

Network Working Group

M. Boucadair

(Editor)

Internet Draft

France Telecom R&D

Document: [draft-boucadair-qos-bgp-spec-01.txt](#)

July 2005

Category: Experimental

QoS-Enhanced Border Gateway Protocol
< [draft-boucadair-qos-bgp-spec-01.txt](#) >

Status of this Memo

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

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress"

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

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

This Internet-Draft will expire on January 2006.

Abstract

This memo describes a proposal for exchanging QoS-enabled reachability information between service providers. It defines new BGP attributes to be used in order to convey QoS performance characteristics between service providers and it describes how to use these QoS values in order to select the best path. An example of route selection process that takes into account the QoS performance values enclosed in BGP messages is also detailed.

Table of Contents

1.	Introduction.....	2
2.	Conventions used in this document.....	3
3.	Terminology.....	3
4.	Inter domain QoS delivery services taxonomy.....	3
5.	q-BGP requirements and behaviors.....	4
5.1.	Requirements.....	4
5.2.	q-BGP behaviors.....	5
5.3.	How to set QoS-related information carried in q-BGP messages?.....	5
6.	q-BGP specifications.....	5
6.1.	Attributes.....	6
6.1.1.	QoS Service Capability.....	6
6.1.2.	QoS_NLRI attribute.....	7
6.1.3.	QoS_NLRI attribute for Group-1.....	7
6.1.4.	QoS_NLRI attribute for Group-2.....	9
6.1.5.	Multiple paths.....	11
6.2.	Processing q-BGP attributes.....	11
6.3.	q-BGP Route Selection Process.....	12
6.3.1.	Group-1 Route Selection Process.....	12
6.3.2.	Group-2 Route Selection Process.....	13
6.3.2.1.	QoS comparison logic.....	14
6.3.2.2.	Priority-based route selection process.....	14
6.3.2.2.1.	Processing of route with QoS holes.....	16
7.	Security Considerations.....	18
8.	References.....	18
9.	Acknowledgments.....	19
10.	Contributors.....	19
11.	Author's Address.....	19

[1.](#) Introduction

QoS delivery services are seen as critical future Internet services [IAB]. The introduction of such services requires updating network infrastructures, including network devices capabilities, protocols, and management tools, with appropriate technologies. From a reachability perspective, modifications need to be brought to existing signaling and routing protocols in order to convey richer information in addition to best effort one. In other words, reachability information should provide routers with pertinent

information to help the route selection decision-making process. Such information could be for instance the QoS that will be experienced along a given path. Therefore, Network Providers have to evolve and update the protocols deployed in their domains in order to meet the new requirements and then to be able to offer new sophisticated added value services. As far as inter-domain is concerned, the issue is crucial since all providers should deploy means to convey QoS-related information between their domains so that QoS-based services could

have a world-wide scope and then be accessible for a large set of customers.

The Border Gateway Protocol (BGP) protocol [BGP] is the de facto protocol deployed by network providers in order to interconnect their autonomous systems with the ones of their peers. This memo describes how BGP could be used as a means to convey QoS-related information between adjacent autonomous systems and then allow deployment of new inter-domain QoS delivery services. The proposed solution is generic and could be applied to any kind of inter-domain QoS delivery solution that is based on the exchange of QoS-related information.

[2.](#) Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [RFC2119].

[3.](#) Terminology

This memo makes use of the following terms:

- o Local QoS Class: A QoS transfer capability within a single SP domain, characterized by a set of QoS performance parameters.
- o Extended QC: A QoS transfer capability provided using both the local domain and neighbor domains. An extended QC is provided by combining the local QC of local domain with the ones of adjacent domains. The topological scope of an extended QC extends the boundaries of local domain.
- o Domain: within this document it denotes an Autonomous System

- o q-BGP (QoS-enhanced BGP): an enhanced BGP that takes into account QoS information it carries in its messages as an input to its route selection process.

[4.](#) Inter domain QoS delivery services taxonomy

Internet is a collection of domains interconnected together and that exchange reachability information between them. In order to introduce inter domain QoS delivery services, providers should exchange QoS information in addition to best effort one. This information will be used as an input for the route selection process and will guide the selection of optimal paths that will be used to carry traffic belonging to inter-domain QoS services. The QoS-related information exchange occurs either at the service level or at the routing level.

Boucadair

Experimental - Expires January 2006

[Page 3]

Internet Draft

QoS-enhanced Border Gateway Protocol

July 2005

Two groups of QoS delivery services have been identified and are detailed hereafter:

- o The first group of services requires propagating only an identifier that has been agreed between adjacent peers. This identifier could be the DSCP value used to signal an extended QoS class. Additional QoS performance characteristics could be negotiated but not exchanged in the routing level. In the rest of this document, we will refer to this group as "group-1".
- o The second group requires the propagation of a set of QoS performance characteristics associated with an identifier. The nature of the QoS-related information to be exchanged has to be agreed between adjacent peers. We will refer to this group in the rest of this document as "group-2".

[5.](#) q-BGP requirements and behaviors

[5.1.](#) Requirements

As stated in the introduction, BGP is the inter domain routing protocol used to interconnect domains. This protocol is widely deployed and activated in a big range of network nodes. One of the risks to be taken into account when introducing new services like inter domain QoS ones is to preserve backward compatibility. Therefore, in order to ease introduction of these value added

services, it is recommended to reuse existing protocols and systems. Nevertheless, these protocols should be evolved and enhanced with additional features in order to achieve these new service objectives.

As far as inter domain routing is concerned, we will reuse BGP and define new features in order to convey QoS-related information. This information could only be reduced to a DS code point or be a set of QoS performance characteristics. In the rest of this document we will refer to this enhanced BGP as q-BGP (QoS-Enhanced BGP) protocol. When designing extensions to be added to classical BGP, the following requirements are to be taken into account:

- o The q-BGP protocol should, as far as possible, be able to operate independently of the inter-domain QoS delivery service it serves. Nevertheless, q-BGP should be able to detect the group it serves.
- o q-BGP should be able to propagate topology changes without any significant impact on the existing best-effort based network infrastructure;
- o q-BGP should be able to support all kind of services based on an exchange of QoS-related information especially to serve the two

groups described above.

[5.2.](#) q-BGP behaviors

q-BGP could have several behaviors depending on the nature of the QoS-related information carried by its messages. If q-BGP messages carry only a QC identifier (this identifier could be a DS code-point or a proprietary identifier), offline traffic engineering functions are certainly complex but the q-BGP route selection process complexity is reduced. This complexity increases when a set of QoS characteristics are associated with each QC identifier. The route selection process can use either the QC-identifier for all services that take part of group-1 or the QC-identifier and QoS performance characteristics for solutions belonging to group-2.

[5.3.](#) How to set QoS-related information carried in q-BGP messages?

q-BGP can carry QoS performance characteristics that could be

advantageously taken into account by the q-BGP route selection process to select an optimal path. This would enable to tune the route selection process in order to select routes according to more sophisticated routing policies (e.g route with highest available rate and lower delay). The QoS information inserted in q-BGP messages could be of different nature. It could be (1) administratively enforced. In that case it would not change too frequently. Or, it could (2) be much more dynamic (result of an active measurement for instance). In that case the frequency of changes could be much higher.

Administrative setting of QoS values could be achieved either statically or periodically. If these values are set statically, the behavior of q-BGP will be static and the route selection process will choose the same route. The QoS-related information doesn't bring major added-value to the final behavior of the route decision making process and freezes the state of the inter-domain routing. Nevertheless, in the case of QoS performance characteristics values are set periodically or dynamically, providers will deploy mechanisms that monitor the network and then guide the setting of these values. q-BGP will be provided by accurate information in order to select the optimal path. The frequency between two q-BGP router configuration operations in an administrative scheme should not be too small and could be very small in the dynamic scheme. In case of dynamic setting scheme, the risk is to impact routing table stability and probably introduce oscillation phenomena.

[6.](#) q-BGP specifications

In this section we detail the specification of q-BGP. This specification encloses new attributes introduced in order to convey QoS performance characteristics between distinct domains. Also, we

Boucadair Experimental - Expires January 2006 [Page 5]

Internet Draft QoS-enhanced Border Gateway Protocol July 2005

detail the processing of QoS_NLRI attribute and the use of QoS related information enclosed in BGP update message in order to select an optimal path towards a given prefix. A new capability attribute that allows negotiation of QoS service capability is also defined.

[6.1.](#) Attributes

[6.1.1.](#) QoS Service Capability

In order to prevent from BGP session closure due to reception of a

non supported optional parameter, IDR WG has adopted a mechanism called capability negotiation [CAP]. Capability exchange is achieved thanks to the specification of a new optional parameter.

As far as q-BGP is concerned, it is useful for a q-BGP peer to know the capabilities of a q-BGP neighbor with respect to the BGP protocol extensions. Thus, a new parameter is included in the optional parameters of the OPEN message of a q-BGP session to indicate that this peer supports q-BGP related attributes.

In order to indicate that a given inter-domain QoS delivery solution belongs to a given group (either group-1 or group-2) a new parameter called QoS Service Capability is introduced. A q-BGP speaker SHOULD use this capability advertisement in order to indicate the group to which an offered inter-domain QoS delivery solution belongs to, so that q-BGP peers can deduce if they can use the 'QoS service'-related attributes with a q-BGP speaker.

The fields of this optional parameter are set as follows:

- o The capability code field is set to a value between 128 and 255 as described in [CAP];
- o The capability length is set to 2;
- o The capability value field is encoded as shown in Figure below:




Figure 1. QoS service capability attribute.

- o The first octet (G1QS, Group-1 QoS Service) is set to 0xFF if an offered inter-domain QoS delivery solution that belongs to group-1 is supported;

- o The second octet (G2SQ, Group-2 QoS Service) is set to 0xFF if an offered inter-domain QoS delivery solution that belongs to group-2 is supported.

[6.1.2.](#) QoS_NLRI attribute

In order to convey QoS-related information, we adopt the [\[QOSNLRI\]](#) proposal that consists at introducing a new optional attribute called QOS_NLRI attribute as the starting point. This attribute is modified in order to support the requirements of two groups defined above. The format of the QoS_NLRI is different depending on the group (group-1 or group-2) it serves.

The modifications are twofold:

- o Information carried by this attribute:
 - o The [\[QOSNLRI\]](#) proposal allows to send only one QoS performance characteristic per announcement. This limitation has been relaxed within this specification since it might be necessary to carry a list of QoS performance characteristics in a single q-BGP UPDATE message;
 - o Unlike the [\[QOSNLRI\]](#) proposal, this specification allows to propagate information about extended QCs that are pre-negotiated between service peers;
 - o The [\[QOSNLRI\]](#) proposal adopts the multiple paths [\[Walton\]](#). The basic q-BGP specification focuses on single path;
 - o The PHB identifier has been removed from the list of possible "QoS Information Code" because of the existence of "QoS Class identifier".
- o The format of the QoS_NLRI attribute:
 - o Add a new field called "QoS Information length": the purpose of this field is to indicate the number of QoS performance characteristics that are enclosed in a q-BGP UPDATE message.
 - o The lengths of "QoS Information code" and "QoS Information Sub-code" have been reduced to 4 bits in order to reduce the total length of the QOS_NLRI attribute. This is also motivated by the fact that 2^4 values are sufficient to indicate this information.

[6.1.3.](#) QoS_NLRI attribute for Group-1

As described above, both group-1 and group-2 solutions need to exchange a QC identifier. This identifier is used to differentiate between the extended QCs. Especially this is the unique additional information that must be carried by q-BGP messages. Therefore, we define a new QoS_NLRI attribute for group-1 solutions. The format of this QoS_NLRI attribute is as follows:

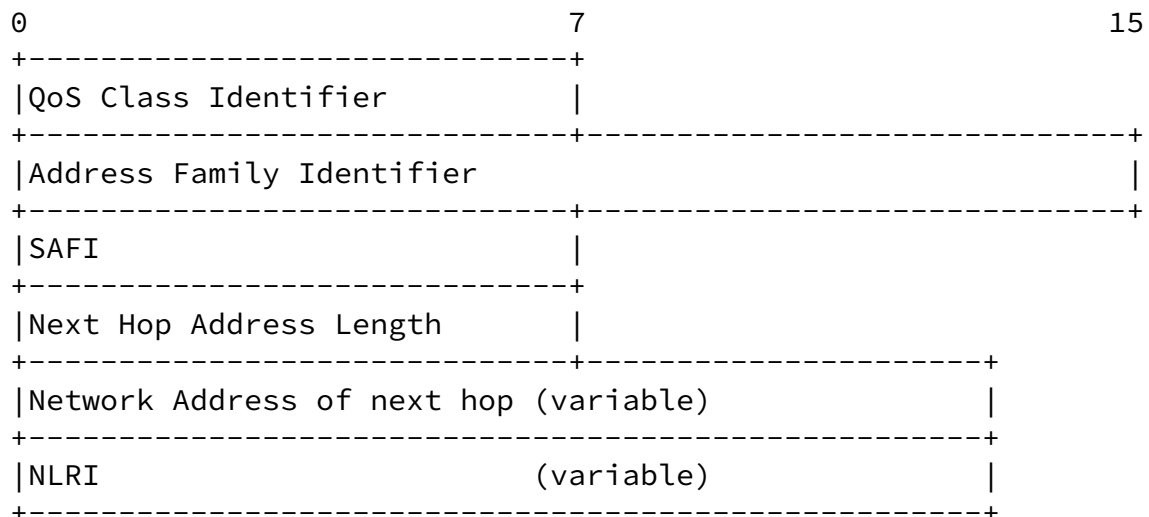


Figure 2. QoS NLRI attribute for Group-1.

The meaning of the fields of the group-1 QoS_NLRI attribute is as follows:

- o QoS class identifier (1 octet): this field indicates the QC identifier as described in [\[DS\]](#);
- o Address Family Identifier (AFI) (2 octet): this field carries the identity of the Network Layer protocol associated with the Network Address that follows;
- o Subsequent Address Family Identifier (SAFI) (1 octet): this field provides additional information about the type of the prefix carried in the QoS_NLRI attribute;
- o Next Hop Network address length (1 octet): the length of next hop network address;
- o Network address of Next Hop: this field contains the network address of the next router on the path to the destination prefix;
- o Network Layer Reachability Information: This variable length field lists the NLRI information for the feasible routes that

are being advertised by this attribute. The next hop information carried in the QOS_NLRI path attribute defines the Network Layer address of the border router that should be used as the next hop

to the destinations listed in the QOS_NLRI attribute in the UPDATE message.

[6.1.4.](#) QoS_NLRI attribute for Group-2

As described above, group-2 solutions need to exchange both QC identifier and a set of QoS performance characteristics. Therefore, we define a new QoS NLRI attribute that carry several QoS values for a given extended QC. The format of this new attribute is as follows:

0	3	7
+-----+-----+		
QoS Information length		
+-----+-----+		
QoS Information Code	QoS Information Sub-Code	
+-----+-----+		
QoS Information value		
+		+
+-----+-----+		
QoS Information length		
+-----+-----+		
QoS Information Code	QoS Information Sub-Code	
+-----+-----+		
QoS Information value		
+		+
+-----+-----+		
.		
+-----+-----+		
QoS Information length		
+-----+-----+		
QoS Information Code	QoS Information Sub-Code	
+-----+-----+		
QoS Information value		
+		+

+-----+	+-----+	
QoS Class Identifier		
+-----+	+-----+	
Address Family Identifier		
+-----+	+-----+	
SAFI		
+-----+	+-----+	
Next Hop Address Length		
+-----+	+-----+	
Network Address of next hop (variable)		

Boucadair

Experimental - Expires January 2006

[Page 9]

Internet Draft

QoS-enhanced Border Gateway Protocol

July 2005

+-----+		
NLRI	(variable)	
+-----+		

Figure 3. QoS NLRI attribute for Group-1.

The meaning of the fields of the group-2 QOS_NLRI attribute is defined below:

- o QoS information length: this field carries the number of the QoS information Code that will be sent by the q-BGP speaker in a single q-BGP UPDATE message.
- o QoS information Code: this field identifies the type of QoS information:
 - o (0) Reserved
 - o (1) Packet rate
 - o (2) One-way delay metric
 - o (3) Inter-packet delay variation
- o QoS information Sub-code: this field carries the sub-type of the QoS information. The following sub-types have been identified:
 - o (0) None
 - o (1) Reserved rate

- o (2) Available rate
- o (3) Loss rate
- o (4) Minimum one-way delay
- o (5) Maximum one-way delay
- o (6) Average one-way delay
- o QoS information value: this field indicates the value of the QoS information. The corresponding units depend on the instantiation of the QoS information code.
- o QoS class identifier: this field indicates the QC identifier as described in [\[DS\]](#).
- o Address Family Identifier (AFI): this field carries the identity of the Network Layer protocol associated with the Network

Boucadair

Experimental - Expires January 2006

[Page 10]

Internet Draft

QoS-enhanced Border Gateway Protocol

July 2005

Address that follows.

- o Subsequent Address Family Identifier (SAFI): this field provides additional information about the type of the prefix carried in the QOS_NLRI attribute.
- o Length of Next Hop's Network address: this field carries the length of next hop's network address;
- o Network address of Next Hop: this field contains the network address of the next router on the path to the destination prefix.
- o Network Layer Reachability Information: This variable length field lists the NLRI information for the feasible routes that are being advertised by this attribute. The next hop information carried in the QOS_NLRI path attribute defines the Network Layer address of the border router that should be used as the next hop to the destinations listed in the QOS_NLRI attribute in the UPDATE message.

[6.1.5.](#) Multiple paths

[Walton] proposes a mechanism that allows the advertisement of multiple paths for the same prefix without the new path implicitly replacing any previous ones. This is achieved thanks to the use of an arbitrary identifier that will identify (in addition to the prefix) a given path. The QoS_NLRI attributes could support this mechanism by adding: Flags, Identifier, Length, and Prefix in the QoS NLRI attribute. The meaning of these fields is described in [WALTON].

6.2. Processing q-BGP attributes

As described above, q-BGP peers could exchange their respective capabilities through capability negotiation procedure. As a consequence, q-BGP peers will indicate if they both support QoS_NLRI attribute or not. If a q-BGP speaker does not support capability negotiation feature, it will be hard to know in advance its behavior when receiving QoS_NLRI attribute. Therefore, three scenarios should be examined in order to describe the processing of QoS_NLRI attribute by a q-BGP speaker:

- o Scenario 1: If a q-BGP speaker does not support capability feature, QoS_NLRI could be sent to this peer;
- o Scenario 2: If a q-BGP speaker does not support QoS Service Capability, no QoS_NLRI should be sent to this peer;

- o Scenario 3: Both q-BGP peers support QoS Service Capability. In this case, q-BGP peers could use QoS_NLRI attribute. The variant of QoS_NLRI attribute that will be used depends on the nature of the deployed inter-domain QoS delivery solution, either it is a group-1 or group-2.
 - o When sending a QoS_NLRI attribute, the local q-BGP speaker should set the QC identifier field to the identifier of extended QC on the corresponding inter-domain link. In addition, if it is a group-2 solution and if the q-BGP peer supports group-2 QoS delivery solution, the local q-BGP speaker should set the value(s) of "QoS Information value" field(s).
 - o When receiving a QoS_NLRI attribute, q-BGP speaker applies

its inbound policies to grant the received announcement depending on QC binding list. The local q-BGP speaker gets the value of the "QoS Class Identifier" enclosed in the QoS_NLRI of the received announcement and checks if there is a binding entry.

- o If there is no entry in the binding list: the local q-BGP speaker drops the received announcement.
- o If there is an entry in the binding list: the local q-BGP speaker updates the values of "QoS Information value" fields enclosed in the QoS_NLRI with the ones of bound QC.

[6.3.](#) q-BGP Route Selection Process

As far as QoS-related information is conveyed in BGP UPDATE messages, the route selection process should take into account this information in order to make a choice and make a tie-break between equal paths and determine the one(s) to be stored in the local FIB. This process could differ between solutions that belong to group-1 or group-2.

[6.3.1.](#) Group-1 Route Selection Process

For inter-domain QoS-delivery solution that belongs to group-1 only the identifier of the extended QC is to be taken into account in order to choose a path that will be stored in the local RIB. The unique modification to be added to the classical route selection process is to identify routes that serve the same destination with similar extended QCs. Local policies could be configured by each INP in their ASBRs.

From this perspective, the pseudo code of the modified route selection process will be as follows:

Boucadair Experimental - Expires January 2006 [Page 12]

Internet Draft QoS-enhanced Border Gateway Protocol July 2005

=====

1. Identify the received routes that serve the same destination
2. Consider the routes with similar extended QCs

3. Apply local policies (e.g. prefer a given origin AS, cost,à).
4. If only one route has been returned
Store this route in the RIB
5. If more than one route has been returned
Apply the classical BGP route selection process.

=====

6.3.2. Group-2 Route Selection Process

For inter-domain QoS-delivery solutions that belong to group-2, q-BGP UPDATE messages carry QoS performance characteristics together with a QC identifier. q-BGP route selection process should exploit enclosed QoS performance characteristics in order to determine the path that will be stored in the local RIB. Modifications that should be added to the classical route selection process are at least as follows:

- o Identify routes that serve the same destination in the same QC plane;
- o Select a route that optimises QoS performance characteristics.

Therefore, the pseudo code of the new route selection process becomes:

=====

1. Identify routes that serve the same destination
2. Consider routes that have the same QoS class identifier
3. Compare the QoS performance characteristics associated with resulting routes with respect to a given comparison logic
4. Return the route that optimises the QoS performance characteristic
5. if more than one route has been returned, apply the classical BGP route selection process

=====

6.3.2.1. QoS comparison logic

In this section, we discuss several approaches for comparing sets of QoS values enclosed in q-BGP messages. Consider two QoS tuples X and Y. These tuples consist of both the attributes (e.g. delay, jitter, loss rate) and their values. Let the tuples consist of QoS attributes A, B and C, and let the QoS tuple X have the values (Ax, Bx, Cx) and let QoS tuple Y have the values (Ay, By, Cy).

Then to compare the two QoS tuples X and Y, a number of mechanisms can be adopted. To generalise the discussion, here we assume that $P > Q$ means that P is better than Q, irrespective of whether we are comparing bandwidth values (where a higher numerical value represents a better level of QoS) or delay values (where a lower numerical value represents a better level of QoS). The proposed mechanisms are as follows:

- o Lexicographical ordering method: the QoS attributes are compared in strict order. Thus if $A_x > A_y$ then X is better than Y, irrespective of the relative values of Bx, By, Cx or Cy. If $A_x = A_y$ then the second QoS attributes are compared: if $B_x > B_y$ then X is said to be better than Y. We refer to the route selection process that uses this QoS comparison logic as priority-based route selection process.
- o Simultaneous comparison method: X is better than Y if $A_x > A_y$ and $B_x > B_y$ and $C_x > C_y$. Similarly, Y is better than X if $A_y > A_x$ and $B_y > B_x$ and $C_y > C_x$. This approach does not define a result if some of the QoS attributes A, B, C of one tuple are better than the second tuple but some of the QoS attributes are worse. For example, if $A_x > A_y$ while $B_y > B_x$ then the result of the comparison of X with Y is undefined. This approach is not recommended to be used as QoS comparison logic in route selection process implementations.
- o Weighted ordering method: the QoS attributes are normalised to create dimensionless values, and summed. This results in a single value for each QoS tuple, which can be compared to determine which tuple is better. The dimensionless values could additionally be weighted so as to prefer one attribute over others. The advantage of this approach is that it potentially allows a wider range of QoS metrics to be fairly compared, for example a low delay route with reasonable bandwidth.

6.3.2.2. Priority-based route selection process

When a route selection process implements lexicographical comparison

characteristics. Then, the comparison of available routes should be based on the use of the priority value that has been affected to each QoS performance characteristic. The priority ordering of the QoS performance characteristics could be commonly understood by service providers or only configured by each provider. The philosophy of this method is: "find the route that optimises the highest priority QoS-related information".

Therefore, the pseudo code of the priority-based route selection process algorithm is as follows:

=====

1. Identify routes that serve the same destination
2. Consider routes that have the same QoS class identifier
3. Consider the QoS performance characteristic that has the highest priority, and return the routes that optimise that QoS performance characteristic
 - i. If only one route is returned store this route in the local RIB
 - ii. If more than one route are returned
 1. Exclude the QoS performance characteristic that has been used in the step 3 from the list of QoS performance characteristics.
 - a. If there is no remaining QoS performance characteristics, go to step 4
 - b. Else, go to step 3
4. if more than one route has been returned, apply the classical BGP route selection process

=====

Example: Suppose that a q-BGP listener has received the following routes for reaching the same destination. Each of these routes is

associated with a set of QoS performance values as follows:

- o R1: minimum one way delay=150ms, Loss rate=5%
- o R2: minimum one way delay=90ms, Loss rate=2%
- o R3: minimum one way delay=100ms, Loss rate=3%
- o R4: minimum one way delay=200ms, Loss rate=1%

If the q-BGP router is configured to prioritize minimum one way delay, the selected route is R2. But if the q-BGP router is configured to prioritize loss rate, the selected route is R4.

o Processing of route with QoS holes

Let suppose now that a q-BGP router has received from its peers, the following routes for reaching the same destination D1. The received routes enclose the following QoS performance values as detailed below:

- o Route R1: QoS1=90ms, QoS3=150ms, QoS4=5%
- o Route R2: QoS2=30ms, QoS3=153ms, QoS4=1%
- o Route R3: QoS1= 120ms, QoS2=100ms, QoS3= 60ms, QoS4=3%
- o Route R4: QoS2=90ms, QoS3= 50ms, QoS4=8%

The aforementioned routes enclosed different QoS performance characteristics. The issue is how to compare these routes?

This problem could be caused by a mis-configuration or because provider does not support some QoS parameters. This risk in both cases is that providers will not advertise QoS-related information since if only one AS in the chain does not implement a QoS parameter it will introduce a "QoS hole" in all routes that it will advertise and then impact the announcement of this route for downstream ASes.

In order to solve this issue, two solutions could be considered:

- o Solution 1: Add a new control level in the definition of l-QC.

This consists in defining mandatory and optional attributes. A "Mandatory QoS information" is a parameter that must be present in the QoS_NLRI. In the case it is missing the announcement will be dropped by q-BGP receiver. The announcement is not dropped if an "Optional QoS Information" is missing. Nevertheless, the problem of ensuring route selection consistency when optional parameters are missing is unsolved. For solving this case, one of options details under Solution 2 bullet could be implemented. It is clear that all providers should have the same understanding of the mandatory and the optional parameters in order to preserve a consistent and coherent treatment in crossed ASes.

- o Solution 2: No additional control level is introduced in the local QoS Class definition (all parameters are optional). In this second solution, the risk is that service providers won't advertise routes with required QoS information that should be

Boucadair

Experimental - Expires January 2006

[Page 16]

Internet Draft

QoS-enhanced Border Gateway Protocol

July 2005

used as guidance to meet service needs. As a consequence, group-2 solutions become as group-1 ones because there is no control regarding the enclosed QoS information. When a QoS parameter is missing, two options could be considered.

- o Option 1: Discard unvalued routes but keep them all if they are all unvalued. In this case, the priority criterion is respected and the comparison between routes is consistent. But, the risk is that some destinations could be unreachable if received routes do not enclose higher priority QoS performance characteristics.
- o Option 2: Always keep routes with unvalued parameter, but perform selection for the remaining routes. The route selection process compares between routes that have valued the QoS parameter used as criterion in a given step. If there still have equal routes, all routes are considered and the route selection process checks for the route that encloses the next QoS parameter to be used as criterion of its selection even if these routes were not present in the previous step. The inconvenient of this option is that the priority criterion is not satisfied.

As a consequence, the updated pseudo code of the route selection process is as follows:

- =====
1. Identify routes that serve the same destination
 2. Group routes having the same QoS class identifier
 3. For each route group,
 - i. If the number of remaining routes is not null, check for each route, if all mandatory QoS performance characteristics are valued
 - i. If yes, add this route to remaining routes
 - ii. If no, drop this route
 4. Consider the remaining routes
 - i. If the number of remaining routes is not null,
 1. Consider the QoS performance characteristic that has the highest priority, and return the routes that optimise that QoS performance characteristic

Boucadair

Experimental - Expires January 2006

[Page 17]

Internet Draft

QoS-enhanced Border Gateway Protocol

July 2005

- a. If only one route is returned store this route in the local RIB
 - b. If more than one route is returned
 - o Exclude the QoS performance characteristic that has been used in the step 4.i.1 from the list of QoS performance characteristics.
 1. If there is no remaining QoS performance characteristics, go to step 5
 2. Else, go to step 4.i.1
 - ii. If the number of remaining routes is null, go to step 5
6. if more than one route has been returned, apply the classical

=====

7. Security Considerations

This document does not change the underlying security issues in the BGP protocols specifications. The security recommendations related to BGP should be considered [[BGPSEC](#)].

8. References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997

[Walton] Walton, D., et al., "Advertisement of Multiple Paths in BGP", [draft-walton-bgp-add-paths-01.txt](#), Work in Progress, November 2002

[QOSNLRI] Cristallo, G., Jacquenet, C, " Providing Quality of Service Indication by the BGP-4 Protocol: the QOS_NLRI attribute", <[draft-jacquenet-qos-nlri-00.txt](#)>, work in progress

[DS] Nichols, K., Blake, S., Baker, F. and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", [RFC 2474](#), December 1998

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

Boucadair Experimental - Expires January 2006 [Page 18]

Internet Draft QoS-enhanced Border Gateway Protocol July 2005

[BGPSEC] Heffernan, A., "Protection of BGP sessions via the TCP MD5 Signature Option", [RFC 2385](#), August 1998.

[CAP]Chandra, R., Scudder, J., Capabilities Advertisement with BGP-4, [RFC 3392](#) November 2002

[IAB]Atkinson, R., Floyd, S., "IAB Concerns & Recommendations Regarding Internet Research & Evolution", [RFC 3869](#), August 2004

9. Acknowledgments

The authors would also like to thank all the partners of the MESCAL (Management of End-to-End Quality of Service Across the Internet At Large, <http://www.mescal.org>) project for the fruitful discussions.

10. Contributors

- o Pierrick Morand (France Telecom)
- o Pierre Levis (France Telecom)
- o Thibaut Coadic (France Telecom)
- o Hamid Asgari (Thales Research and Technology)
- o Panagiotis Georgatsos (Algonet)
- o David Griffin (University College London)
- o Jason Spencer (University College London)
- o Micheal Howarth (University of Surrey)

11. Author's Address

Mohamed Boucadair
France Telecom R & D
42, rue des Coutures
BP 6243
14066 Caen Cedex 4
France

Email: mohamed.boucadair@francetelecom.com

Intellectual Property Statement

Boucadair	Experimental - Expires January 2006	[Page 19]
-----------	-------------------------------------	-----------

Internet Draft	QoS-enhanced Border Gateway Protocol	July 2005
----------------	--------------------------------------	-----------

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to

pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

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

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

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