Inter-Provider QoS Peering for IP Service Offering Across Multiple Domains

Abolghasem (Hamid) Asgari[1], Mohamed Boucadair[2], Richard Egan[1], Pierrick Morand[2], David Griffin[3], Jonas Griem[3], Panagiotis Georgatsos[4], Jason Spencer[3], George Pavlou[5], Michael P. Howarth[5]

[1]Thales Research and Technology (TRT) Limited, Worton Drive, Reading RG2 0SB, UK, Hamid.Asgari@thalesgroup.com

[2]France Telecom R&D, 42 rue des coutures, BP 6243, 14066 Caen, France

[3]University College London, Dept. of EEE, Torrington Place, London WC1E 7JE, UK

[4]Algonet S.A, 206 Syngrou Ave., Kallithea 176 72, Athens, Greece

[5]University of Surrey, Guildford GU2 7XH, UK

***Abstract* - Providers must interact more closely with each other for providing end-to-end QoS for the global Internet. The focus of this paper is to propose new techniques for providing these interactions in propagating QoS-based agreements among the set of providers involved in the chain of inter-domain service delivery. These providers' interactions occur at the service layer that result in the establishment of QoS-aware service agreements between providers. This needs to be supported by appropriate service management and traffic engineering capabilities per provider domain as well as by BGP-based interactions at the IP layer. The concepts of how to build such agreements are discussed and models for inter-domain QoS peering are proposed. By taking into account the loosely coupled structure of the Internet, we adopt a hop-by-hop, cascaded model in which each provider establishes agreements with its directly attached peers. Once these agreements are in place and networks are provisioned and engineered, QoS-based services can be offered. Targeted value-added services are identified and the methods of supporting the bi-directionality for these services are also proposed.**

***Key-Words* - Inter-domain, Peering, Cascade, QoS, SLA, SLS, Network Provider.**

1 Introduction

In the best-effort Internet, there exist two forms of distinct relationships between providers for traffic exchange, underlined by respective business agreements: peering and transit [1]. Peering is termed as the business relationship whereby providers reciprocally provide access only to each other's customers. Peering is a mutual non-transitive relationship between providers to exchange data between themselves, normally for no fee. Transit is the business relationship whereby one transit provider offers access to all destinations in its routing table to another provider for a charge. It should be clarified that the term 'peering' used throughout this paper is to denote that two providers interact with each other for the purpose of expanding the topological scope of their offered services, under any business relationship which may govern this interaction.

Deployment of QoS-based services across the Internet requires a large set of providers to co-operate and to establish peering agreements for expanding the scope of each other's QoS services. In a QoS-enabled Internet, enhanced agreements are required to determine the QoS levels, the traffic quantities and the destinations to be reached across the pre-existing inter-domain links, together with the agreed financial settlement terms. The intention is to enable a provider to extend its QoS over multiple domains, thus enabling the provider to offer reachability to networks beyond its own domain. Eurescom specified a number of organizational models for these interactions between providers and the support of inter-operator IP-based services [2]. Interactions between providers and in particular service layer interactions are required when offering QoS services across multiple domains. These interactions result in the establishment of service agreements between providers aggregating customer service traffic, which need to be supported by appropriate service management and traffic engineering capabilities per provider domain [3] as well as by BGP-based interactions [4] at the IP layer for QoS inter-domain routing purposes. There are the following essential processes for setting-up an end-to-end inter-domain QoS path: QoS-based capability discovery, mapping and binding of QoS capabilities of the domains, and putting in effect the selected QoS capabilities by using appropriate means such as traffic engineering mechanisms.

This paper investigates the inter-domain QoS peering between providers. We discussed in [5] the issues related to the financial settlement for inter-domain QoS services. This paper is organized as follows. Section 2 defines the terms used throughout this paper. Section 3 describes the

inter-domain QoS peering arrangements emphasizing on the hop-by-hop cascaded model. The issues related to the bi-directional service offering is also discussed in section 4. In section 5, we propose the target services that can be constructed on top of these peering agreements. Section 6 concludes the paper.

2 Definitions and terms

A '***Customer***' (subscriber) denotes an entity, which has the ability to subscribe to QoS-based services offered by 'Providers'. Customers are the target recipients of QoS-based services that are offered on the basis of respective agreements, i.e., Service Level Agreements (SLAs).

An '***IP Network Provider***' (INP[1])' offers IP connectivity services, that is services that provide reachability between hosts in the IP address space. Such 'Providers' own and administer an IP network infrastructure. For the purpose of expanding the geographical span of the offered connectivity services, INPs interact with each other, on a one-to-one peering relationship basis.

'***QoS-based service***' term denotes a service that offers a type of added value to customers, e.g. matching application and customer usage requirements. The current trend in QoS-based service offering is agreement (contract)-based. An SLA denotes such an agreement that describes the characteristics of a service offering and the mutual responsibilities of the customers/providers for using/providing the offered service. The *Service Level Specification (SLS)* is an integral part of a SLA that denotes the technical characteristics of a service offered. Two types of SLS (and subsequently of SLAs) are distinguished here:

- *cSLS (customer SLS)*, established between end-customers and providers, and
- *pSLS (peer SLS)*, established between providers for exchanging traffic in the Internet.

'***Connectivity Service***' is an IP layer transport service for reaching particular destination(s) from specific source(s) in the IP address space. QoS-based connectivity services offered by INPs are divided into elementary services (point-to-point and unidirectional) and complex connectivity services (multipoint-to-multipoint and bi-directional). TEQUILA project specified an SLS template for intra-domain QoS-based elementary connectivity services [6].

2.1 QoS-Classes

A '***QoS-class (QC)***' denotes a basic network-wide *QoS transfer capability* of a provider domain. A QoS transfer capability is a set of attribute-value pairs, where the attributes express various packet transfer performance parameters such as one-way transit delay, packet loss and inter-packet delay variation (jitter), and their particular values. Considering a provider domain, its QoS-classes can be distinguished as *local-QoS-classes* and *extended-QoS-classes* [7] in addition to *Meta-QoS-Class (m-QC)* as an abstract concept. *m-QCs* rely on global understanding of QoS requirements of well-known applications. This distinction is required for capturing the notion of 'QoS capabilities' across domains, upon which QoS-based Internet services could be built. These QoS-classes are explained below.

- A l*ocal-QoS-class (l-QC)* denotes a basic QoS transfer capability that can be provided by means employed in the provider domain itself. Evidently, the topological scope of an *l-QC* is restricted within a domain.
- An *extended-QoS-class (e-QC)* denotes a basic network-wide QoS transfer capability that can be provided by means employed not only in the provider domain but also utilising appropriate means in other provider domains. In other words, an *e-QC* is provided by concatenating the QoS transfer capabilities (QoS-classes) of the provider domain with appropriate capabilities (QoS-classes, *l-QC* or *e-QC*) of other provider domains. The topological scope of an *e-QC* could therefore be outside the boundaries of the provider domain.
- A *Meta-QoS-Class (m-QC)* denotes an abstract QoS-class, where the 'meta' term refers to a quantitative or qualitative range of values of the QoS-class performance parameters. A QoS-class with delay value of 'very low' and loss value of 'very low', or delay-sensitive-QoS-class with delay value of 'low' and loss value of 'any' are typical examples of *m-QCs*. Unlike *e-QC*s that are defined end-to-end with distinct performance characteristics, *m-QCs* do not imply a predefined/engineered end-to-end-QoS [8]. *m-QCs* lays the foundations for a set of "parallel networks", one of which may be suitable for specific service/s.

From a service offering perspective, QoS-classes correspond to the performance (transfer quality) guarantees expressed in c/*pSLS* along with other aspects such as bandwidth, grade of service guarantees and topological scope. From a service provisioning perspective, QoS-classes segregate the network QoS-space into a number of distinct classes, aggregating user QoS traffic accordingly. For a provider domain wishing to provide QoS-classes from its domain to destinations outside its domain, a number of *QC-operations* need to be performed. Further details about QC-operations are provided in [7].

3 Inter-domain QoS Peering Models

Eurescom organizational models [2] are strongly influenced by experience in the telecommunications

[1] The terms INP and AS (Autonomous System) are used interchangeably throughout the paper to denote a business entity owning a network and being responsible for its operation and the provision of Internet connectivity aspects.

industry of provision of international telephony and other services for which network interconnection is a requirement, both in commercial and regulatory terms. We extend these organizational models and build the concepts to establish a set of inter-domain QoS peering models in order to set-up *pSLSs* and consequently to construct end-to-end QoS-based services across the Internet at large scale. The challenge is to propose methodologies to fulfil the requirements of new IP QoS-based services, which will be attractive both to customers and providers.

There are many models for the interconnection and service-layer interactions between providers' for offering QoS services across multiple domains. The type of inter-domain peering impacts the service negotiation procedures, the required signaling protocols, the path discovery through QoS binding, and path selection. The following peering models are considered: hub, centralized, cascaded, and hybrid.

- The *hub* model where the Service Provider (SP), as a distinct entity from INP, is the central point that negotiates and establishes *pSLSs*. Here, the SP takes the responsibility for the overall service management of any given customer IP QoS service instance. This is achieved by making *pSLS* contracts with a chain of INPs so as to create an end-to-end service.
- The *centralised* model, which is similar to the hub model, where an INP negotiates *pSLSs* directly with an appropriate number of downstream providers to construct an end-to-end QoS service. With this model, service peers are not necessarily BGP peers.
- The *cascaded* model where an INP only negotiates *pSLS*s with its immediate neighbouring provider/s to construct an end-to-end QoS service. With this model, service peers can be BGP peers.
- The *hybrid* model is combination of the centralised and the cascaded models. The value of the hybrid model is that it combines the benefits of the cascaded and centralised models. However, it also suffers from the limitations of both models [9].

Any solution for QoS peering should function effectively and in a scalable manner. The two most significant models (centralized and cascaded) are explained in detail below.

3.1 Centralized Model

The centralized model disassociates *pSLS* negotiations from the existing BGP peering arrangements. The originating domain (central point) knows the end-to-end topology of the Internet and directly establishes *pSLS*s with a set of potential domains (neighbor, transit, and distant ASs) in order to reach a set of destinations, to offer an end-to-end QoS-based service. Each AS is responsible for the connection inside its domain and its inter-domain link interfaces. Fig. 1 shows the centralized model where AS1 as the central point establishes *pSLSs* so that its customers can reach destinations in AS3. The scope (the source point and the reachable destinations) of the c/*pSLS*s and the QC binding operation are also shown in this figure. As an example, the intra-domain QoS capability of $l\text{-}QC_3$, $l\text{-}QC_2$, and $l\text{-}QC_1$ are supported and advertised by AS3, AS2, and AS1 respectively. AS1 as the central point discovers these QoS capabilities and constructs the $e\text{-}QC_1$ through QC mappings and bindings of the above stated *l-QC*s. This $e\text{-}QC_1$ is only known by the AS1 and is advertised only to its customers and not to other providers.

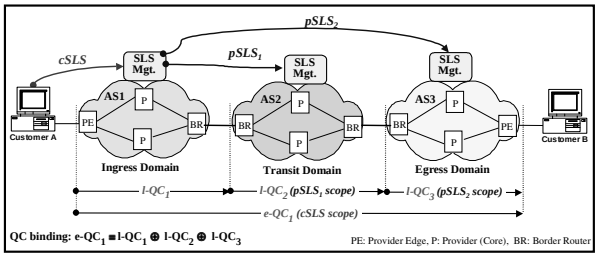

Fig. 1: Centralised QoS peering model.

The central point requires an up-to-date topology of the Internet including the existence and operational status of the physical links between ASs for discovering and selecting end-to-end routes. In addition it needs to know the domains' advertised *l-QCs* in order to perform mapping and binding of these *l-QCs* to form *e-QC*s (e.g., $e\text{-}QC_1$). The combination of the above information is used for selecting the appropriate ASs as well as negotiating and establishing *pSLS* agreements. *pSLS* agreements are therefore tailored to the central point requirements. Although it is possible to find and set-up the optimal routes to the destinations since the central point has access to the overall QoS-based topology, the need for accurate topological and QoS related information of the Internet is a major drawback of this model. This may be feasible for a relatively small number of domains, but it raises scalability concerns when a large number of networks are involved. The central point may end up with many *pSLSs* to manage, which can be in the order of $O(N_d^2)$ where N_d is the number of domains in the Internet.

3.2 Cascaded Model

In the cascaded model, each INP makes *pSLS* contracts with the immediately adjacent interconnected INPs. Thus, the QoS peering agreements are between adjacent neighbors, but not between providers more than "one hop away". This type of peering agreement provides the QoS connectivity from a customer to reachable destinations that may be several domains away. There are two flavors of cascaded model that are described below.

3.2.1 Strict Scope Cascaded Approach

Setting-up *pSLS*s with defined scope and distinct performance characteristics between adjacent INPs is the

compelling feature of this approach. For QoS-Class discovery and selection, each INP in the chain needs to know its adjacent neighbors and the status of related interconnection links. In addition, each INP needs to know the *e-QC*s advertised by its neighboring domains for binding with its own *l-QC*s in order to implement its own *e-QC*s to be advertised to its customers and upstream domains. This is true for every INP involved in the chain in order to implement its *e-QC*s. Fig. 2 gives an overview of the operations in this model. $l\text{-}QC_3$, $l\text{-}QC_2$, and $l\text{-}QC_1$ are supported and advertised by AS3, AS2, and AS1 respectively. AS2 discovers the AS3 capabilities and negotiates a contract (*pSLS2*) with AS3 for enabling its customers to reach destinations in AS3 with an $e\text{-}QC_2$. AS2 constructs (i.e., through QC mapping, binding, and implementation) $e\text{-}QC_2$ and advertises it. This type of process is repeated recursively to enable AS1 customers to also reach destinations in AS3, but at no point do AS1 and AS3 negotiate directly. In each step of the cascade, the upstream provider acts in the consumer role to the provider immediately downstream. It is each provider's responsibility to make appropriate *pSLS*s with the immediate downstream provider making it possible for individual customer IP QoS services to be created and managed along the entire route.

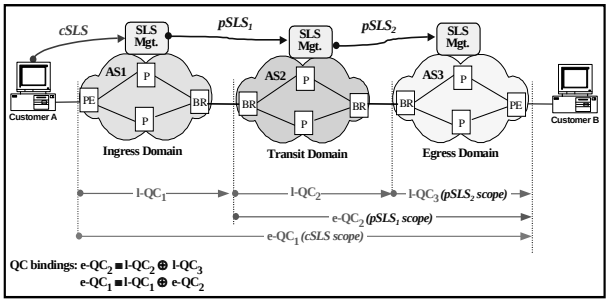

Fig. 2: Strict scope cascaded QoS peering approach.

With this model, a provider is able to aggregate traffic demands by establishing a single *pSLS* with its adjacent provider's domain if that traffic enters the provider's domain from the same ingress point, has the same QoS requirements, and is destined for the same destination point in spite of the fact that the traffic originated from different sources. Therefore, *pSLS* agreements can be tailored at nearly optimum level by aggregating the customers' traffic demands. As traffic demand aggregation can happen at *pSLS* level, each INP may only have a limited number of *pSLSs* to manage. Thus, an order of $O(N_d)$ of *pSLSs* needs to be established by an INP making the cascaded model more scalable. A detailed comparison between centralized and cascaded models is reported in [9].

3.2.2 Open Scope Cascaded Approach

This approach relies on the cascaded model and the use of the *m-QC* concept. Setting-up *pSLS*s with open scope (i.e., no explicit reachability information) and no distinct quantitative performance characteristics but simple compliance with well-known m-QC behaviors between adjacent INPs having qualitative performance characteristics is the compelling feature of this model. In this approach, there is no end-to-end QoS guarantees defined and consequently there is no need to build *e-QC*s, which are the fundamental differences between this approach and 'strict scope cascaded approach'. This approach does not provide any end-to-end bandwidth guarantees because it enables any destination to be reached, without prior explicit indication in the *pSLS*. Each domain is engineered to support a number of local QoS classes (i.e. *l-QC*s). These *l-QC*s are mapped to globally well-known *m-QC*s. Each AS advertises the *m-QC*s that it supports in its administrative domain. Other domains can make *pSLS* arrangement in cascaded fashion with this domain to make use of offered *m-QC*s. Although, inter-domain routing is *pSLS* constrained, each domain can find out whether it can reach certain destinations in an *m-QC* plane through a BGP-like protocol (qBGP) [10]. The core functions of this approach are as follows:

- A globally-known identifier identifies each *m-QC*.
- Establishing *pSLS*s sets off the exchange of inter-domain connectivity information per *m-QC* between service peer domains.
- Each AS announces to its service peers the network prefixes that can be reached within each *m-QC* plane through qBGP.

This results in QoS-enabled Internet that can be viewed as a set of parallel planes each offering service levels associated with a specific *m-QC* and running distinct instances of qBGP. The *pSLS* agreement for an *m-QC* allows the AS to benefit from their neighbor's inter-domain QoS capabilities and enabling it to reach anywhere in the QoS-Internet of that specific *m-QC*.

4 Bi-directionality Support for Complex Connectivity Services

We discussed in previous sections QoS peering models that provide elementary connectivity services. Supporting bi-directionality provides the means in offering complex connectivity services. The primary challenge is in constructing the QoS-enabled reverse path for return traffic. This section discusses providing bi-directionality using the peering models. This section identifies the issues, presents a discussion of the resulting implications and provides methods for resolving them.

4.1 Bi-directionality Support in the Centralized Model

Since the central point has access to the overall QoS-based topology and is able to establish the appropriate SLSs in forward and return directions with a set of potential domains, it can constructs e-QCs in both directions and offer a bi-directional end-to-end QoS-based service.

4.2 Bi-directionality Support in the Strict Scope Cascaded Approach

In the 'strict scope cascaded approach', the scope of the desired *e-QC*s for the forward direction is part of the c/*pSLS* during the negotiation phase. However, when the reverse direction is considered from the destination AS's point of view there are the following apparent issues that must be addressed when constructing bi-directional services.

Firstly, from the destination AS's point of view (AS3 in Fig. 2), the destination for the traffic in the reverse direction is not known. This is due to the fact that the *cSLS* is between Customer A and AS1 and AS1 knows the 3-tuple (Source Customer, Destination Customer, *e-QC : A, B, e-QC1*) whereas AS3 is unaware of it. Thus, AS3 cannot find/verify whether there are *e-QC*s formed through QC binding operations to reach the desired return destination (e.g., Customer A) or not. Therefore, how does AS3 find the scope for the reverse direction (e.g., AS1) in order to see if any *e-QC* exists to reach AS1 customers? Secondly, every time another upstream AS forms an *e-QC* that utilizes the *l-QC* of the destination AS (AS3), the scope of the return paths for AS3 extends and AS3 will not know this. Thirdly, if asymmetric QoS is required (as it is often the case), which QoS class (*l-QC*) at each AS (e.g., AS3) should be used for return traffic and how should this *l-QC* be mapped to an *e-QC* offered by the upstream AS (e.g., AS2)? The destination AS (AS3) has no explicit information to answer the above questions.

Two methods are proposed to tackle the problem of providing QoS enabled path in the reverse direction. The first method extends the single cascade with bi-directional capabilities. The second method employs unidirectional cascades in forward and reverse directions to build bi-directional services.

4.2.1 Single Cascade

One possible solution for setting up a reverse path is to negotiate *pSLSs* in the reverse direction between peer ASs with an open destination scope. This allows the upstream AS (e.g., AS4 in Fig. 3) to offer the QCs to further upstream ASs (e.g., AS3) without the need for amending the scope of pre-existing downstream *pSLS*s every time the scope changes. This implements bi-directional QoS-enabled services by employing *e-QC* enabled *c/pSLSs* in forward direction and *l-QC* enabled *pSLSs* with no explicit *e-QC* binding in reverse direction. While each ISP provides the QoS environment by provisioning its network and allocating resources in the forward direction, it can provide a similar environment for return traffic in the reverse direction within its own domain. This potentially solves the bi-directionality problem at the *pSLS* level, but there are still some issues such as implementing the *e-QC*s and invoking the service, e.g. it may not provide the fine-tuned desired *e-QC* for return traffic. To provide fine-tuned desired QC as discussed in the next section, it requires the use of a signaling mechanism for communicating between the Source and Destination ASs in order to inform the Destination AS the desired QC level for return traffic. This is problematic and violates the cascaded concept, which implies relationships only between cascade neighbors, because there is no direct business relationship between the two remote ASs.

4.2.2 Multiple Uni-directional Cascades

This method allows suitable *e-QC*s to be set-up separately by the source and destination ASs. This method would potentially provide the environment for having bi-directional services using the cascaded approach in both directions. Fig. 3 shows the cascaded implementation for forward direction in which the sources from AS1 (e.g., Customer A), AS2, AS3, AS4 and AS5 can reach Customer C in AS5 with the specified desired *e-QC* quality. Fig. 3 shows the establishment of cascade for reverse direction. However, there are a few implications/issues, arising from this approach.

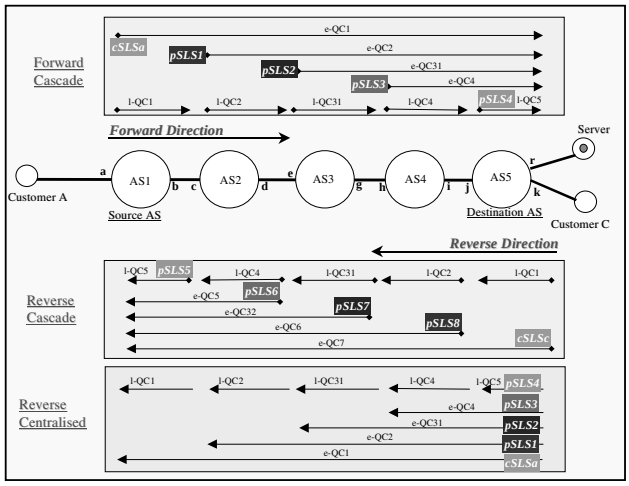

Fig. 3: Bi-directionality support: *e-QC* bindings and SLS set-up in both directions.

The most obvious issue is the requirement to have a suitable reverse path *e-QC*. Since the existence of this path depends on the willingness of the Destination AS to set it up, its existence and availability can be determined either during the invocation phase of the service or via some other pre-set-up means of communication between the involved ASs. This implies that there is a negotiation between Customers/ASs/third parties before invocation has taken place in order to agree the forward and reverse QoS levels. These are then mapped to *e-QC*s at each end according to the inter-domain capabilities (and pre-existing *cSLSs*) at each Source and Destination ASs.

However, additional complexity arises from multiple reverse direction cascades. As shown in Fig. 3, sources from AS1 (e.g., Customer A), AS2, AS3, AS4 and AS5 can reach Customer C in AS5 with the specified desired *e-QC* quality (see Forward Cascade in Fig. 3). For reverse direction, can there be only a reverse path *e-QC* for every

forward path *e-QC* using the cascaded approach in order to allow bi-directional QoS offering? As it is depicted in Fig. 3, it may not possible to achieve this using the cascaded approach starting from Destination AS. This is only possible by constructing *e-QC*s in reverse direction by using centralized approach. As an example (bottom of Fig. 3), *e-QC*$_4$ and *l- QC*$_{31}$ need to be combined to form *e-QC*$_{31}$ in reverse direction and this can only be achieved using the centralized approach. The reverse cascade block in Fig. 3 shows the source AS as the starting point for constructing *e-QC*s in the reverse direction and not the destination AS.

In general, it is necessary to build multiple reverse cascades in the reverse direction to transport return traffic. In the forward path, the *e-QC* paths merge as they approach the Destination AS. This provides the opportunity for constructing *pSLSs* downstream based on the aggregated traffic demands. However, the paths in the reverse direction de-merge as they depart from a Destination AS towards Source ASs. Consequently, for the Destination AS to provide return service to its upstream domains, multiple cascades of *e-QC*s and *pSLSs* set-ups are required. Thus, for a single forward direction cascade, there must be multiple cascades in the reverse direction, depending on the number of Source ASs that are served by a Destination AS.

4.3 Bi-directionality Support in the Open Scope Cascaded Approach

pSLSs for specific *m-QCs* agreed between two domains are not tied to specific destinations and have no further qualifications beyond the domain's boundaries. Hence, as *pSLSs* are uni-directional and they are established for transporting traffic in forward direction, p_rSLS can be established for transporting traffic in reverse direction. The boundary for handling QoS of these two *pSLSs* are the same i.e., *m-QC* support across the domain. The performance targets for these two types of *SLSs* within the domain are also the same, e.g., *m-QC*$_1$. As an alternative option, the reverse direction can be part of the *pSLS* negotiation for the forward path. Although *pSLS* & p_rSLS negotiation can be merged, the reverse direction p_rSLS can have a different bandwidth requirement, which may be specified as part of p_rSLS negotiation.

The path for forward traffic and return traffic between two remote ASs may be different depending on the qBGP updates but the *SLS* agreements between all involved ASs are in place to handle the traffic in both directions irrespective of the paths traffic may take in forward and reverse directions. There might be a different *m-QC* requirement in the reverse direction than the forward direction. To address this, there can be application level communication between the two parties (customers) involved in order to specify the QoS requirements in either direction. This may also require having a trust relationship between the two involved ASs.

5 Target Services

Residential and corporate customers differ both at the level of the performance and traffic guarantees and geographical scope of the services they require. Residential customers may subscribe to IP services to reach any available destination at any time with better-than-best-effort service levels. The duration of the communications may be short and the frequency of interactions can be sparse. Corporate customers, on the other hands, may request specific, strong guarantees with hard upper bounds on QoS parameters and a constant bandwidth for supporting particular mission- or safety-critical applications and services such as IP VPNs in order to reach a limited set of destinations. Obviously, a range of customers could be identified between these two extreme cases, requiring hard upper bounds on delay to a large but limited set of destinations with statistically guaranteed throughput.

The 'strict scope cascaded approach' can be used for services that require QoS performance guarantees for reaching specific destinations and allows end-to-end bandwidth guarantee within statistical bounds. It is able to provide a qualitative QoS service, although quantitative services can also be offered where values for packet delay and loss are specified. Inter-domain QoS services are created by constructing paths across domains that are able to statistically guarantee the required QoS. QoS services can be constructed to meet specific quantified QoS constraints.

The 'open scope cascaded approach' can used to offer better Internet connectivity services with some QoS levels, but doesn't offer any strong guarantees. It enables a provider to offer differentiated transport services, where each differentiated service is related to a *m-QC*. It is envisaged that providers throughout the Internet will implement a small number of well-known *m-QC*s. Inter-domain QoS services are then formed by using advertised paths across those domains that support a particular *m-QC*.

6 Conclusion

In order to provide access to the global Internet, providers must interact with each other; there cannot be a single provider offering global Internet coverage. The aim of this paper was to propose the models of interactions between providers best suited for inter-domain QoS service delivery. The concepts to establish a set of inter-domain QoS peering approaches are discussed in order to construct end-to-end QoS-based services across the Internet at large scale. Two different models are explained. A single point of control for the service instances is the compelling feature of the centralized model. The scalable cascaded model makes it possible to build IP QoS services on a global basis while only maintaining contractual relationships with adjacent providers. Hence, the cascaded model is more scalable than the centralized model. It reflects the loosely coupled

structure of Internet and the current behavior of BGP. A limitation of the cascaded model is that it gives the service initiator less control of the whole IP service path. In the 'strict scope cascaded approach', *pSLSs* are established between adjacent INPs with defined scope and distinct performance characteristics, while the 'open scope cascaded approach', *pSLSs* are set-up between adjacent INPs with open scope and no distinct performance characteristics.

We also discussed the complexity for supporting for bi-directional services. The main issue in the cascaded model is how to construct the QoS-enabled reverse path for return traffic. We identified and discussed some issues with 'strict scope cascaded approach' in order to provide bi-directional services and the resulting implications. These are to do with finding the source/s for return traffic, changing/extending of the scope of the return path, and the selection and mapping of *l-QC/e-QC* for the return direction. We proposed two methods for enabling providers to offer bi-directional services. Providing bi-directionality in using 'open scope cascaded approach' causes less complication. We also discussed the target services that can be supported by using these cascaded approaches. Finally, we are aiming to perform tests both at testbed and simulation environments, evaluating the performance of open scope cascaded model and its relevant aspects at service and control planes. Various kinds of performance assessment tests are being considered, as appropriate for the aspect under test.

Acknowledgements

The work described in this paper has been carried out in the IST MESCAL project, which is partially funded by the Commission of the European Union. The authors would like to acknowledge the contribution of their colleagues in the MESCAL project for their contributions to the ideas presented in this paper.

References

[1] William B. Norton, "Internet Service Providers and Peering", available at http://www.nanog.org/mtg-0302/norton.html.

[2] Project P1008, "Inter-operator interfaces for ensuring end-to-end IP QoS", Deliverable 2, Selected Scenarios and requirements for end-to-end IP QoS management, January 2001.

[3] P. Trimintzios, et al., "Service-driven Traffic Engineering for Intra-domain Quality of Service Management", IEEE Network, special issue on Network Management of Multiservice, Multimedia, IP-based Networks, Vol. 17, No. 3, pp. 29-36, May/June 2003.

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

[5] P. Georgatsos, J. Spencer, D. Griffin, P. Damilatis, H. Asgari, J. Griem, G. Pavlou and P. Morand, "Provider-level Service Agreements for Inter-domain QoS delivery", to appear in proceeding of Fourth International Workshop on Advanced Internet Charging and QoS Technologies (ICQT04), Springer, September 2004.

[6] D. Goderis, et al., "Service level specifications semantics, parameters, and negotiation requirements" IETF Internet-Draft, draft-tequila-sls-02.txt, Feb. 2002.

[7] P. Flegkas, et al., Mescal Deliverable D1.1, "Specification of Business Models and a Functional Architecture for Inter-domain QoS Delivery", available at: http://www.mescal.org/, June 2003.

[8] P. Levis, M. Boucadair, P. Morand, P. Trimintzios, "The Meta-QoS-Class concept: a step towards global QoS inter-domain services", to appear in proceeding of IEEE International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2004), Oct. 2004.

[9] H. Asgari, et al., Mescal Deliverable D1.4, "Issues in MESCAL Inter-Domain QoS Delivery: Technologies, Inter-operability, and Financial Settlement", available at: http://www.mescal.org/, Jan. 2004.

[10] M. Howarth, et al., Mescal Deliverable D1.2, "Initial specification of protocols and algorithms for inter-domain SLS management and traffic engineering for QoS-based IP service delivery and their test requirements," available at: http://www.mescal.org/, Jan. 2004.