Considerations on inter-domain QoS and Traffic Engineering issues through a utopian approach

Pierre Levis

France Telecom R&D, Caen France
pierre.levis@francetelecom.com

Abolghasem (Hamid) Asgari

Thales Research & Technology (TRT) UK Ltd, Reading United Kin gdom
hamid.asgari@thalesgroup.com

Panos Trimintzios

C.C.S.R. University of Surrey, Guildford United Kingdom
P.Trimintzios@eim.surrey.ac.uk

Abstract. End-to-end QoS has been seldom studied in its inter-domain aspects, particularly within the scope of the global Internet and from an engineering perspective. This paper is intended to be a contribution in this direction. Starting from a utopian model that achieves any kind of engineering operations without any specific constraint, we deduce the open issues andproblems to be solved by a viable inter-domain QoS model that could be deployed in operational networks. We provide some guidelines that will help the design of a QoS-based inter-domain Traffic Engineering solution. This paper is part of the work conducted within the IST MESCAL project.

1. Introduction

Despite the years-long effort put into IP QoS (Internet Protocol Quality of Service), inter-domain aspects have been seldom studied, particularly within the scope of the global Internet and from an engineering perspective. However, the fact that both services and customers are spread all over the world, together with the inherent multi-domain nature of the Internet, requires a solution to handle inter-domain QoS delivery. Some work has been conducted in this direction within the scope of the IST MESCAL[1] (Management of End to end quality of Service in the internet At Large) project. Preliminary studies have shown that all potential solutions, besides their own pros and cons, present a set of common issues. The intent of this paper is to share this experience by highlighting the major issues a QoS-enabled Internet solution would have to solve. The novel feature of this paper is to consider all the problems a packet

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

can encounter in an end-to-end inter-domain QoS path, on what we can call a pure transportation level.

This paper is organised as follows. Section 2 defines the terms used in the paper. Section 3 introduces a utopian model, while Section 4 describes an example of usage scenario. Section 5 highlights the inter-domain QoS issues and provides some guidelines. Section 6 briefly explains the utopian model limitations and the current work being carried out.

2. Definitions

This paper makes use of the following definitions:

– **{D, J, L}:** a triple that refers to the metrics defined in the IPPM (IP Performance Metrics) working group in the IETF. "D" refers to one-way delay [1], "J" refers to IP packet delay variation (a.k.a. Jitter) [2] and "L" refers to one-way packet loss [3].
– **Autonomous System (AS):** a collection of routers under a single administrative authority enforcing a common routing policy.
– **Service Level Specification (SLS):** a set of technical parameters negotiated for a given service [4].
Customer: an entity that negotiates a service on behalf of one or several end-users.
– **Customer SLS (cSLS):** an SLS established between a provider and a customer.
– **Provider SLS (pSLS):** an SLS established between two providers.
– **QoS Class (QC):** a basic QoS transfer capability expressed in terms of {D, J, L}.
– **Local QC (l-QC):** a QC that spans a single AS (notion similar to Per Domain Behaviour (PDB) [5]).
– **Extended QC (e-QC):** a QC that spans several ASs. It consists of an ordered set of l-QCs.

3. Utopian end-to-end QoS model assumptions

The "utopian" qualifier expresses the fact that we can take the liberty of selecting and activating any network function regardless of the way it could be actually implemented. The intention is to simplify the discussion by relaxing non-essential technical constraints. From this perspective, the "utopian" model is practically unfeasible and consequently non viable. Nevertheless, it has proven to be an effective tool to identify and qualify the issues that any realistic model should consider, in order to be operationally deployed in real networks.

We consider the whole set of ASs with Diffserv-like QoS capabilities [6] in the Internet. These capabilities are provider-specific and can differ: by the number of l-QC deployed, by the respective QoS characteristics of each l-QC and by the way they have been locally implemented. Thus, we don't put any constraints on the intra-domain traffic engineering policies and the way they are enforced. When crossing an AS, a traffic requesting a particular QoS treatment experiences conditions constrained by the values of the QoS triple {D, J, L} corresponding to the l-QC applied by the provider. We assume service peering agreements exist between adjacent ASs enabling the providers to benefit from the QCs implemented within the neighbouring domain, and thus allowing the building of a set of e-QCs. We assume that a specific QoS negotiated between a customer and a provider is described by a {D, J, L} value in the corresponding cSLS.

The model is based on the following utopian fundamental assumptions:

– From any source to any destination we are able to compute the resulting {D, J, L} for all AS paths and l-QC combinations.
– From any source to any destination we are able to force any given AS path and a given l-QC in each AS along this path.

It should be noted that we never force an AS to support any arbitrary QC, but we leave the definition of l-QCs as a local decision. We always use the QoS capabilities of the ASs as they are. Therefore if a client requests a specific {D, J, L} and there is no AS chain matching exactly this request, we don't strive to re-engineer some ASs in order to exactly fulfil the request. However, if there is at least one AS chain that satisfies the requested QoS, we can select and activate the appropriate e-QC.

4. Usage scenario

This section describes how the utopian model operates. For the sake of clarity we assume that the Internet is dwindled down to the size as shown in Fig.1:

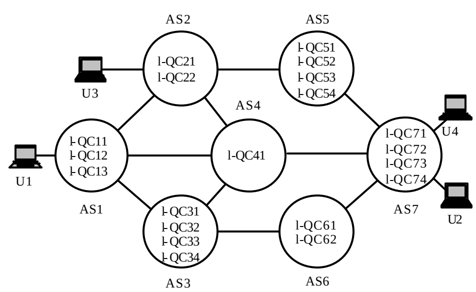

Fig. 1. A simplified view of the Internet

We focus on two pairs of end-users: (U1, U2) attached respectively to AS1 and AS7 and (U3, U4) attached to AS2 and AS7. QoSi denotes a QoS capability number i requested by a customer. l-QCxy means l-QC number y implemented in AS number x

A customer asks for QoS1 and QoS2 for traffic T1 and T2 directed from U1 to U2. Similarly, another customer asks for QoS3 for traffic T3 directed from U3 to U4. From U1 to U2 we investigate all possible AS path combinations (loop free combinations only) and compute the resulting e-QCs. The end-to-end {D, J, L} value is returned for each AS path combination. The fact that we rely on a utopian model enables us to enumerate all possible values and choose the best. This approach would simply not scale in an operational environment, even for moderately small AS-level networks due to the (hyper-) exponential growth of alternatives [2] The whole set of {D, J, L} values can be computed as follows:

(AS1, AS2, AS5, AS7): 3*2*4*4 = 96 values.
(AS1, AS2, AS4, AS3, AS6, AS7): 3*2*1*4*2*4 = 192 values.
(AS1, AS2, AS4, AS7): 3*2*1*4 = 24 values.
(AS1, AS3, AS4, AS2, AS5, AS7): 3*4*1*2*4*4 = 384 values.
(AS1, AS3, AS4, AS7): 3*4*1*4 = 48 values.
(AS1, AS3, AS6, AS7): 3*4*2*4 = 96 values.
(AS1, AS4, AS7): 3*1*4 = 12 values.
(AS1, AS4, AS2, AS5, AS7): 3*1*2*4*4 = 96 values.
(AS1, AS4, AS3, AS6, AS7): 3*1*4*2*4 = 96 values.
Total = 1044 values.

We compare these 1044 results with the requested QoS1 and QoS2 and say we deduce:

- The best path for QoS1 is eQC1:(l-QC11, l-QC21, l-QC51, l-QC71) (see Fig.2 solid line).
- The best path for QoS2 is eQC2:(l-QC12, l-QC21, l-QC54, l-QC74) (see Fig.2 dashed line).

Likewise, we compute from U3 to U4 all possible combinations of AS and QC. The {D, J, L} end-to-end value is returned for each combination. We compare these results with the requested QoS and say we deduce:

- The best path for QoS3 is e-QC3:(l-QC21, l-QC41, l-QC72) (see Fig.2 dotted line).

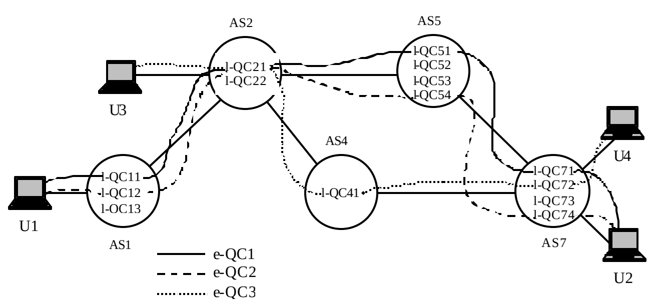

Fig. 2. Paths

Finally, we route T1, T2 and T3 traffics so that they respectively experience the selected e-QC1, e-QC2 and e-QC3.

[2] The problem we are trying to solve here is actually a well known NP-complete problem, i.e. routing based on multiple metrics and constraints, e.g. see [7] and [8] (though our problem is slightly different because the metrics are on the vertices instead of the edges, it can be shown that computationally it is exactly the same problem).

5. Inter-domain QoS issues

Apart from the complexity in e-QC computation, the scenario described above enables to highlight some important issues raised by the utopian model, namely: l-QC splitting (which is a new notion introduced by this paper), AS path selection criteria and impact on IP routing. In each case, we provide a description of the problem within the context of the utopian model, a definition of the issue and some guidelines for the design of viable (non-utopian) models.

5.1. l-QC splitting

In the utopian model, several e-QCs can share one l-QC within a single AS; traffic using this l-QC may then need to be split into different l-QCs in a downstream AS:

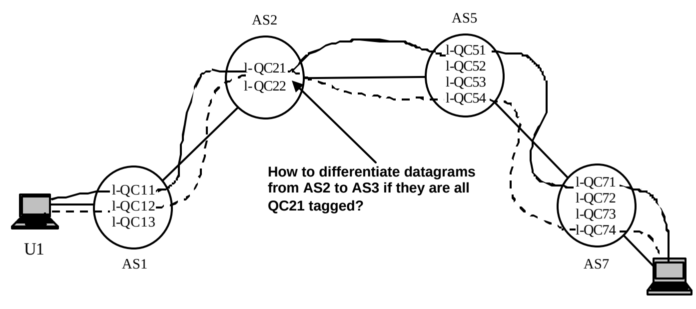

Fig. 3. l-QC splitting problem

In Fig.3, T1 and T2 traffics following e-QC1 and e-QC2 both use l-QC21 in AS2, but they need to be split into l-QC51 and l-QC54 respectively in AS5. The problem that arises is as follows: what should be the appropriate DSCP (Differentiated Services Code Point) marking for the datagrams forwarded to AS5? AS2 could for instance, from a utopian standpoint, rely on an internal mapping table from e-QC label to next DSCP. This label would be tagged in each packet (e.g. IPv4 option or DSCP fields, or IPv6 flow label [9] field) by the source or by a device close to the source. It would indicate the e-QC each datagram belongs to. This label could be a global value agreed by all ISPs or a local value understandable by two adjacent ASs (then it is likely to be modified by any AS). This solution would assume AS2 knows all cSLSs crossing AS2, that is to say all end-to-end QoS communications crossing AS2!

We define l-QC splitting as the process of binding one l-QC of an AS to several l-QCs in the next AS. The issue is to select for each datagram the appropriate downstream l-QC. Any viable solution should explicitly state how it deals with the l-QC splitting problem. It should either disallow l-QC splitting or strongly limit its scope in order to avoid time consuming processes and to control the amount of state information that must be stored at each border router.

5.2. AS path selection criteria

The core function of the utopian model consists in the selection of the AS path whose QoS characteristics are the closest to the requested QoS. This requires the huge knowledge of all possible AS path combinations with their corresponding QoS characteristics (i.e. all possible e-QCs between the source and destination). This selection could be based on the definition of formal rules (e.g. the definition of a distance between two {D, J, L} vectors) and on the customers' preferences.

The AS path selection criteria are the basic criteria used by the inter-domain routing process in order to select the AS path. In current non-QoS IP inter-domain routing the main criterion is the smallest number of ASs. Any viable solution should explicitly state on what criteria it selects an AS path. It should carefully evaluate what knowledge it needs and how it can retrieve the appropriate information.

5.3. Impact on routing

Since a QoS-enabled AS path is likely to be different from the best effort path, it does not rely on the traditional BGP (Border Gateway Protocol) [10] choice to build routing tables. Making use of BGP would indeed require some extensions [11]. The upstream AS must know the next selected downstream AS and the exit interface leading to it. With an obvious problem of scalability, this could be achieved thanks to an internal mapping table from e-QC label to egress interface. This label would be tagged in each packet (see l-QC splitting paragraph). This table could be populated by a centralised management entity.

Two packets bearing the same destination address may have to be routed towards two different egress points if they are carrying traffic of different eQCs. Therefore we cannot solely rely on the traditional intra-domain routing process. In the worst case, two datagrams entering an AS with the same (source address, destination address) pair and the same DSCP, may be routed to two different egress points. This is the case when they are originated from two sessions of the same application invoking two different end-to-end QoS, assigned to two different e-QC, that happen to use the same l-QC in the given AS but transiting two different downstream ASs. With an obvious problem of scalability, the intra-domain routing could be achieved thanks to an IGP (Interior Gateway Protocol) that builds in each router an internal mapping table from the aforementioned e-QC label to the next router.

Routing is the core process of any IP network. It has proven to be extremely efficient so as to allow the widespread uptake of the Internet. A particular care must be put to preserve its ease of use and its scalability. Addition of QoS capabilities to the Internet should have very limited impact on existing routing. This appears to be one of the most challenging issues for the success of a massive deployment of QoS-based services. We define inter-domain routing as the process that selects for each datagram the appropriate AS egress point. We define intra-domain routing as the process that selects the route (the suite of routers) to convey each datagram from the AS ingress point to the AS egress point selected by the inter-domain routing process. It should be noted that this latter definition applies to a transit AS only and therefore it is a restricted definition of intra-domain routing. Any viable model should explicitly state what protocol and mechanisms it uses for inter-domain and intra-domain routing. It should clearly state if it has any impact on the existing BGP protocol [12]. It should clearly state if it has any impact on IGP protocols already deployed within ASs.

6. Utopian approach limitations and future work

It is obvious that the utopian approach suffers from some shortcomings, since it does not address certain issues and behaviours like bandwidth reservation, multicast issues and security considerations. The reader will find some additional considerations on these topics in [13].

The utopian approach cannot address issues that are on the management or on the business plane, like accounting and charging or business relationship between service entities. We recognise, of course, these issues to be of primary importance for a Service Provider. The utopian approach, because it is an engineering approach, simply does not apply. We firmly believe, however, that this packet transportation approach is also of primary importance. Indeed, it has no meaning to construct a clever management and business model, if underneath, there are no real QoS network capacities working to transmit the users' packets from one end-station to another end-station.

Moreover, if this paper has limited itself to the description of issues and guidelines, the work conducted in MESCAL has gone a step further by providing some possible viable approaches to these issues. Next steps carried out within MESCAL are detailed description, specification and demonstration of a viable model.

7. Conclusion

In this paper we have presented a utopian approach as a tool for investigating the issues of end-to-end QoS delivery in the Internet at large scale. We have emphasised on the e-QC computation and AS path selection complexity and scalability. We have introduced the new notion of l-QC splitting. We have highlighted the problems that arise when confronted with l-QC splitting, AS path selection criteria and inter- and intra-domain routing. We have provided guidelines for each of these major issues. This work has the intention of providing guidance for future work in this domain and particularly for the design of viable models, which address requirements expressed by the providers and customers involved in concrete business practices and with well-defined business relationships.

8. Acknowledgements

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] G. Almes, S. Kalidindi, M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.

[2] C. Demichelis, P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.

[3] D. Goderis, D. Griffin, C. Jacquenet, G. Pavlou, " Attributes of a Service Level Specification (SLS) Template ", <draft-tequila-sls-03.txt>, Oct. 2003, work in progress.

[4] G. Almes, S. Kalidindi, M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.

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

[6] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998

[7] M. R. Garey, D. S. Johnson, "Computers and Intractability - A Guide to the Theory of NP-Completeness", Freeman, California, USA, 1979.

[8] Z. Wang, J. Crowcroft "Quality of Service Routing for Supporting Multimedia Applications", IEEE Journal of Selected Areas in Communications (JSAC), 1996.

[9] J. Rajahalme, A. Conta, B. Carpenter, S. Deering, "IPv6 Flow Label Specification", draft-ietf-ipv6-flow-label-07.txt, April 2003, work in progress

[10] Y. Rekhter, T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.

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

[12] L. Xiao, K.-S. Lui, J. Wang, K. Nahrstedt, "QoS extension to BGP", 10th IEEE International Conference on Network Protocols (ICNP'02), November 2002.

[13] Paris Flegkas, et al., "D1.1: Specification of Business Models and a Functional Architecture for Inter-domain QoS Delivery", http://www.mescal.org