

Network Working Group
Request for Comments: 5160
Category: Informational

P. Levis
M. Boucadair
France Telecom
March 2008

Considerations of Provider-to-Provider Agreements
for Internet-Scale Quality of Service (QoS)

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

IESG Note

This RFC is not a candidate for any level of Internet Standard. The IETF disclaims any knowledge of the fitness of this RFC for any purpose and notes that the decision to publish is not based on IETF review apart from IESG review for conflict with IETF work. The RFC Editor has chosen to publish this document at its discretion. See RFC 3932 for more information.

Abstract

This memo analyzes provider-to-provider Quality of Service (QoS) agreements suitable for a global QoS-enabled Internet. It defines terminology relevant to inter-domain QoS models. It proposes a new concept denoted by Meta-QoS-Class (MQC). This concept could potentially drive and federate the way QoS inter-domain relationships are built between providers. It opens up new perspectives for a QoS-enabled Internet that retains, as much as possible, the openness of the existing best-effort Internet.

Table of Contents

1. Introduction	2
2. Assumptions and Requirements	3
3. Terminology	4
4. Weaknesses of Provider-to-Provider QoS Agreements Based on SP Chains	5
4.1. IP Connectivity Services	6
4.2. Similarity between Provider and Customer Agreements	6
4.3. Liability for Service Disruption	7
4.4. SP Chain Trap Leading to Glaciation	7
5. Provider-to-Provider Agreements Based on Meta-QoS-Class	7
5.1. Single Domain Covering	8
5.2. Binding l-QCs	9
5.3. MQC-Based Binding Process	10
6. The Internet as MQC Planes	12
7. Towards End-to-End QoS Services	13
8. Security Considerations	15
9. Acknowledgements	15
10. References	16
10.1. Normative References	16
10.2. Informative References	16

1. Introduction

Three different areas can be distinguished in IP QoS service offerings. The first area is the single domain where a provider delivers QoS services inside the boundaries of its own network. The second area is multiple domains where a small set of providers, with mutual business interests, cooperate to deliver QoS services inside the boundaries of their network aggregate. The third area, which has very seldom been put forward, is the Internet where QoS services can be delivered from almost any source to any destination. Both multiple domains and Internet areas deal with inter-domain aspects. However, they differ significantly in many ways, such as the number of domains and QoS paths involved, which are much higher and dynamic for the Internet area. Multiple domains and Internet areas are therefore likely to differ in their respective solutions. This memo is an attempt to investigate the Internet area from the point of view of provider-to-provider agreements. It provides a framework for inter-domain QoS-enabled Internet.

[MESCAL] provides a set of requirements to be met by any solution aiming to solve inter-domain QoS issues. These requirements are not reproduced within this memo. Readers are invited to refer to [MESCAL] for more elaborated description on the requirements.

This memo shows that for the sake of scalability, providers need not be concerned with what occurs more than one hop away (from their Autonomous System) when they negotiate inter-domain QoS agreements. They should base their agreements on nothing but their local QoS capabilities and those of their direct neighbors. This analysis leads us to define terminology relevant to inter-domain QoS models. We also introduce a new concept denoted by Meta-QoS-Class (MQC) that drives and federates the way QoS inter-domain relationships are built between providers. The rationale for the MQC concept relies on a universal and common understanding of QoS-sensitive applications needs. Wherever end-users are connected, they experience the same QoS difficulties and are likely to express very similar QoS requirements to their respective providers. Globally confronted with the same customer requirements, providers are likely to design and operate similar network capabilities.

MQC brings up a simplified view of the Internet QoS capabilities as a set of MQC planes. This memo looks at whether the idea of MQC planes can be helpful in certain well-known concrete inter-domain QoS issues. The focus, however, is on the provider-to-provider QoS agreement framework, and the intention is not to specify individual solutions and protocols for a full inter-domain QoS system. For discussion of a complete architecture based on the notion of parallel Internets that extends and generalizes the notion of MQC planes, see [AGAVE].

Note that this document does not specify any protocols or systems.

2. Assumptions and Requirements

To avoid a great deal of complexity and scalability issues, we assume that provider-to-provider QoS agreements are negotiated only for two adjacent domains that are directly accessible to each other. We also assume, because they exchange traffic, that these neighbors are BGP [RFC4271] peers. This pairwise peering is logical, therefore it can be supported not only on physical point-to-point connections but also on Internet exchange points (IXPs), where many operators connect to each other using a layer 2 switch.

The QoS solutions envisaged in this document are exclusively solutions suitable for the global Internet. As far as Internet-wide solutions are concerned, this document assumes that:

- o Any solutions should apply locally in order to be usable as soon as deployed in a small set of domains.

- o Any solutions should be scalable in order to allow a global deployment to almost all Internet domains, with the ability to establish QoS communications between any and all end-users.
- o Any solutions should also maintain a best-effort service that should remain the preeminent service as a consequence of the end-to-end argument [E2E].
- o If there is no path available within the requested QoS to reach a destination, this destination must remain reachable through the best-effort service.

This memo does not place any specific requirements on the intra-domain traffic engineering policies and the way they are enforced. A provider may deploy any technique to ensure QoS inside its own network. This memo assumes only that QoS capabilities inside a provider's network can be represented as local-QoS-Classes (l-QCs). When crossing a domain, traffic experiences conditions characterized by the values of delay, jitter, and packet loss rate that correspond to the l-QC selected for that traffic within that domain. 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.

3. Terminology

(D, J, L)

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

Domain

A network infrastructure composed of one or several Autonomous Systems (AS) managed by a single administrative entity.

IP connectivity service

IP transfer capability characterized by a (Destination, D, J, L) tuple where Destination is a group of IP addresses and (D, J, L) is the QoS performance to get to the Destination.

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].

L-QC binding

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.

L-QC thread

Chain of neighboring bound l-QCs.

Meta-QoS-Class (MQC)

An MQC provides the limits of the QoS parameter values that two l-QCs must respect in order to be bound together. An MQC is used as a label that certifies the support of a set of applications that bear similar network QoS requirements.

Service Provider (SP)

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

SP chain

The chain of Service Providers whose domains are used to convey packets for a given IP connectivity service.

4. Weaknesses of Provider-to-Provider QoS Agreements Based on SP Chains

The objective of this section is to show, by a sort of *reductio ad absurdum* proof, that within the scope of Internet services, provider-to-provider QoS agreements should be based on guarantees that span a single domain.

We therefore analyze provider-to-provider QoS agreements based on guarantees that span several domains and emphasize their vulnerabilities. In this case, the basic service element that a provider offers to its neighboring providers is called an IP connectivity service. It uses the notion of SP chains. We first define what an IP connectivity service is, and then we point out

several weaknesses of such an approach, especially the SP chain trap problem that leads to the so-called Internet glaciation era.

4.1. IP Connectivity Services

An IP connectivity service is a (Destination, D, J, L) tuple where Destination is a group of IP addresses reachable from the domain of the provider offering the service, and (D, J, L) is the QoS performance to get from this domain to Destination. Destination is typically located in a remote domain.

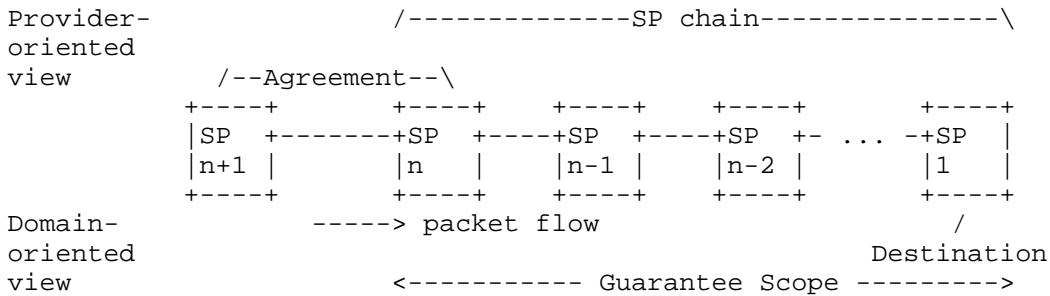


Figure 1: IP connectivity service

In Figure 1, Provider SP_n guarantees provider SP_{n+1} the level of QoS for crossing the whole chain of providers' domains (SP_n, SP_{n-1}, SP_{n-2}, ..., SP₁). SP_i denotes a provider as well as its domain. The top of the figure is the provider-oriented view; the ordered set of providers (SP_n, SP_{n-1}, SP_{n-2}, ..., SP₁) is called an SP chain. The bottom of the figure is the domain-oriented view.

4.2. Similarity between Provider and Customer Agreements

This approach maps end-users' needs directly to provider-to-provider agreements. Providers negotiate agreements to a destination because they know customers are ready to pay for QoS guaranteed transfer to this destination. As far as service scope is concerned, the agreements between providers will resemble the agreements between customers and providers. For instance, in Figure 1, SP_n can sell to its own customers the same IP connectivity service it sells to SP_{n+1}. There is no clear distinction between provider-to-provider agreements and customer-to-provider agreements.

In order to guarantee a stable service, redundant SP chains should be formed to reach the same destination. When one SP chain becomes unavailable, an alternative SP chain should be selected. In the context of a global QoS Internet, that would lead to an enormous number of SP chains along with the associated agreements.

4.3. Liability for Service Disruption

In Figure 1, if SPn+1 sees a disruption in the IP connectivity service, it can turn only against SPn, its legal partner in the agreement. If SPn is not responsible, in the same way, it can only complain to SPn-1, and so on, until the faulty provider is found and eventually requested to pay for the service impairment. The claim is then supposed to move back along the chain until SPn pays SPn+1. The SP chain becomes a liability chain.

Unfortunately, this process is prone to failure in many cases. In the context of QoS solutions suited for the Internet, SP chains are likely to be dynamic and involve a significant number of providers. Providers (that do not all know each other) involved in the same SP chain can be competitors in many fields; therefore, trust relationships are very difficult to build. Many complex and critical issues need to be resolved: how will SPn+1 prove to SPn that the QoS level is not the level expected for a scope that can expand well beyond SPn? How long will it take to find the guilty domain? Is SPn ready to pay SPn+1 for something it does not control and is not responsible for?

4.4. SP Chain Trap Leading to Glaciation

In Figure 1, SPn implicitly guarantees SPn+1 the level of QoS for the crossing of distant domains like SPn-2. As we saw in Section 4.2, SP chains will proliferate. A provider is, in this context, likely to be part of numerous SP chains. It will see the level of QoS it provides guaranteed by many providers it perhaps has never even heard of.

Any change 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 large number of external (to its domain) agreements that depend on it. This is what is referred to as the "SP chain trap" issue. This solution is not appropriate for worldwide QoS coverage, as it would lead to glaciation phenomena, causing a completely petrified QoS infrastructure, where nobody could renegotiate any agreement.

5. Provider-to-Provider Agreements Based on Meta-QoS-Class

If a QoS-enabled Internet is deemed desirable, with QoS services potentially available to and from any destination, any solution must resolve the above weaknesses and scalability problems and find alternate schemes for provider-to-provider agreements.

5.1. Single Domain Covering

Due to the vulnerabilities of the SP chain approach, we assume provider-to-provider QoS agreements should be based on guarantees covering a single domain. A provider guarantees its neighbors only the crossing performance of its own domain. In Figure 2, the provider SPn guarantees the provider SPn+1 only the QoS performance of the SPn domain. The remainder of this document will show that this approach, bringing clarity and simplicity into inter-domain relationships, is better suited for a global QoS Internet than one based on SP chains.

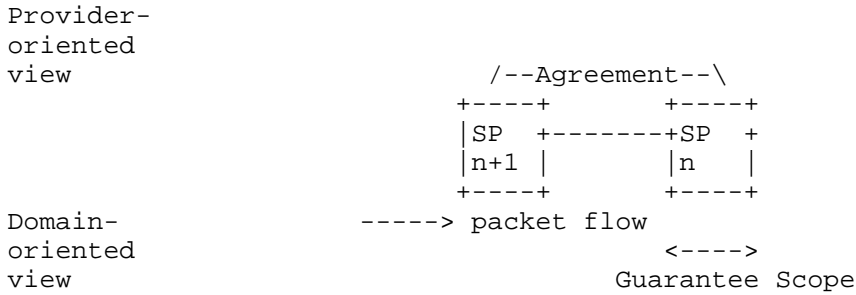


Figure 2: provider-to-provider QoS agreement

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 end-to-end guarantees for communications.

The simple fact that SP chains do not exist makes the AS chain trap problem and the associated glaciation threat vanish.

The liability issue is restricted to a one-hop distance. A provider is responsible for its own domain only, and is controlled by all the neighbors with whom it has a direct contract.

5.2. Binding l-QCs

When a provider wants to contract with another provider, the main concern is deciding which l-QC(s) in its own domain it will bind to which l-QC(s) in the neighboring downstream domain. The l-QC binding process becomes the basic inter-domain process.

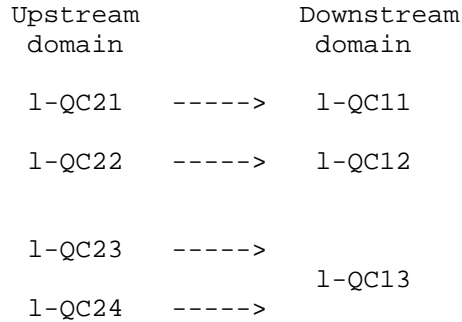


Figure 3: l-QC Binding

If one l-QC were to be bound to two (or more) l-QCs, it would be very difficult to know which l-QC the packets should select. This could imply a flow classification at the border of the domains based on granularity as fine as the application flows. For the sake of scalability, we assume one l-QC should not be bound to several l-QCs [Lev]. On the contrary, several l-QCs can be bound to the same l-QC, in the way that l-QC23 and l-QC24 are bound to l-QC13 in Figure 3.

A provider decides the best match between l-QCs based exclusively on:

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

It does not use any information related to what is happening more than one domain away.

Despite this one-hop, short-sighted approach, the consistency and the coherency of the QoS treatment must be ensured on an l-QC thread formed by neighboring bound l-QCs. Packets leaving a domain that applies a given l-QC should experience similar treatment when crossing external domains up to their final destination. A provider should bind its l-QC with the neighboring l-QC that has the closest performance. The criteria for l-QC binding should be stable along any l-QC thread. For example, two providers should not bind two l-QCs to minimize the delay whereas further on, on the same thread, two other providers have bound two l-QCs to minimize errors.

Constraints should be put on l-QC QoS performance parameters to confine their values to an acceptable and expected level on an l-QC thread scale. These constraints should depend on domain size; for example, restrictions on delay should authorize a bigger value for a national domain than for a regional one. Some rules must therefore be defined to establish in which conditions two l-QCs can be bound together. These rules are provided by the notion of Meta-QoS-Class (MQC).

5.3. MQC-Based Binding Process

An MQC provides the limits of the QoS parameters two l-QCs must respect in order to be bound together. A provider goes through several steps to extend its internal l-QCs through the binding process. Firstly, it classifies its own l-QCs based on MQCs. An MQC is used as a label that certifies the support of a set of applications that bear similar network QoS requirements. It is a means to make sure that an l-QC has the appropriate QoS characteristics to convey the traffic of this set of applications. Secondly, 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 benefit from it. Thirdly, it contracts an agreement with its neighbor to send some traffic that will be handled according to the agreed MQCs.

The following attributes should be documented in any specification of an MQC. This is not a closed list, other attributes can be added if needed.

- o A set of applications (e.g., VoIP) the MQC is particularly suited for.
- o Boundaries or intervals of a set of QoS performance attributes whenever required. For illustration purposes, we propose to use in this document attribute (D, J, L) 3-tuple. An MQC could be focused on a single parameter (e.g., suitable to convey delay sensitive traffic). Several levels could also be specified depending on the size of the network provider; for instance, a small domain (e.g., regional) needs lower delay than a large domain (e.g., national) to match a given MQC.
- o Constraints on traffic (e.g., only TCP-friendly).
- o Constraints on the ratio: network resources for the class / overall traffic using this class (e.g., less resources than peak traffic).

Two l-QCs can be bound together if, and only if, they conform to the same MQC.

Provider-to-provider agreements, as defined here, are uni-directional. They are established for transporting traffic in a given direction. However, from a business perspective, it is likely that reverse agreements will also be negotiated for transporting traffic in the opposite direction. Note that uni-directional provider-to-provider agreements do not preclude having end-to-end services with bi-directional guarantees, when you consider the two directions of the traffic separately.

Two providers negotiating an agreement based on MQC will have to agree on several other parameters that are outside the definition of MQC. One such obvious parameter is bandwidth. The two providers agree to exchange up to a given level of QoS traffic. This bandwidth level can then be further renegotiated, inside the same MQC agreement, to reflect an increase in the end-user QoS requests. Other clauses of inter-domain agreements could cover availability of the service, time of repair, etc.

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 business). A given l-QC can be suitable for several MQCs (even outside the same hierarchy). Several l-QCs in the same domain can be classified as belonging to the same MQC. There is an MQC with no specific constraints called the best-effort MQC.

There is a need for some form of standardization to control QoS agreements between providers [RFC3387]. Each provider must have the same understanding of what a given MQC is about. We need a global agreement on a set of MQC standards. The number of classes to be defined must remain very small to avoid overwhelming complexity. We also need a means to certify that the l-QC classification made by a provider conforms to the MQC standards. So the standardization effort should be accompanied by investigations on conformance testing requirements.

The three notions of PDB, Service Class [RFC4594], and MQC are related; what MQC brings is the inter-domain aspect:

- PDB is how to engineer a network;
- Service Class is a set of traffic with specific QoS requests;
- MQC is a way to classify the QoS capabilities (l-QCs, through Diffserv or any other protocols or mechanisms) in order to reach agreements with neighbors. MQCs are only involved when a provider

wants to negotiate an agreement with a neighboring provider. MQC is completely indifferent to the way networks are engineered as long as the MQC QoS attribute (e.g., (D, J, L)) values are reached.

6. The Internet as MQC Planes

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 according to the same MQC. An MQC plane can have holes and isolated domains because QoS capabilities do not cover all Internet domains. When an l-QC maps to several MQCs, it belongs potentially to several planes.

When a provider contracts with another provider based on the use of MQCs, it simply adds a logical link to the corresponding MQC plane. This is basically what current traditional inter-domain agreements mean for the existing Internet. Figure 4a) depicts the physical layout of a fraction of the Internet, comprising four domains with full-mesh connectivity.

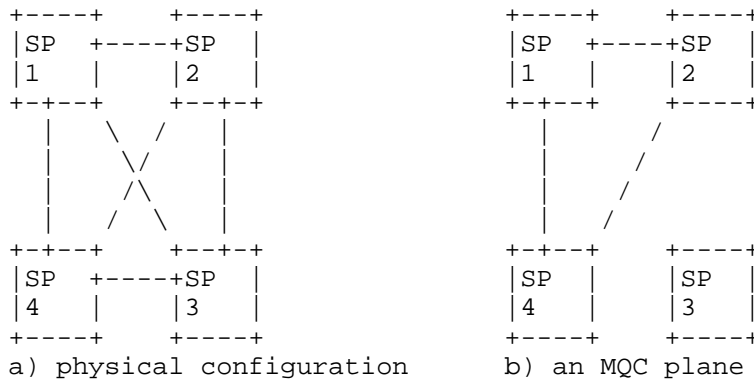


Figure 4: MQC planes

Figure 4 b) depicts how these four domains are involved in a given MQC plane. SP1, SP2, and SP4 have at least one compliant l-QC for this MQC; SP3 may or may not have one. A bi-directional agreement exists between SP1 and SP2, SP1 and SP4, SP2 and SP4.

MQC brings a clear distinction between provider-to-provider and customer-to-provider QoS agreements. We expect a great deal of difference in dynamicity between the two. Most provider-to-provider agreements should have been negotiated, and should remain stable, before end-users can dynamically request end-to-end guarantees. Provider agreements do not directly map end-users' needs; therefore, the number of provider agreements is largely independent of the

number of end-user requests and does not increase as dramatically as with SP chains.

For a global QoS-based Internet, this solution will work only if MQC-based binding is largely accepted and becomes a current practice. 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 can be extended. Therefore, incremental deployment is possible, and leads to incremental benefits. For example, in Figure 4 b), as soon as SP3 connects to the MQC plane it will be able to benefit from the SP1, SP2, and SP4 QoS capabilities.

The Internet, as a split of different MQC planes, offers an ordered and simplified view of the Internet QoS capabilities. End-users can select the MQC plane that is the closest to their needs, as long as there is a path available for the destination. One of the main outcomes of applying the MQC concept is that it alleviates the complexity and the management burden of inter-domain relationships.

7. Towards End-to-End QoS Services

Building end-to-end QoS paths, for the purpose of QoS-guaranteed communications between end-users, is going a step further in the QoS process. The full description of customer-to-provider QoS agreements, and the way they are enforced, is outside the scope of this memo. However, in this section, we will list some important issues and state whether MQC can help to find solutions.

Route path selection within a selected MQC plane can be envisaged in the same way as the traditional routing system used by Internet routers. Thus, we can rely on the BGP protocol, basically one BGP occurrence per MQC plane, for the inter-domain routing process. The resilience of the IP routing system is preserved. If a QoS path breaks somewhere, the routing protocol enables dynamic computation of another QoS path, if available, in the proper MQC plane. This provides a first level of QoS infrastructure that could be conveniently named "best-effort QoS", even if this is an oxymoron.

On this basis, features can be added in order to select and control the QoS paths better. For example, BGP can be used to convey QoS-related information, and to implement a process where Autonomous Systems add their own QoS values (D, J, L) when they propagate an AS path. Then, the AS path information is coupled with the information on Delay, Jitter, and Loss, and the decision whether or not to use the path selected by BGP can be made, based on numerical values. For

example, for destination N, an AS path (X, Y) is advertised to AS Z. During the propagation of this AS path by BGP, X adds the information concerning its own delay, say 30 ms, and Y adds the information concerning its own delay, say 20 ms. Z receives the BGP advertisement (X, Y, N, 50 ms). One of Z's customers requests a delay of 100 ms to reach N. Z knows its own delay for this customer, say 20 ms. Z computes the expected maximum delay from its customer to N: 70 ms, and concludes that it can use the AS path (X, Y). The QoS value of an AS path could also be disconnected from BGP and computed via an off-line process.

If we use QoS routing, we can incorporate the (D, J, L) information in the BGP decision process, but that raises the issue of composing performance metrics in order to select appropriate paths [Chau]. When confronted by multiple incompatible objectives, the local decisions made to optimize the targeted parameters could give rise to a set of incomparable paths, where no path is strictly superior to the others. The existence of provider-to-provider agreements based on MQC offers a homogenous view of the QoS parameters, and should therefore bring coherency, and restrict the risk of such non-optimal choices.

A lot of end-to-end services are bi-directional, so one must measure the composite performance in both directions. Many inter-domain paths are asymmetric, and usually, some providers involved in the forward path are not in the reverse path, and vice versa. That means that no assumptions can be made about the reverse path. Although MQC-based provider-to-provider agreements are likely to be negotiated bi-directionally, they allow the inter-domain routing protocol to compute the forward and the reverse paths separately, as usual. The only constraint MQC puts on routing is that the selected paths must use the chosen MQCs throughout the selected domains. There might be a different MQC requirement in the reverse direction than in the forward direction. To address this problem, we can use application-level communication between the two parties (customers) involved in order to specify the QoS requirements in both directions.

We can go a step further in the control of the path to ensure the stability of QoS parameters such as, e.g., enforcing an explicit routing scheme, making use of RSVP-TE/MPLS-TE requests [RFC3209], before injecting the traffic into an l-QC thread. However, currently, several problems must be resolved before ready and operational solutions for inter-domain route pinning, inter-domain TE, fast failover, and so forth, are available. For example, see the BGP slow convergence problem in [Kushman].

Multicast supports many applications such as audio and video distribution (e.g., IPTV, streaming applications) with QoS

requirements. Along with solutions at the IP or Application level, such as Forward Error Correction (FEC), the inter-domain multicast routing protocol with Multiprotocol Extensions for BGP-4 [RFC4760], could be used to advertise MQC capabilities for multicast source reachability. If an inter-domain tree that spans several domains remains in the same MQC plane, it would be possible to benefit from the consistency and the coherency conferred by MQC.

Note that the use of some QoS parameters to drive the route selection process within an MQC plane may induce QoS deterioration since the best QoS-inferred path will be selected by all Autonomous System Border Routers (ASBRs) involved in the inter-domain path computation (i.e., no other available routes in the same MQC plane will have a chance to be selected). This problem was called the QoS Attribute-rush (QA-rush) in [Grif]. This drawback may be avoided if all involved ASes in the QoS chain implement some out-of-band means to control the inter-domain QoS path consistency (MQC compliance).

To conclude this section, whatever the protocols we want to use, and however tightly we want to control QoS paths, MQC is a concept that can help to resolve problems [WIP], without prohibiting the use of any particular mechanism or protocol at the data, control, or management planes.

8. Security Considerations

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

Particular attention should be paid to giving access to MQC resources only to authorized traffic. Unauthorized access can lead to denial of service when the network resources get overused [RFC3869].

9. Acknowledgements

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) and AGAVE (A lightweight Approach for Viable End-to-end IP-based QoS Services) projects. The authors would like to thank all the other partners for the fruitful discussions.

We are grateful to Brian Carpenter, Jon Crowcroft, and Juergen Quittek for their helpful comments and suggestions for improving this document.

10. References

10.1. Normative References

- [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.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

10.2. Informative References

- [AGAVE] Boucadair, et al., "Parallel Internets Framework", IST AGAVE project public deliverable D1.1, September 2006.
- [Chau] Chau, C., "Policy-based routing with non-strict preferences", Proceedings of the ACM SIGCOMM 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Pisa, Italy, pp 387-398, September 2006.
- [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.
- [Grif] Griffin, D., Spencer, J., Griem, J., Boucadair, M., Morand, P., Howarth, M., Wang, N., Pavlou, G., Asgari, A., and P. Georgatsos, "Interdomain routing through QoS-class planes [Quality-of-Service-Based Routing Algorithms for Heterogeneous Networks]", IEEE Communications Magazine, Vol 45, Issue 2, pp 88-95, February 2007.
- [Kushman] Kushman, N., Kandula, S., and D. Katabi, "Can You Hear Me Now?! It Must Be BGP", ACM Journal of Computer and Communication Review CCR, November 2007.
- [Lev] Levis, P., Asgari, H., and P. Trimintzios, "Consideration on Inter-domain QoS and Traffic Engineering issues Through a Utopian Approach", SAPIR-2004 workshop of ICT-2004, (C) Springer-Verlag, August 2004.

- [MESCAL] Flegkas, et al., "Specification of Business Models and a Functional Architecture for Inter-domain QoS Delivery", IST MESCAL project public deliverable D1.1, May 2003.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [RFC3086] Nichols, K. and B. Carpenter, "Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification", RFC 3086, April 2001.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC3387] Eder, M., Chaskar, H., and S. Nag, "Considerations from the Service Management Research Group (SMRG) on Quality of Service (QoS) in the IP Network", RFC 3387, September 2002.
- [RFC3869] Atkinson, R., Ed., Floyd, S., Ed., and Internet Architecture Board, "IAB Concerns and Recommendations Regarding Internet Research and Evolution", RFC 3869, August 2004.
- [RFC4594] Babiarez, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", RFC 4594, August 2006.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [WIP] Deleuze, G. and F. Guattari, "What Is Philosophy?", Columbia University Press ISBN: 0231079893, April 1996.

Authors' Addresses

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

EEmail: pierre.levis@orange-ftgroup.com

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

EEmail: mohamed.boucadair@orange-ftgroup.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78 and at <http://www.rfc-editor.org/copyright.html>, and except as set forth therein, the authors retain all their rights.

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, THE IETF TRUST 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.

Intellectual Property

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.

