Provider-level Service Agreements for Inter-domain QoS delivery

Panos Georgatsos[1], Jason Spencer[2], David Griffin[2], Takis Damilatis[1], Hamid Asgari[3], Jonas Griem[2], George Pavlou[4], Pierrick Morand[5]

[1]Algonet SA, Athens, Greece, {pgeorgat, pdamil}@egreta.com
[2]University College London, UK, {dgriffin, jsp, jgriem}@ee.ucl.ac.uk
[3]Thales Research and Technology Ltd., Reading, UK, Hamid.Asgari@thalesgroup.com
[4]University of Surrey, Guildford, UK, g.pavlou@eim.surrey.ac.uk
[5]France Telecom R&D, Caen, France, pierrick.morand@rd.francetelecom.com

Abstract. In the current Internet, business relationships and agreements between peered ISPs do not usually make specific guarantees on reachability, availability or network performance. However, in the next generation Internet, where a range of Quality of Service (QoS) guarantees are envisaged, new techniques are required to propagate QoS-based agreements among the set of providers involved in the chain of inter-domain service delivery. In this paper we examine how current agreements between ISPs should be enhanced to propagate QoS information between domains, and, in the absence of any form of central control, how these agreements may be used together to guarantee end-to-end QoS levels across all involved domains of control/ownership. Armed with this capability, individual ISPs may build concrete relationships with their peers where responsibilities may be formally agreed in terms of topological scope, timescale, service levels and capacities. We introduce a new concept of QoS-proxy peering agreements and propose a cascade of inter-domain Service Level Specifications (SLSs) between directly attached peers: each ISP meeting the terms of the SLSs agreed with upstream peers by being responsible for its own intra-domain service levels while relying on downstream peers to fulfill their SLSs.

1 Introduction

There is a growing trend towards IP-based services, not only from a technical perspective e.g. VoIP (Voice over IP) services, but also from business perspectives, e.g. the emergence of Internet-based Application Service Providers. As more and more network-performance-sensitive services migrate to IP networks, best-effort networks no longer meet their QoS requirements. To this end, services require differentiation to provide different QoS levels for different applications over the Internet at large.

The issue of provisioning end-to-end QoS in the Internet is currently being investigated by both research and standardisation communities. Requirements from ISP perspectives and proposals targeted at building MPLS-based inter-domain tunnels have been recently submitted to the IETF [9], [10], [11] by key players in the field. As a pure layer 3 (routed) solution, the European-funded IST research project MESCAL [7] is investigating solutions targeted at building and maintaining a QoS-aware IP layer spanning across multiple domains, and considers architecture, implementation and the business models associated with it [4].

To provide end-to-end QoS across the Internet closer co-operation between multiple ISPs is required. The business relationships between ISPs must be considered and this is the main focus of this paper, although, we do not consider accounting and data collection methods, charging, rating and pricing models in this paper. A number of assumptions are made when approaching the co-operation problem. Firstly, we assume that there is no global controlling entity over all ISPs; secondly, that the problem of intra-domain is already solved and that different domains offer a range of *local Quality Classes* (l-QCs) between their edge routers.

With these assumptions in place we describe and analyse various modes of distributed interaction between ISPs and propose suitable business models to provide both loose, qualitative and statistical, quantitative end-to-end QoS guarantees. However, to achieve this ISP interaction, it is seen that a strengthening of business relationships is required with explicit QoS information being part of these agreements. To achieve this we use *pSLSs* (provider-Service-Level-Specifications) to describe QoS attributes to given destinations. These pSLSs are then used to concatenate l-QCs to form *extended Quality Classes* (e-QCs) to remote destinations. We analyse various ways this concatenation can be achieved and the implications on scalability.

Once the pSLSs have been created, the further problem of financial settlement must be considered. Drawing on concepts from the current Internet and the global PSTN (Public Switched Telephone Network) we propose new peering agreements and examine the flow of monies across the new QoS-enabled Internet.

The paper is organised as follows. Section 2 describes the current arrangements between ISPs and describes the business relationships between international PSTN operators. Section 3 analyses source-based and cascaded models for QoS-peering. Section 4 considers business cases for both loose and statistically guaranteed performance levels and the associated financial settlement issues between ISPs. Finally, section 5 presents our conclusions.

2 Business Relationships and Financial Settlements in the Current Internet and PSTN Networks

The global Internet is a collection of independently operated networks whose organisation in retrospect has been modelled by a three-tiered hierarchy [6] (Figure 1). The connectivity and position in the tier model is dependent on the size of the ISP, its geographic reach, capacity (in terms of link speeds and routing capability) and the available reachable prefixes. While this model is not strictly accurate it serves to demonstrate the variety of ISPs and their relationships within the Internet.

Currently, in the best-effort Internet, there exist two forms of distinct relationships between ISPs for traffic exchange, underlined by respective business agreements:

peer-to-peer and *transit* (customer-provider). A transit relationship is where one provider provides forwarding (transit) to destinations in its routing table (could be to the global Internet) to another ISP for a charge. Usually, this type of business relationship is between ISPs belonging to different tiers of the three-tier Internet model (lower tier ISP being a customer of the upper tier ISP). Peer-to-peer is the business relationship whereby ISPs reciprocally provide only access to each other's customers and it is a non-transitive relationship. It is a kind of 'short-cut' to prevent traffic flowing into the upper tiers and allows for the direct flow of traffic between the peer-to-peer ISPs.

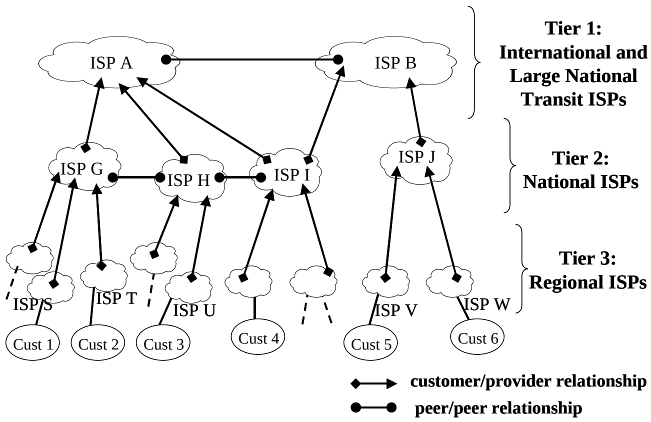

Fig. 1. A post-hoc approximation of a three-tier Internet model with peering/transit agreements.

Financial Settlements

The financial settlements between ISPs primarily depend on their business relationship. In the *service-provider settlement*, a customer (end-customer or ISP) pays a flat rate or a usage-based amount to the provider ISP for reachability to networks, which the provider ISP can reach through its peers, customers or through its own provider ISPs. The customer will always pay whether the traffic is being sent or received.

In the *negotiated-financial settlement*, the traffic volume in each direction is monitored and then payment is made on the net flow of traffic.

In the *settlement-free agreement* (also known as the Sender-Keeps-All, or SKA agreement in the PSTN) neither ISP pays the other for traffic exchange, and they usually split the physical layer costs between them. This settlement is a special case of the negotiated-financial settlement, because either the traffic is symmetric or because the perceived gain to each party is considered worth the agreement.

PSTN Networks

An existing network that requires close business relationships to provide end-to-end better-than-best-effort communication is the PSTN (Public Switched Telephone Network). The inter-connection of international PSTN networks has a number of traits similar to the inter-domain QoS problem: such as resources must be reserved to provide the required level of service but direct interconnection between all networks is not possible and therefore trust relationships are required. When crossing international boundaries the problem of trust becomes more acute, and the transit network topology is a function of political and financial considerations. These peering agreements are overseen by the International Telecommunications Union (ITU) and the World Trade Organisation (WTO) and usually result in a financial settlement like the Accounting Revenue Division Procedure [1]. Here the originating caller at the originating network is charged at a customer rate, and the originating network pays the terminating network to terminate the call at a previously negotiated rate, the settlement rate. The transfer of monies is then usually performed only on the net flow in a negotiated-financial settlement style agreement. If, however, the perceived value to each party is similar, as it may be when the edge operators adjust their customer rates, the inter-operator agreements approach the SKA (sender-keeps-all) agreement.

These ideas could be drawn into a QoS-enabled Internet business model, albeit with the following differences: PSTN has a single QoS traffic unit – the call-minute, whereas a QoS-enabled Internet could have many QoS levels. Also, the PSTN, especially at the international level, is more flat than hierarchical and has organisations in loose control to mediate and arbitrate settlements. Another issue that must be considered is which traffic end-point is the initiator, who gains from the data flow, and who is charged for it [3], as well as methods of charging for multicast traffic. Some of these issues have been considered in the split-edge pricing work in [2].

3 QoS Peering Approaches

In the current Internet, ISPs form business relationships with one another and deploy links between their networks, either directly or through Internet exchange points (IXPs). BGP policies are then deployed to determine which prefixes will be advertised to adjacent providers and subsequently best-effort traffic may be routed across the Internet.

We view that in a QoS-enabled Internet additional agreements are needed to determine the QoS levels, traffic quantities and the destinations to be reached across the pre-existing inter-domain links, together with the agreed financial settlement terms, penalty clauses, etc. The technical aspects of the QoS agreements between ISPs are contained in pSLSs, introduced in section 1. Once the pSLSs are in place, BGP will announce the existence of destinations tagged with the agreed QoS levels, and traffic conforming to the pSLSs may be forwarded to remote destinations and receive the appropriate treatment to meet the agreed performance targets. pSLSs are meant to support aggregate traffic, and they are assumed to be in place prior to any agreements with end customers (via customer-SLSs, or *cSLSs*) or upstream ISPs (via pSLSs) to use services based on them. They are negotiated according to the business policies of the provider and the outcome of the deployed off-line inter-domain traffic engineering algorithms that use forecasts of customer traffic as input. pSLSs are considered as semi-permanent, only changing when traffic forecasts for aggregate traffic alter significantly, or when business policies are modified. They should not be seen as dynamic entities to be deployed for individual customer flows.

There are many models for the interconnection and service-layer interactions between ISPs. Such models are required not only to establish a complete end-to-end customer service, but also to provide and maintain the hierarchy of related management information flows. Eurescom specified organisational models for the support of inter-operator IP-based services [8]. The models may be grouped into three configurations known as the *cascade model*, the *hub model*, and the *mixture model*: a combination of the two. The type of inter-domain peering impacts the service negotiation procedures, the required signalling protocols, the QoS binding, and path selection. In the following we give an overview of the cascade and source-based model and analyse their pros and cons.

Source-Based Approach

The source-based approach (similar to the hub model from Eurescom) disassociates pSLS negotiations from the existing BGP peering arrangements. The originating domain knows the end-to-end topology of the Internet and establishes pSLSs with a set of adjacent and distant domains in order to reach a set of destinations, with a particular QoS.

As shown in Figure 2, the originating domain (AS1) has the responsibility for managing the overall requested QoS service/connection. To manage customer requests, the provider (AS1) directly requests peering agreement ($pSLS_1$ and $pSLS_2$) with providers AS2 and AS3 and with any other network provider involved in order to create an e-QC (from AS1 to AS3).

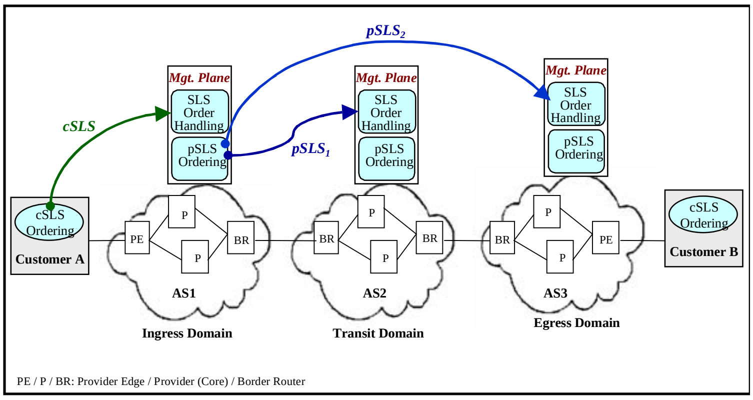

Fig. 2. Source-Based Approach

Cascaded Approach

In the cascaded approach, each ISP established pSLSs only with adjacent ISPs, i.e. those ISPs with whom there are existing BGP peering relationships. Figure 3 gives an overview of the operations in this approach. The domain AS3 supports an intra-domain QoS capability (l-QC_1). AS2 supports an intra-domain QoS capability (l-QC_2)

and is a BGP peer of AS3. AS2 and AS3 negotiate a contract ($pSLS_2$) that enables customers of AS2 to reach destinations in AS3 with a QoS (e-QC_1). This process can be repeated recursively to enable AS1 to also reach destinations in AS2 and AS3, but at no point do AS1 and AS3 negotiate directly.

There are no explicit end-to-end agreements in the cascaded approach, each domain may build upon the capabilities of adjacent downstream ASs to form its own e-QCs to the required destinations. This recursive approach results in an approximation of end-to-end agreements with the merit, as discussed in the following section, of being more scalable.

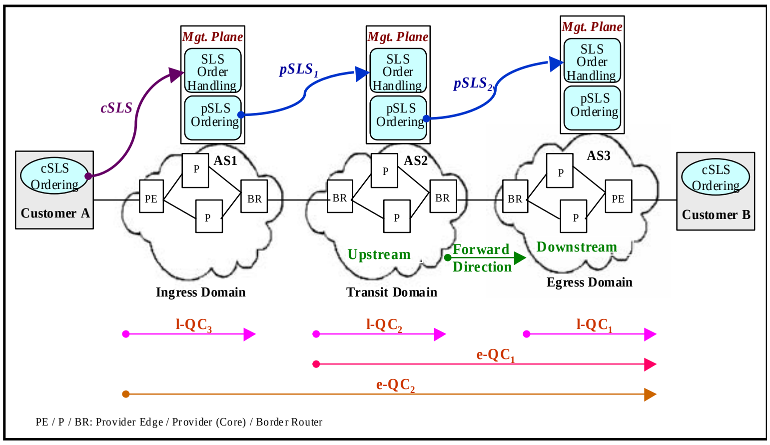

Fig. 3. Cascaded Approach

Strengths and Limitations of the Source-Based and Cascaded Approaches

The originating domain in the source-based approach requires an up-to-date topology of the Internet including the existence and operational status of every physical link between ASs. Whereas, in the cascaded approach, each ISP in the chain only needs to know its adjacent neighbours and the status of related interconnection links.

In both approaches inter-domain routing is pSLS constrained, i.e. traffic will only pass through ASs where pSLS agreements are already in place. Since the originating domain in the source-based approach has topological knowledge of all domains and their interconnections, it is possible to exercise a finer degree of control over the chain of pSLSs through to the destinations. In the cascaded approach each AS participating in the chain does not have all topology data and the initiator is obliged to use the e-QCs previously constructed by the downstream domain. Therefore there is less flexibility and control of the whole IP service path, which may result in sub-optimal paths, although the QoS constraints of the traffic will still be met.

In conclusion, a single point of control for the service instances is the compelling feature of the source-based approach. However it would be difficult to manage for more than a few interconnected ISPs and it is expected that most providers would prefer the cascaded approach, which reflects the loosely coupled structure of the Internet. The cascaded approach makes it possible to build IP QoS services on a global basis while only maintaining contractual relationships with adjacent operators. Hence, this approach is more scalable than the source-based approach. This also reflects the current behaviour of BGP.

4 Inter-Domain Business Relationships

Considering a hop-by-hop, cascaded approach for interactions between providers, the following business cases are proposed. Figure 4 depicts the business case, which directly corresponds to the business model of the Internet as it stands today. The business relationships (transit and peer-to-peer) need to be supported by appropriate pSLSs to allow the exchange of QoS traffic. Due to the bi-directional nature of these business relationships and their broad topological scope, only services with loose, qualitative QoS guarantees may be supported in this case.

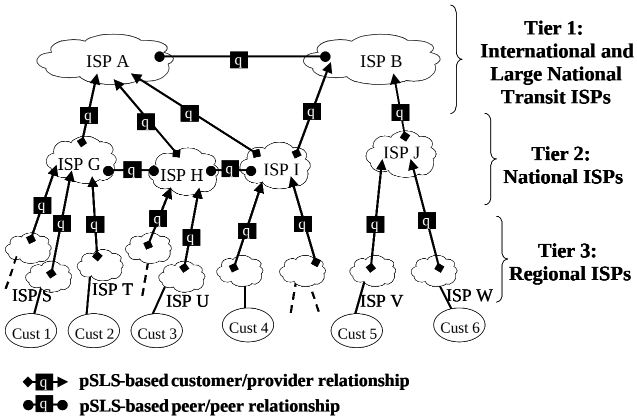

Fig. 4. Case A: Provisioning services with loose QoS guarantees

To provision services with statistical guarantees on quantitative bandwidth and performance metrics a so-called *upstream-QoS-proxy* (or simply *QoS-proxy*) relationship needs to exist between ISPs. Figure 5 illustrates this second business case.

In the QoS-proxy relationship, either of the ISPs may agree with the other ISP to provide a transit QoS-based connectivity service to (a subset of) the destinations it can reach with this QoS level. The ISP offering the transit QoS service would have built its QoS reach capabilities based on similar agreements with its directly attached ISPs and so on.

Two things are worth noting about this business relationship: first, its liberal, unidirectional nature, where either ISP can use the other as a QoS proxy; and, second, its strong collaborative and transitive nature, which is built in a cascaded fashion. The QoS-proxy business relationship differs from the peer-to-peer and customer-provider relationships in the current best-effort Internet in the connotation of the established agreements on traffic exchange and subsequently in the directionality of the traffic flows. In customer-provider relationships, agreements are established for transporting traffic from/to the customer ISPs to/from the provider ISPs and in the peer-to-peer relationships, agreements are established for the ISPs to exchange traffic on a mutual basis; in QoS-proxy relationships, agreements may be established independently in either direction, as each ISP wishes.

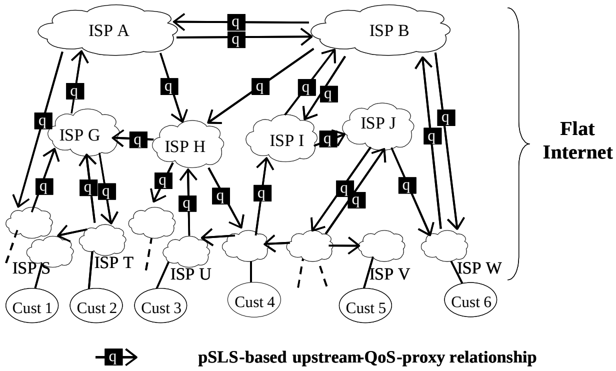

Fig. 5. Case B: Provisioning services with statistical QoS guarantees

The differences between the two business models, as discussed above, are attributed to the diverse types of services each model is set-up to provide. In a best-effort or loose-QoS-based connectivity service Internet, geographical coverage is clearly the strongest selling point; hence, the customer-provider and peer-to-peer relationships. In a statistical-QoS-based service Internet, provided that there is a global demand for related services, the ability to provide such QoS levels becomes clearly an asset. As such, what is required, is to seek for suitable peers to deliver such QoS to desired destinations; hence, the QoS-proxy relationship, which equally applies to all ISPs - regardless of their size.

The business model based on the QoS-proxy relationship resembles the current practice in the traditional telecom service world. The QoS-proxy business agreements could be seen as corresponding to call-termination agreements established between telephony operators/providers. Furthermore, in the telco's world, synergies between operators are built mainly on grounds of reliability and competitiveness, much as the flat Internet business model implies. This resemblance is not surprising; telephony calls and IP services based on statistically guaranteed quantitative QoS metrics, provided that there is a global demand for them, are very similar in that they are both commodities, which need to be widely offered at a certain quality.

Financial Settlements

First, it should be made clear that the financial settlements in a QoS-aware Internet are *in addition* to the settlements made for best-effort connectivity. Broadly speaking, the following two principles govern these settlements: the ISP who requests the pSLS pays the other; and in the cases where pSLSs exist in both directions or the pSLSs have the connotation of mutual agreements, payment reconciliation may take place.

Table 1. Financial Settlements in the QoS-aware Internet

Type of business relationship	Type of financial settlement
customer-provider	service-provider settlement
peer-to-peer	negotiated-financial, or, settlement-free agreement
QoS-proxy	service-provider settlement, if only one ISP requests pSLSs, or, negotiated-financial or settlement-free agreement, if ISPs request pSLSs from each other

5 Conclusions

We propose in this paper that agreements for QoS traffic exchange need to be in place between ISPs in addition to agreements for best-effort connectivity. As with all agreements between ISPs, the establishment of these pSLSs is instigated by the business objectives of the ISP. By augmenting the peer-to-peer and transit (customer-provider) relationships of today's Internet with QoS-based pSLSs we may achieve loose inter-domain QoS guarantees. For provisioning harder QoS guarantees to specific destinations a new QoS-proxy business relationship is proposed. This type of business relationship could be thought as being the QoS Internet counterpart of call-termination agreements in the PSTN or VoIP business world.

The following points are worth mentioning regarding the proposed business relationships. First, they are built on top of existing best-effort agreements and they are not exclusive: e.g. ISP A may act as transit for ISP B, while using ISP C as a QoS-proxy, if this best fits its business objectives. The proposed business models therefore offer an incremental, best-effort compatible migration path towards a QoS-capable Internet.

Second, a key aspect of the proposed business relationships is that they are based on pSLSs, i.e. service agreements, implying legal obligations. On one hand, this offers a tangible lever for ensuring trust between ISPs, which after-all is the bottom-line of any inter-domain solution. On the other hand, this might discourage ISPs. However, given that QoS-based services are offered to end-customers on the basis of Service Level Agreements, requiring the same for ISPs seems reasonable and fair.

Third, the feasibility of the realisation of the proposed business relationships is currently being undertaken by the IST MESCAL project [7]. The MESCAL solutions encompass service management and traffic engineering functions and rely on interactions between adjacent ISPs both at the service layer - for pSLS establishment - and the IP layer - for determining, configuring and maintaining suitable inter-domain QoS routes. The MESCAL validation work covers:

- information models for describing pSLSs under each of the identified business relationships,
- the logic and protocols for pSLS negotiation,
- mechanisms for parsing and extracting the necessary traffic engineering information from the pSLSs,
- extensions to the BGP protocol and associated route selection algorithms to allow the exchange and processing of QoS routing information based on the established pSLS,

- off-line traffic engineering algorithms for determining the required set of pSLSs based on anticipated QoS traffic demand and dimensioning network resources accordingly.

Inevitably, the price to be paid is mainly the increase of the size of the Internet routing tables – growing linearly with the number of the distinct levels of QoS offered. Further simulations and testbed experimentation work [5] for assessing inter-domain routing performance aspects is currently underway.

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. Andreas Thuswaldner, "The international revenue settlement debate", Telecommunication Policy Volume 24 Issue 1 (2000) pp 31-50.
2. Bob Briscoe, “The Direction of Value Flow in Connectionless Networks”, Proceedings of Networked Group Communication: First International COST264 Workshop, NGC’99, Pisa, Italy, November 17-20, (1999).
3. David D. Clark, “Combining Sender and Receiver Payments in the Internet”, Tele-communications Research Policy Conference (1996).
4. Paris Flegkas, et al., "D1.1: Specification of Business Models and a Functional Architecture for Inter-domain QoS Delivery", http://www.mescal.org/, (2003).
5. Michael Howarth, et al., "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", http://www.mescal.org/, (2004).
6. G. Huston, "ISP Survival Guide", John Wiley and Sons 1998. ISBN 201-3-45567-9
7. http://www.mescal.org
8. Eurescom 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, (2001).
9. JP Vasseur, Arthi Ayyangar, "Inter-area and Inter-AS MPLS Traffic Engineering", IETF Internet Draft, (2004).
10. JL Le Roux, JP Vasseur, Jim Boyle, "Requirements for Inter-area MPLS Traffic Engineering", Internet Draft, TEWG Working Group, (2004).
11. Raymond Zhang, JP Vasseur, "MPLS Inter-AS Traffic Engineering requirements", Internet Draft, TEWG Working Group, (2004).