +49 30 3463 8000
smirnow@fokus.gmd.de

Key-Theme: T1

Abstract

IP networks and in particular the Internet are becoming ever more widespread. Many services/applications with different requirements use these networks, but presently there are no global methodologies and tools for managing their Quality of Service (QoS). Moreover, it is not clear which QoS parameters to use for expressing QoS at the user level (QoS perceived) and how to correlate them to and network performance (NP) parameters. Finally, services/applications in these networks (including the Internet) are not presently charged on the basis of user QoS requirements and QoS offered by service/network providers.

This paper addresses the main issues that, in our opinion, have to be dealt with in order to be able to “offer” end-to-end quality to the end user. In particular, we present a QoS model that allows the specification of *Quality Classes* for the end user and the mapping of them into network performance levels. We also present the charging aspects related to the QoS model.

1. Issues and challenges of end-to-end quality in IP networks

Quality of Service (QoS) is a key factor for differentiating service offers in a competitive market. However, measurement, management and charging of QoS are not yet fully deployed and, above all, are not engineered yet in an integrated way. Many mechanisms exist nowadays for controlling performance in IP networks (e.g., IETF diff-serv and int-serv) but it is not clear how these mechanisms can be used in order to provide end-to-end quality, i.e., a quality that can be perceived by the end user.

Customers are faced with defining their QoS requirements, somehow. Can this be done, perhaps, by selecting among a set of predefined Quality Classes (Figure 1)?

Moreover, once the customer’s QoS requirements are defined, the next step is to translate these into QoS parameters that can be measured and managed. Finally, measurements need to be converted into data that can be used for charging.

Another dimension of the problem is when more than one provider is involved. The requirements need to be negotiated among the providers in the end-to-end chain. Measurements need to be made along the chain and data for charging need to be transferred between providers.

It should be noted that the main costs related to QoS provisioning are not on implementing QoS but on measuring, managing and charging it [1]. To alleviate some of these issues, the following requirements should be imposed:

- the number of Quality Classes a user can select among should be minimised;
- the number of QoS parameters used should be minimised;
- the number of target values for those parameters should be minimised.

Finally, note that in order to provide end-to-end quality to the user, the user’s domain (LANs, terminals, applications) must also be taken into account in the end-to-end transmission chain because they play an important role in the quality that is going to be actually perceived by the user.

Figure 1. User's quality classes

2. Related work

We made a survey of some exising QoS models by MMCF [3], ETSI (TIPHON [4], ETR 03 [5], ETR 138 [6]), ITU-T (E.800 [7], I.350 [8], I.380 [9]), EURESCOM (P616 [10], P806 [11]) and IETF (IntServ [12], DiffServ [13]) in order to evaluate how they deal with the issues discussed above. In particular, we evaluated their ability to deal with a number of aspects that can be described as: *End-to-end QoS, Quality Classes, Charging, Measurement/Management and Service Specific(S)/Service Independent(I).* The following table summarises the findings.

	End-to-end QoS	Quality Classes	Charging	Measurement/ Management	Specific (S)/ Independent (I)
MMCF	Y	Y	N	Y	S
ETSI	Y	Y (only [4])	N	Y	S
ITU	Y	N	N	Y	I
EURESCOM	Y	N	N	Y	I
IETF	N	N	N	Y	I

Table 1. Analysis of some existing QoS models (Y=Yes, N=No).

From this table the following conclusions can be drawn:

1. Only a few models consider the concept of Quality Classes.
2. None of the models handles charging!
3. Some models are Service (Application) specific.

From our survey (though limited), it seems that a work on QoS that contemporarily deals with all the different aspects indicated in Table 1 has not been yet developed.

3. A proposed solution for end-to-end quality

In this section we describe a possible approach for the definition, offer and charging of end-to-end quality in IP networks. We believe that the two following major components are needed:

- A QoS model that allows the provider to make an "offer" to the end user in terms of "Quality Classes" for the types of services/applications (expressed in terms of "Service Categories") he/she is using. This model also considers mapping between Quality Classes and Service Categories to Network Performance Levels. The user is presented with an offer and the charge related to the offer (Figure 2).

- A QoS framework that considers the network infrastructure to support the QoS model as well as the solutions for QoS management, measurement and charging.

In the next two subsections we discuss the QoS model and the QoS charging aspects in more details. The presented results (as well as the QoS survey discussed above) stem from an ongoing EURESCOM project [2], named "QUASIMODO".

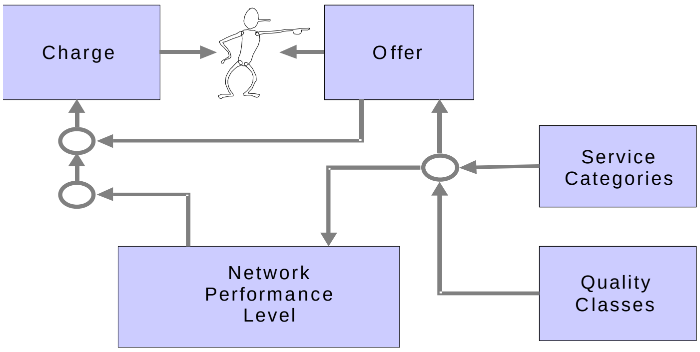

Figure 2. The main components of the proposed QoS model.

3.1 A QoS model

As said above, we want to be able to offer end-to-end quality to the end user. To this end, we consider the scenario reported in Figure 3. This scenario is based on the realistic assumption that the user will create and control his/her domain on his/her own.

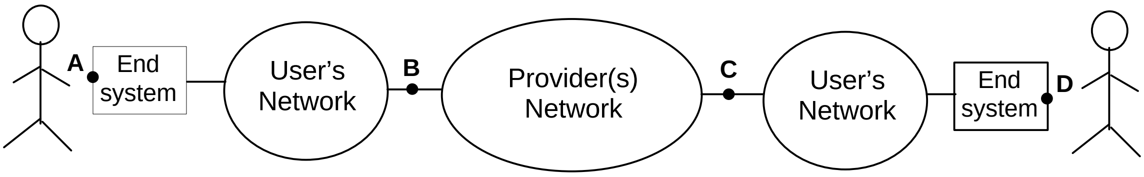

Figure 3. Scenario for the offer of end-to-end quality.

In this scenario a network provider will offer some network performance levels at the edges of its network (B and C points in Figure 3). Moreover it will characterize the user's domain (e.g., LANs, terminals, applications) in terms of a minimal number of characteristics and performance levels.

In doing so, we make sure that the user domain does not become the bottleneck of the whole transmission chain and that the user can actually perceive a different quality in his/her services/applications (points A and D in Figure 3) when the network performance levels (between B and C in Figure 3) change. Of course, if the provider is also providing part or all of the user's domain, its task will be even easier because it will directly control (part or all of) the user's domain.

It should be noted that, in the case of a multiprovider environment, the points B and C of Figure 3 will indicate the edges of all the providers' networks involved in the data trasmission between the two user domains. Clearly, in this case, agreements must be signed between providers in order to provide the requested end-to-end network performance.

Given the above scenario, the main characteristics of the QoS model can be summarised as follows:

- A set of "Service Categories" (SC) is considered. A Service Category groups services/applications that have similar needs (in terms of network performance) in order to work properly (i.e., with a minimal quality level).
- Users, for a give SC, can select between a number of "Quality Classes" (QC) (depending on his/her requirements) which have different charging implications. The Quality Classes

are created such that there is always a perceivable difference between one class and the "next" higher class.

- A set of network parameters are to be managed in order to provide the desired Quality Class to the user for a certain Service Category.

Since Quality Classes and Service Categories are unrelated we can represent their relationship in a table (Table 2).

	SC_1	**SC_2**	**….**	**SC_n**
QC_1	NPL_{11}	NPL_{12}	...	NPL_{1n}
QC_2	NPL_{21}	NPL_{22}	...	NPL_{2n}
…	...	...	...	...
QC_m	NPL_{m1}	NPL_{m2}	...	NPL_{mn}

Table 2: The QoS model

Each combination (QC, SC) corresponds to a Network Performance Level (NPL). A NPL is a (small) set of network perormance parameters (e.g., delay, jitter, loss), each with a value range and a guarantee that the parameter value will stay inside the range. This represents the target for the network in order to provide the desired Quality Class when a service from a certain Service Category is used. Each NPL can be seen as a possible "flow" in the network and must be trated accordingly.

The table and description above represents the ideal case. However, the number of SCs and QCs respectively must be decided. Since the number of NPLs increases exponentially with the number of SCs and QCs, these should be kept small enough to be manageble.

In practice, we believe that a practical QoS model could contain two Quality Classes: "Gold" and "Basic" and two Service Categories: "SC1" and "SC2" relative to services/applications characterised by significantly different performance requirements (e.g., real time vs. non real time). The combination of the two QCs with the two SCs results at most in four possible data flow aggregates to be treated differently in the network in order to preserve the performance requirements of the relative applications.

Figure 4 reports an example. The NPLs are defined in terms of delay, jitter and loss parameters, upper bounds on the values of those parameters and related guarantees (e.g., % of time the parameter value is below the threshold).

	SC 1, (non-Real Time)	**SC 2,** (Real Time)
Gold	D<= 1 s, J<= 0.5 s, L<= 10^{-3} Guarantee >= 90%	D<= 100 ms, J<= 10 ms, L<= 10^{-3} Guarantee >= 90%
Basic	D<= 10 s, J<= 5s, L<= 10^{-1} Guarantee >= 80%	D<= 300 ms, J<= 30 ms, L<= 10^{-2} Guarantee >= 80%

Figure 4. An example of QoS model (Numbers are purely indicative).

A potential offer for the user/customer, beside the NPL corresponding to the QC and SC chosen by the user, will contain the description of the user traffic that can be injected in the network, the characterization of the user domain (as discussed above) and the cost of that offer for the user. This offer can be part of a Service Level Agreement (SLA) to be agreed between the user and the provider.

Reconsidering the scenario of Figure 3, the offer just described will then contain a NPL to be provided between points B and C (the provider's network). This means that the provider will have to implement QoS control and resource management mechanisms in its network to assure the proper treatment of data flows. Moreover, measurements have to be performed between points B and C in order to verify the compliance of the actual performance with the NPL specified and sold to the user.

3.2 QoS Charging in the QoS Model

We argue that in a QoS aware IP network, *QoS charging* becomes a service with 5 logical layers – billing, charging, accounting, assembly and metering (Fig. 5b). Each of the layers can be organised by a provider in a proprietary structure given that the QoS Charging interworking with neighbour domains is assured - QoS charging is essentially an end-to-end service. A number of possible interactions between charging service components within a domain and between domains can be efficiently implemented through policy control functions as proposed in [15].

Since in our model we differentiate between quality classes and service categories we assume QoS charging differentiation support as an important requirement as well. The motivation for doing so is merely the observation that *QoS Charging becomes a service of its own* with service differentiation. Fig. 5a represents this viewpoint. Basic charging may mean flat rate charging (for non real time basic service) or usage based (i.e. resources usage metering) charging which is hidden from an end-user (does not appear as a service component) at a user service interface, but of course appears in the SLA with appropriate ranges of parameters.

Premium charging differs from the basic one in the amount of control and configuration capabilities a user might have over its mechanisms. Gold quality class assumes that QoS charging becomes a service of its own - a user paying more for the QoS should benefit additionally from the fact that many of QoS charging layers (from billing to possibly metering) are under his/her partial control. Thus, premium charging appears not only as a component of an SLA but also as a user service interface component. As Figure 5a shows, this raises the complexity of implementation and maintenance of QoS charging.

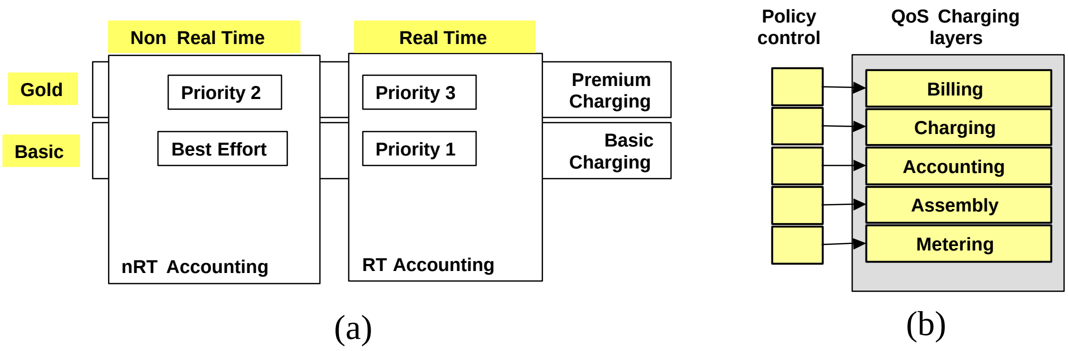

Figure 5. QoS charging in the QoS model (a) and charging layers (b)

Bearing in mind that next generation of IP routers operates already at Gigabit speeds, it becomes obvious that overcomplicated QoS charging schemes will not be feasible. We argue, however, that QoS charging being integrated into the overall design of QoS aware IP internetworking as its essential component can cope with all its challenges.

4. Conclusions and future work

When our QoS model is analysed using the same criteria as for the existing models (reported above) we see that it complies with all the requirements. In particular:

1. Using Quality Classes that end-users understand, the model is truly end-to-end.
2. The QoS model considers Quality Classes aimed towards the end user.
3. The QoS model deals with charging and proposes a charging service differentiation for quality classes.
4. The QoS model deals with measurement/management. The network parameters can be measured and controlled by different mechanisms (e.g. policing, shaping, metering, QoS routing, etc.)

5. The QoS model is service independent but s*ervice aware*! This means that for a certain Quality Class (e.g. Gold) the Network Performance Levels are different depending on which Service Category is used.

To our best knowledge, this is one of the first attempts to create a single framework that provides end-to-end quality (to the user) and, at the same time, deals with such different aspects as measuring, managing and charging QoS.

Of course many problems have still to be solved. First of all, the practical applicability of the proposed solution has to be shown through experiments. We are planning to implement the QoS model and perform such experiments. Nevertheless, we believe that the flexibility of our model and yet its simplicity should lead to an easy implementation of the model while getting satisfying results in terms of the quality perceived by the user when considering different quality classes and different types of services/applications.

5. Acknowledgements

The ideas described in this paper stem from a EURESCOM project (P906 – QUASIMODO). Information about this project, and access to deliverables can be found on the EURESCOM web site http://www.eurescom.de/Public/Projects/p900-series/P906/P906.htm.

References

[1] R. Guerin “Quality of Service: Too little or too much?” – Proc. IEEE Sixth International Workshop on Quality of Service (IWQoS ’98), Napa (CA), May, 1998.

[2] EURESCOM P906: QUASIMODO – QUAlity of ServIce MethODOlogies and solutions within the service framework: Measuring, Managing and Charging QoS.

[3] MMCF (MultiMedia Communication Forum) Part I: Framework; Part II: Multimedia desktop collaboration requirements Final MMCF document 1995

[4] ETSI TR 101329 V1.2.5 Telecommunication and Internet Protocols Harmonisation Over Networs (Tiphon) General Aspects of Quality of Service 1998-10

[5] ETSI ETR 003, General Aspects of Quality of Service (QoS) and Network Performance (NP), October 1994.

[6] ETSI ETR 138, Quality of service indicators for Open Network Provision (ONP) of voice telephony and Integrated Services Digital Network (ISDN), December 1997.

[7] ITU-T Rec. E.800 - “Terms and definitions related to Quality of Service and Network Performance including dependability” August 1994

[8] ITU-T Rec. I.350 - “General aspects of Quality of Service and Networks Performance in digital networks, including ISDN” March 1993

[9] ITU-T Rec. I.380 – “Internet Protocol Data Communication Service – IP Packet Transfer and Availability Performance Parameters.1999.

[10] EURESCOM P616 Enhanced ATM implementation issues 1998

[11] EURESCOM P806 EQoS - A Common Framework for QoS/Network Performance in a multi-Provider Environment 1999

[12] IETF RFC 1633. Integrated Services in the Internet Architecture: an Overview, June 1994.

[13] IETF RFC 2475. An Architecture for Differentiated Services

[14] Y. Bernet, et al, “A Framework for Differentiated Services”, IETF, work in progress document draft-ietf-diffserv-framework-02.txt, February, 1999

[15] G. Carle, F. Hartanto, M. Smirnov, T. Zseby “Charging and Accouting for QoS enhanced IP Multicast“- Proc. IFIP Sixth International Workshop on Protocols For High-Speed Networks (PfHSN '99), Salem, MA, August, 1999.