



INTERNATIONAL TELECOMMUNICATION UNION

# ITU-T

TELECOMMUNICATION  
STANDARDIZATION SECTOR  
OF ITU

## Series E

### Supplement 8

(11/2009)

SERIES E: OVERALL NETWORK OPERATION,  
TELEPHONE SERVICE, SERVICE OPERATION AND  
HUMAN FACTORS

---

## Guidelines for Inter-Provider Quality of Service

### **CAUTION !**

### ***PREPUBLISHED RECOMMENDATION***

This prepublication is an unedited version of a recently approved Recommendation. It will be replaced by the published version after editing. Therefore, there will be differences between this prepublication and the published version.

## FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

## NOTE

In this publication, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this publication is voluntary. However, the publication may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the publication is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the publication is required of any party.

## INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this publication may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the publication development process.

As of the date of approval of this publication, ITU [had/had not] received notice of intellectual property, protected by patents, which may be required to implement this publication. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2009

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

## **E.800SerSup8 – Supplement 8 to the E.800-series of Recommendations**

### **Guidelines for Inter-Provider Quality of Service**

#### **Preamble**

This Supplement was originally prepared as a white paper by the Interconnection Working Group of the Communications Futures Program (CFP) of the Massachusetts Institute of Technology (MIT). Following discussions at a meeting of the Quality of Service Development Group (QSDG, formerly associated with ITU-T Study Group 2, but now Study Group 12) this document was brought into Study Group 12, where some interest was shown.

Because some conflicts were identified with the content of existing SG 12 Recommendations in the Y.154X-series, it was agreed to publish it as a Supplement, with the understanding that in any areas of conflict, the relevant ITU-T Recommendations shall take precedence.

#### **Communications Futures Program (CFP)**

CFP is a partnership between university and industry at the forefront of defining the roadmap for communications and its impact on adjacent industries. CFP's mission is to help our industry partners recognize the opportunities and threats from these changes by understanding the drivers and pace of change, building technologies that create discontinuous innovation and building the enablers for such innovation to be meaningful to our partners. Further information about CFP, see <http://cfp.mit.edu>.

## Table of Contents

|       |                                                         |    |
|-------|---------------------------------------------------------|----|
| 1     | Introduction.....                                       | 6  |
| 1.1   | Scope.....                                              | 6  |
| 1.2   | Relationship to standards .....                         | 6  |
| 2     | Reference model & terminology.....                      | 7  |
| 2.1   | Definitions.....                                        | 7  |
| 2.2   | Reference approach.....                                 | 9  |
| 2.3   | Basic reference case .....                              | 10 |
| 2.3.1 | Customer Managed CE model .....                         | 10 |
| 2.3.2 | Provider Managed CE model.....                          | 11 |
| 2.4   | Inter-provider reference case .....                     | 11 |
| 2.4.1 | Provider Managed CE model with multiple providers ..... | 12 |
| 2.5   | DSCP Marking.....                                       | 12 |
| 2.6   | Routing.....                                            | 13 |
| 2.7   | Measurement and management.....                         | 13 |
| 3     | Service Class.....                                      | 13 |
| 3.1   | Network Service Classes.....                            | 13 |
| 3.1.1 | Class 5.....                                            | 14 |
| 3.1.2 | Class 0.....                                            | 14 |
| 3.1.3 | Other Service Classes.....                              | 14 |
| 3.2   | Customer to Provider Interface (CPI) Behavior .....     | 14 |
| 3.2.1 | Marking of Customer Traffic .....                       | 14 |
| 3.2.2 | Policing and Re-Marking .....                           | 15 |
| 3.3   | Network to Network (NNI) Behavior .....                 | 15 |
| 3.3.1 | Marking of Traffic at NNI .....                         | 15 |
| 3.3.2 | Policing and Re-Marking at the NNI.....                 | 16 |
| 3.4   | Definitions of metrics.....                             | 16 |
| 3.4.1 | Initial Considerations .....                            | 16 |
| 3.4.2 | Measurement and reporting considerations .....          | 17 |
| 3.4.3 | One-way Packet Transfer Delay .....                     | 17 |
| 3.4.4 | One-way IP Packet Delay Variation (IPDV) .....          | 18 |
| 3.4.5 | IP Packet Loss Ratio (IPLR).....                        | 18 |
| 3.5   | Class 0 Service definition .....                        | 19 |
| 3.6   | Impairment Budgets.....                                 | 19 |
| 3.6.1 | Introduction.....                                       | 19 |

|         |                                                                               |    |
|---------|-------------------------------------------------------------------------------|----|
| 3.6.2   | Division of end to end Class 0 budget .....                                   | 19 |
| 3.6.3   | The challenge of budget division and subsequent concatenation .....           | 20 |
| 3.7     | Proposed allocation approach .....                                            | 21 |
| 3.8     | Assessment reference model .....                                              | 21 |
| 3.8.1   | Access Segment .....                                                          | 22 |
| 3.8.2   | Number of providers .....                                                     | 23 |
| 3.9     | Budget allocation .....                                                       | 23 |
| 3.9.1   | Budget Allocation for Providers offering services between a UNI and NNI ..... | 24 |
| 3.9.2   | Budget allocation for IPDV and IPDV concatenation .....                       | 24 |
| 3.9.2.1 | The “2 point promise” .....                                                   | 24 |
| 3.9.2.2 | IPDV performance bands .....                                                  | 24 |
| 3.9.2.3 | IPDV core budget thresholds and probabilities .....                           | 25 |
| 3.9.2.4 | IPDV access budget thresholds and probabilities .....                         | 25 |
| 3.9.2.5 | Estimation of IP Path IPDV .....                                              | 26 |
| 3.9.3   | Impairment allocation budget for IPTD and IPLR .....                          | 27 |
| 3.10    | Admission management .....                                                    | 28 |
| 4       | QoS Measurement .....                                                         | 28 |
| 4.1     | QoS Measurement Requirements .....                                            | 29 |
| 4.1.1   | Service Provider Measurement Agreements .....                                 | 29 |
| 4.2     | QoS Measurement Methodologies .....                                           | 30 |
| 4.3     | QoS Measurement Protocols .....                                               | 31 |
| 4.3.1   | OAM based active measurement .....                                            | 31 |
| 4.4     | Measurement considerations for VPN services .....                             | 31 |
| 4.4.1   | Measurement of each service provider’s network performance .....              | 31 |
| 4.4.2   | Measurement of individual customer VPNs .....                                 | 32 |
| 4.5     | Recommended measurement approach .....                                        | 32 |
| 4.6     | Performance measurement metrics .....                                         | 32 |
| 4.6.1   | IP Delay and Delay Variation measurement metrics .....                        | 32 |
| 4.6.2   | IP Loss Ratio measurement metrics .....                                       | 33 |
| 4.7     | Reporting of Measurement results .....                                        | 33 |
| 4.7.1   | Proposal for reporting of measurement results .....                           | 34 |
| 4.7.2   | Monitoring and comparison to threshold .....                                  | 34 |
| 4.7.3   | Reporting the measurement results .....                                       | 35 |
| 4.8     | QoS Measurement Security Considerations .....                                 | 35 |
| 5       | Routing .....                                                                 | 36 |

|         |                                                                    |    |
|---------|--------------------------------------------------------------------|----|
| 5.1     | Current BGP Capabilities .....                                     | 36 |
| 5.2     | BGP Considerations.....                                            | 36 |
| 5.3     | Routing recommendation.....                                        | 37 |
| 6       | Securing QoS .....                                                 | 37 |
| 6.1     | Motivation.....                                                    | 37 |
| 6.2     | Areas which need to be secure.....                                 | 37 |
| 6.3     | Provisioning Security .....                                        | 38 |
| 6.3.1   | Goals .....                                                        | 38 |
| 6.3.2   | Attacks .....                                                      | 38 |
| 6.3.2.1 | Eavesdropping.....                                                 | 39 |
| 6.3.2.2 | Replay .....                                                       | 39 |
| 6.3.2.3 | Message Insertion .....                                            | 39 |
| 6.3.2.4 | Deletion.....                                                      | 39 |
| 6.3.2.5 | Modification.....                                                  | 39 |
| 6.3.2.6 | Denial of Service.....                                             | 39 |
| 6.3.3   | Security of Provider-Provisioned CE Devices.....                   | 40 |
| 6.3.4   | Carrier of Carriers Issues .....                                   | 40 |
| 6.4     | Service Security .....                                             | 40 |
| 6.5     | Security Guidelines .....                                          | 41 |
| 7       | Operational Issues.....                                            | 42 |
| 7.1     | Fault .....                                                        | 42 |
| 7.2     | Configuration and Maintenance.....                                 | 43 |
| 7.3     | Accounting.....                                                    | 44 |
| 7.4     | Performance .....                                                  | 45 |
| 8       | References.....                                                    | 45 |
|         | Appendix A. Examples of the application of budget allocations..... | 47 |
|         | A.1 Case 1: Three Core Providers .....                             | 47 |
|         | A.2 Case 2: Transcontinental Service, 5 Core Providers.....        | 48 |

## Executive Summary

This Supplement presents a proposal to enable the deployment of Inter-provider Quality of Service (QoS). The Supplement begins from the observation that QoS based on the Differentiated Services architecture defined in RFC 2475 is now widely deployed within the networks of single providers. This is especially the case for providers of network-based VPNs (see, for example RFC 2547, RFC 4364). Some providers are now beginning to interconnect with each other via "QoS-enabled peering" in an attempt to offer QoS that spans the networks of multiple providers. However, in the absence of appropriate standards and established procedures for management, trouble-shooting, monitoring, etc., such interconnections are likely to be challenging and labour-intensive. This Supplement seeks to identify the key issues that service providers need to agree upon if inter-provider QoS is to be readily deployable.

### This Supplement

- Takes the approach of describing the simplest multi-class, multi-provider network scenario (i.e., a single end to end low loss and low latency class (Class 0) which is offered as a service to customers, in addition to the best effort class). This does not limit the same approach being applied to multi-provider end to end services with IP QoS classes other than the two classes described.
- Assumes that most regional (intra continental) end to end connections will span no more than 3 core provider networks.
- Identifies "common practices" that could ease the deployment of inter-provider QoS if adopted by a critical mass of providers

While there is continued debate about how many service classes need to be supported across multiple providers, it is widely agreed that some moderate number of classes should be commonly supported and consistently defined among providers. This Supplement uses the two class scenario as it is the simplest multi-class service offering. This is done as a way of exposing the issues that must be addressed. The additional service class that is defined is intended to be suitable for real-time voice applications; and is intended to be appropriate for use both in a provider-provisioned VPN context and in public networks. It is noted that, in many cases, providers may internally make use of an additional class of service that is restricted to network control traffic (such as routing protocol traffic and network management traffic).

The key issues that are addressed in this Supplement are:

**Consistent Definitions of Metrics.** To support QoS meaningfully across multiple providers, it is essential the metrics such as delay, delay variation and loss are defined consistently.

**Service Class Definition.** The "Class 0" service class is defined in terms of what the customer must do to receive the service (e.g. mark packets with a certain DSCP, conform to a certain token bucket) and what the provider in turn commits to deliver (e.g. statistical upper bounds on loss and delay). Although this Supplement only outlines detailed criteria for a single class of service beyond best effort, its goal is to remain flexible so that additional classes of service may be added. Furthermore, any individual Service Provider is free to offer additional service classes beyond those defined here.

**Measurement, Monitoring and Reporting.** Because of the multiple parties involved in the delivery of QoS, it is necessary to have defined methods for measurement of QoS, ways to monitor the performance of different network segments, and ways to report performance consistently among providers. Such methods are defined in this Supplement.

**Routing.** It may be necessary to route QoS-sensitive traffic to different providers or along different routes than those followed by best effort traffic. Routing considerations are discussed.

**Provider Responsibilities.** There may need to be some agreed-upon responsibilities and "common practices" to which providers should agree. A set of such practices are proposed that simplify deployment of inter-provider QoS among a large set of providers.

## 1 Introduction

Quality of Service (QoS) technologies based on the Differentiated Services architecture outlined in RFC 2475 is now widely deployed within the networks of many service providers. This is especially the case for providers of network-based VPNs (see, for example RFC 2547, RFC 4364). Some providers are now beginning to interconnect with each other via "QoS-enabled peering" in an attempt to offer QoS that spans the networks of multiple providers. However, in the absence of appropriate standards and established procedures for management, trouble-shooting, monitoring, etc., such interconnections have proven to be challenging and labour-intensive. This Supplement seeks to identify the key issues that service providers need to agree upon if inter-provider QoS is to be readily deployable.

### 1.1 Scope

It is the intent of this Supplement to discuss guidelines that are applicable to the interconnection of VPN's spanning multiple service providers. Because QoS deployment is much better established in private networks than in public networks, VPN provider interconnection is the primary focus, but the intent is to outline solutions that may also be applicable to public network service interconnection as well. Specific guidelines for public network service interconnection are for further study.

This Supplement is primarily concerned with the network support for two service classes across multiple providers' networks. These two classes being;

- A service class suitable for latency and loss insensitive data (Class 5)
- An additional service class, considered more suitable for the transport of conversational Voice over IP (VoIP) and other latency and loss-sensitive applications (Class 0)

A key distinguishing characteristic of private network services, compared to public network services, for services guaranteeing low packet loss and delay, is that the Service Provider(s) may not need to implement any means of avoiding congestion occurring for traffic in this class. This is discussed further in section 3.10.

Within the VPN context, it is likely that many VPN providers will deliver a service based on RFC 4364 (BGP/MPLS VPNs, formerly known as 2547 VPNs<sup>1</sup>). This Supplement will not restrict itself to BGP/MPLS VPNs - any IP VPN service should be supported - but the specific QoS issues of interconnecting providers of BGP/MPLS VPNs are addressed, including the MPLS-based interconnection styles (referred to as options (B) and (C) in RFC 4364).

### 1.2 Relationship to standards

This work draws heavily on the efforts of both the IETF (particularly the IPPM working group) and ITU-T Recs.Y.1540, Y.1541 and Y.1543.

---

<sup>1</sup> RFC 2547, which was the informational RFC that described MPLS/BGP VPNs, was superseded by the standards track RFC 4364.

The focus of this supplement is on practical operational methods of estimating and managing end to end metrics for services across multiple providers' networks. This focus has led to a pragmatic IP Delay Variation concatenation method that differs from that provisionally outlined in ITU-T Rec.Y.1541.

## 2 Reference model & terminology

### 2.1 Definitions

**Access :** That part of an end to end IP path from the customer's side of the CE router to the customer's side of the first PE router.

**ASBR:** Autonomous System Border Router. The router at the edge of an autonomous system (AS), facing towards another AS. ASBRs will typically be located at inter-provider boundaries, and may also be at AS boundaries that are within a single provider when a provider has chosen to divide his network into several ASes.

**CE:** Customer Edge. The router at the edge of a customer's network, usually facing towards a provider.

**Class 0:** An IP class of service for loss and delay sensitive applications and defined in ITU-T Rec.Y1541.

**Core:** That part of a provider's network from the customer side of the PE router to the customer's side of the distant PE router, or the mid point of the ASBR – ASBR network to network interface.

**CPI: Customer to Provider Interface.** The interface where a physical link interconnects a customer network and a single Provider's access network. This may also be referred to as a CE to PE interface or CE-PE link interface. ITU-T Rec. Y1541 refers to this interface as the UNI.

**Customer:** The user of the services provided by a service provider. In the context of IPVPNs, a customer typically exists at multiple physical locations, all of which are under one administrative authority, with each site connecting to one or more VPN Service Providers. In the context of public networks, customers typically connect to a Service Provider at one or more locations.

**Inter-provider Link:** The transmission link between two providers. Such a link typically interconnects a pair of ASBRs.

**Managed CE:** A Customer Edge device that is configured and managed by the provider on behalf of the customer.

**Measurement POP:** A service provider's point of presence (POP) that contains equipment capable of initiating and responding to measurement probes from another location.

**Network to Network Interface.** The point on an Inter-provider Link that represents the agreed demarcation point for service performance responsibility between the two different interconnecting Network Providers.

**Option A (B, C):** Methods for interconnection of MPLS VPNs across service provider (and AS) boundaries, defined in RFC 4364.

**P:** Provider routers. A Backbone Router, within a public network or VPN Service Provider(s) Network, that only attaches to PE's of the same Service Provider.

**PE:** Provider Edge router. The router at the edge of a provider's network, usually facing towards a customer.

**Provider:** a single public network and/or VPN Service Provider. In the context of this Supplement, more than one Provider is required to deliver an end-to-end Quality of Service IP path for the service class(es) defined herein.

**Trust Boundary:** The line between two entities that do not fully trust each other. A CE-PE link is a typical example of a trust boundary because the provider does not trust the customer to configure his equipment correctly or to stay within his SLA parameters. Conversely an internal link inside a single provider's network is usually not a trust boundary.

**Unmanaged CE:** A Customer Edge router that is managed by the customer, rather than by a provider.

**Acronyms used herein include the following:**

|                 |                                                                               |
|-----------------|-------------------------------------------------------------------------------|
| <b>AFI</b>      | Address Family Identifier                                                     |
| <b>AS</b>       | Autonomous System                                                             |
| <b>ASBR</b>     | Autonomous System Border Router                                               |
| <b>ATM PVC</b>  | Asynchronous Transfer Mode Permanent Virtual Circuit                          |
| <b>BGP</b>      | Border Gateway Protocol                                                       |
| <b>CE</b>       | Customer Edge                                                                 |
| <b>CDF</b>      | Cumulative Distribution Function                                              |
| <b>CPI</b>      | Customer to Provider Interface                                                |
| <b>Diffserv</b> | Differentiated services                                                       |
| <b>DSCP</b>     | Differentiated Services Code Points                                           |
| <b>E2E</b>      | End to End                                                                    |
| <b>EF</b>       | Expedited Forwarding                                                          |
| <b>EXP</b>      | Experimental (field in MPLS header that carries Class of Service information) |
| <b>FCAPS</b>    | fault-management, configuration, accounting, performance, and security        |
| <b>FR DLCI</b>  | Frame Relay Data Link Connection Identifier                                   |
| <b>GigE</b>     | Gigabit Ethernet                                                              |
| <b>IP</b>       | Internet Protocol                                                             |

|               |                                           |
|---------------|-------------------------------------------|
| <b>IPDV</b>   | IP Delay Variation                        |
| <b>IPLR</b>   | IP Loss Ratio                             |
| <b>IPPM</b>   | IP Performance Metrics                    |
| <b>IP-QoS</b> | IP Quality of Service                     |
| <b>IPTD</b>   | IP Transfer Delay                         |
| <b>IPVPN</b>  | Internet Protocol Virtual Private Network |
| <b>LL</b>     | Low Latency                               |
| <b>LSP</b>    | Label-Switched Path                       |
| <b>MPLS</b>   | Multi-Protocol Label Switching            |
| <b>MPoP</b>   | Measurement Point of Presence             |
| <b>NNI</b>    | Network to Network Interface              |
| <b>NMS</b>    | Network Management System                 |
| <b>NLRI</b>   | Network Layer Reachability Information    |
| <b>OAM</b>    | Operational and Management                |
| <b>OPSEC</b>  | Operational Security Capabilities         |
| <b>PDU</b>    | Protocol Data Unit                        |
| <b>PE</b>     | Provider Edge                             |
| <b>PHB</b>    | Per Hop Behaviour                         |
| <b>PWE3</b>   | Pseudo-wire emulation edge to edge        |
| <b>RFC</b>    | Request For Comment                       |
| <b>RPSEC</b>  | Routing Protocol Security Requirements    |
| <b>SAFI</b>   | Subsequence Address Family Identifier     |
| <b>SLA</b>    | Service Level Agreement                   |
| <b>SP</b>     | Service Provider                          |
| <b>TOS</b>    | Type Of Service                           |
| <b>TWAMP</b>  | Two-way Active Measurement Protocol       |
| <b>UTC</b>    | Coordinated Universal Time                |
| <b>VC</b>     | Virtual Circuit                           |
| <b>VoIP</b>   | Voice over IP                             |
| <b>VPN</b>    | Virtual Private Network                   |

## **2.2 Reference approach**

The key practice underpinning this Supplement is that DSCP in the IP header is the default definitive indicator to all providers in the end to end path of the underlying QoS treatment an IP

packet is expected to receive. Providers may use parameters in other protocol headers to convey QoS treatment required (e.g. where encapsulation of the IP packet occurs) but in the event that these markings differ from the QoS class indicated by the DSCP parameter, the later will be the definitive indication of the treatment expected by the customer.

### 2.3 Basic reference case

For simplicity, the single provider case depicted in Figure 1 is initially considered. In this model customer sites connect to the provider via a CE (customer edge) device, and the provider's routers that connect to customer sites are PE (provider edge) devices.




**Figure 1. Basic Diffserv Model for Single Provider**

The Customer Interface is the interface between the customer equipment and the CE. The CE can therefore prioritise and police traffic put onto the CE to PE link to differentially treat various classes of traffic. The basic single provider reference model places no restrictions on the mechanisms that are deployed by the provider within his access and core networks. Services are defined in terms of externally measurable performance parameters (e.g. loss, delay), with the mechanisms for achieving those performance targets left to the single provider.

Two CE models typically occur in the marketplace and these impact how and where a provider may manage the performance of its service to its customers even in the single provider case.

#### 2.3.1 Customer Managed CE model

In the customer-managed CE model, it is the responsibility of the customer to ensure that the traffic that traverses the CE to PE link is "correctly" marked before it reaches the PE. "Correct" in this context simply means that the customer needs to ensure that packets are marked in a way that ensures they receive the service desired. For example, if the customer has subscribed to a Class 0

service and the provider/customer contract for this service dictates that packets must be marked "EF" to receive the service, then the customer must decide which of his packets are to receive the Class 0 service and mark them before they arrive at the PE. The selection of packets to receive the Class 0 service is thus entirely up to the policies of the customer.

The PE, and/or other devices in the access transport network, may enforce various aspects of the service performance contract, such as policing the amount of Class 0 traffic received from, or sent to, a given customer. The details of such policing will be an aspect of a provider's service definition that enables the provider's commitments for class 0 to be managed and therefore met.

### 2.3.2 Provider Managed CE model

When the service provider manages the CE devices on behalf of the customer, it is possible to move the trust boundary to the CE. This means that the CE rather than the PE can be responsible for ensuring that the amount of traffic sent on an access IP path for any service class does not exceed the contract for that service class. This may be achieved by policing or shaping of the customer traffic before sending it to the PE.

## 2.4 Inter-provider reference case

Figure 2 illustrates the Diffserv Model extended to a simple inter-provider scenario. Its main difference from Figure 1 is that there are now two providers in the path between the two sites of each customer. The connection between the two providers is referred to as the Inter-provider link..




**Figure 2. Basic Inter-provider Model**

When the problem of delivering a particular service (e.g. the "Class 0" service) to customer P is considered, several issues that were not present in the single provider case must be addressed, including:

- Packets should be "correctly" marked on the inter-provider link to obtain the desired onward service, and the providers may have different markings for a given service.
- It may be desirable to carry that marking in a manner that avoids modification of the customer's data packets, e.g. in an extra header.

- Providers A and B should each offer a service that, when concatenated with the service of the other provider, provides a useful end-to-end service to the customer (e.g., for a service with a fixed maximum delay, the allowable delay may need to be agreed between both providers).
- Monitoring the end-to-end performance experienced by the customer is now likely to involve both providers in the path.

Marking on inter-provider links is the subject of the following section. Service definitions are discussed in Section 3. QoS measurement issues are discussed in Section 4.

It is desirable to support a wide range of interconnection methods. It should be possible to support a simple IP interconnect (which would include "option A" interconnection of RFC4364 VPNs) as well as MPLS interconnects of both option B and option C styles defined in RFC4364. Interconnection using MPLS traffic engineered LSPs should also be possible. It should also be possible to support any sort of layer 2 interconnect (e.g. ATM, Ethernet, etc.). The encapsulation and style of interconnection used at the inter-provider boundary has consequences on the marking and policing requirements, discussed below.

#### 2.4.1 Provider Managed CE model with multiple providers

The managed CE case is more complicated when there are multiple providers as in Figure 2, for example, if customer P purchases a managed service from provider A, who manages all of that customer's CEs, then the link between customer P and provider B still represents a trust boundary, while the link between customer P and provider A does not. In summary, the management of CEs by providers may or may not cause trust boundaries to be different than in the unmanaged CE case.

### 2.5 DSCP Marking

It is recommended that customer packets egress a provider's network with the same DSCP marking they had on ingress. If a network provider modifies any DSCP values within their network then the network provider should implement a means to accurately restore the original value of the packet's marking at the egress of their network.

When a packet ingresses an access provider's network from an unmanaged CE, the packet marking on egress from that providers network at any NNI should match the class that the packet should be carried in by the subsequent network provider(s). Where the ingress access provider remarks the packet in their network (if the packet is non conforming to the service contract and the provider elects to remark) is the access provider's choice.

It is often necessary for the provider to impose a QoS treatment on customer packets that differs from that which might be indicated by the customer's DSCP. For example, a customer may have an SLA that allows him to send traffic with DSCP=X up to a rate  $r$ , with excess packets being downgraded to best effort. However even if the packets are treated as best effort by the provider, the customer wishes to retain the DSCP marking of X for his own use when the packets arrive at his remote site. In the single provider environment, this capability is readily provided by encapsulating the customer's data with a header that exists only in the service provider network, e.g. an MPLS label header. This header is used to carry the Service provider's desired marking for the traffic, while leaving the customer's headers intact.

When there are multiple providers in the path, as in Figure 2, the marking issue is slightly more complex. Packets need to be marked appropriately to receive the desired service from the provider on the receiving side of the link. (That is, packets from provider A need to arrive at the edge of provider B with an appropriate marking for the desired service.) In options B and C, or when MPLS-TE is used across the inter-provider boundary, the MPLS EXP header may be able to be

used to carry the marking, thus leaving the customer header unchanged. In option A or pure IP interconnection it may be possible to encode the marking in a layer-2 dependent way to again leave the customer header unchanged. For example, an Ethernet 802.1q tag may be used to carry the marking across the boundary, or multiple ATM VCs may be used, one per service, with provider A placing the packets on the appropriate VC to receive the desired service from provider B and vice versa. Such border arrangements can be established through bi-lateral agreement between two interconnecting providers, however when three or more providers are in the end to end path such practices become very challenging to reliably support and therefore should be avoided.

## **2.6 Routing**

In a network as simple as that shown in Figure 2, there are no real routing issues since there is only one path between any two customer sites. However it is clear that in a true multi-provider environment there may be many alternate paths between customer sites. The preferred path among providers is typically determined by BGP policies. However, when multiple classes of service exist, it may be desired to route some traffic preferentially via providers who support the enhanced QoS class(es) while best effort traffic takes the conventional route. This issue is discussed further in Section 5.

## **2.7 Measurement and management**

Performance measurement and management is both important and challenging in the inter-provider QoS context because of the number of potential providers in the path between two customer sites. In the single provider case, a customer can conduct performance measurement between CEs; if the performance targets are not met, it can be assumed that the problem lies with the provider (unless of course the customer has overbooked and thus congested the access links). Even in a network as simple as the one shown in Figure 2, there are now five possible locations of performance problems for a given site-site pair: the two access links, within the core networks of each of the two providers, and the inter-provider link.

In order to deal with troubleshooting and performance monitoring issues, QoS measurement needs to be addressed as part of an inter-provider QoS solution. This topic is addressed in detail in Section 4 and draws on ITU-T Rec. Y1543.

# **3 Service Class**

## **3.1 Network Service Classes**

A best effort service is the default service class that is assumed to be available everywhere. This Supplement is primarily concerned with the support for an additional service class.

Service classes are treated here in terms of "black-box" behaviours – that is, the externally visible attributes of the service are considered (e.g. loss, delay) rather than internal network implementation mechanisms (e.g. Diffserv PHBs). In this respect service classes are similar to the Diffserv concept of a Per-Domain Behavior (PDB) defined in RFC 3086, but the scope of a service is not limited to a single domain. To be precise, a service class is defined by the externally visible treatment that the packets in that class receive as they traverse a network (in terms of loss, delay, and delay variation, and potentially transfer capacity). There may be additional aspects of a service class definition such as a default marking for packets in that service class.

The two network classes considered here are two of the classes defined in ITU-T Rec. Y.1541. These are;

- Class 5 (a “default best efforts” class)