

International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Series E.800

Supplement 8
(11/2009)

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

Guidelines for inter-provider quality of service

ITU-T E.800-series Recommendations – Supplement 8

ITU-T E-SERIES RECOMMENDATIONS
OVERALL NETWORK OPERATION, TELEPHONE SERVICE, SERVICE OPERATION AND HUMAN FACTORS

INTERNATIONAL OPERATION

Definitions	E.100–E.103
General provisions concerning Administrations	E.104–E.119
General provisions concerning users	E.120–E.139
Operation of international telephone services	E.140–E.159
Numbering plan of the international telephone service	E.160–E.169
International routing plan	E.170–E.179
Tones in national signalling systems	E.180–E.189
Numbering plan of the international telephone service	E.190–E.199
Maritime mobile service and public land mobile service	E.200–E.229

OPERATIONAL PROVISIONS RELATING TO CHARGING AND ACCOUNTING IN THE INTERNATIONAL TELEPHONE SERVICE

Charging in the international telephone service	E.230–E.249
Measuring and recording call durations for accounting purposes	E.260–E.269

UTILIZATION OF THE INTERNATIONAL TELEPHONE NETWORK FOR NON-TELEPHONY APPLICATIONS

General	E.300–E.319
Phototelegraphy	E.320–E.329
ISDN PROVISIONS CONCERNING USERS	E.330–E.349
INTERNATIONAL ROUTING PLAN	E.350–E.399

NETWORK MANAGEMENT

International service statistics	E.400–E.404
International network management	E.405–E.419
Checking the quality of the international telephone service	E.420–E.489

TRAFFIC ENGINEERING

Measurement and recording of traffic	E.490–E.505
Forecasting of traffic	E.506–E.509
Determination of the number of circuits in manual operation	E.510–E.519
Determination of the number of circuits in automatic and semi-automatic operation	E.520–E.539
Grade of service	E.540–E.599
Definitions	E.600–E.649
Traffic engineering for IP-networks	E.650–E.699
ISDN traffic engineering	E.700–E.749
Mobile network traffic engineering	E.750–E.799

QUALITY OF TELECOMMUNICATION SERVICES: CONCEPTS, MODELS, OBJECTIVES AND DEPENDABILITY PLANNING

Terms and definitions related to the quality of telecommunication services	E.800–E.809
Models for telecommunication services	E.810–E.844
Objectives for quality of service and related concepts of telecommunication services	E.845–E.859
Use of quality of service objectives for planning of telecommunication networks	E.860–E.879
Field data collection and evaluation on the performance of equipment, networks and services	E.880–E.899
OTHER	E.900–E.999

INTERNATIONAL OPERATION

Numbering plan of the international telephone service	E.1100–E.1199
---	---------------

For further details, please refer to the list of ITU-T Recommendations.

Supplement 8 to ITU-T E.800-series Recommendations

Guidelines for inter-provider quality of service

Summary

Supplement 8 to the ITU-T E.800-series of Recommendations 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). It presents an approach to the deployment of inter-provider quality of service (QoS) to enable further consideration of the topic. This Supplement discusses key issues that service providers need to agree upon if inter-provider QoS is to be readily deployable.

Source

Supplement 8 to ITU-T E.800-series Recommendations was agreed on 12 November 2009 by ITU-T Study Group 12 (2009-2012).

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 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 2010

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

CONTENTS

Page

1	Introduction	2
1.1	Scope	2
1.2	Relationship to standards.....	3
2	Reference model and terminology.....	3
2.1	Definitions	3
2.2	Reference approach	5
2.3	Basic reference case	5
2.4	Inter-provider reference case.....	6
2.5	DSCP marking.....	7
2.6	Routing	8
2.7	Measurement and management	8
3	Service class	9
3.1	Network service classes.....	9
3.2	Customer to provider interface (CPI) behaviour	10
3.3	Network-to-network interface (NNI) behaviour	11
3.4	Definitions of metrics	12
3.5	Class 0 service definition.....	14
3.6	Impairment budgets	14
3.7	Proposed allocation approach.....	16
3.8	Assessment reference model	17
3.9	Budget allocation.....	18
3.10	Admission management	21
4	QoS measurement.....	22
4.1	QoS measurement requirements.....	22
4.2	QoS measurement methodologies	24
4.3	QoS measurement protocols.....	24
4.4	Measurement considerations for VPN services.....	25
4.5	Recommended measurement approach	25
4.6	Performance measurement metrics.....	26
4.7	Reporting of measurement results	27
4.8	QoS measurement security considerations.....	28
5	Routing	28
5.1	Current BGP capabilities.....	29
5.2	BGP considerations	29
5.3	Routing recommendation	30
6	Securing QoS.....	30
6.1	Motivation	30
6.2	Areas which need to be secure	30



Documents

Document Sharing Hub

E.800-series – Supplement 8 (11/2009)

iii

	Page
6.3 Provisioning security	30
6.4 Service security	33
6.5 Security guidelines	33
7 Operational issues	34
7.1 Fault	35
7.2 Configuration and maintenance	35
7.3 Accounting	37
7.4 Performance	37
8 References	38
Appendix I – Examples of the application of budget allocations	40
I.1 Case 1: Three core providers	40
I.2 Case 2: Transcontinental service, five core providers	40

Supplement 8 to ITU-T E.800-series Recommendations

Guidelines for inter-provider quality of service

Executive summary

This Supplement presents an approach to the deployment of inter-provider quality of service (QoS). This Supplement begins from the observation that QoS, based on the differentiated services architecture defined in [IETF 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 [IETF RFC 2547], and [IETF 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 discusses key issues that service providers need to agree upon if inter-provider QoS is to be readily deployable.

NOTE – As this Supplement is a vehicle for sharing research topics under consideration, there may be some conflicts between this Supplement and existing ITU-T Recommendations. In such conflicts, the reader is reminded that Supplements are only informative and are therefore not considered to be an integral part of any Recommendation; they do not imply any agreement on the part of the ITU-T.

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 that 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.



Documents

Document Sharing Hub

E.800-series – Supplement 8 (11/2009)

1

- **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 is proposed that simplifies 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 [IETF 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 [IETF RFC 2547], and [IETF 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 VPNs 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 clause 3.10.

Within the VPN context, it is likely that many VPN providers will deliver a service based on [IETF RFC 4364] (BGP/MPLS VPNs, formerly known as 2547 VPNs¹). 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 [IETF RFC 4364]).

¹ [IETF RFC 2547], which was the informational RFC that described MPLS/BGP VPNs, was superseded by the standards track [IETF RFC 4364].

1.2 Relationship to standards

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

The focus of this Supplement is on practical operational methods for 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 Y.1541].

2 Reference model and terminology

2.1 Definitions

This Supplement defines the following terms:

2.1.1 access: That part of an end-to-end IP path from the customer's side of the customer edge (CE) router to the customer's side of the first provider edge (PE) router.

2.1.2 autonomous system border router (ASBR): 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 ASs.

2.1.3 customer edge router (CE router): The router at the edge of a customer's network, usually facing towards a provider.

2.1.4 Class 0: An IP class of service for loss and delay sensitive applications as defined in [ITU-T Y.1541].

2.1.5 core: That part of a provider's network from the customer side of the provider edge (PE) router to the customer's side of the distant PE router, or the mid-point of the ASBR – ASBR network-to-network interface.

2.1.6 customer to provider interface (CPI): The interface where a physical link interconnects a customer's 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 Y.1541] refers to this interface as the UNI.

2.1.7 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.

2.1.8 inter-provider link: The transmission link between two providers. Such a link typically interconnects a pair of ASBRs.

2.1.9 managed CE device: A customer edge device that is configured and managed by the provider on behalf of the customer.

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

2.1.11 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.

2.1.12 option A (B, C): Methods for interconnection of MPLS VPNs across service provider (and AS) boundaries, defined in [IETF RFC 4364].

2.1.13 provider router (P router). A backbone router, within a public network or VPN service provider(s) network, that only attaches to PE routers of the same service provider.

2.1.14 provider edge router (PE router). The router at the edge of a provider's network, usually facing towards a customer.

2.1.15 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.

2.1.16 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.

2.1.17 unmanaged CE router: A customer edge router that is managed by the customer, rather than by a provider.

Abbreviations

This Supplement uses the following abbreviations:

AFI	Address Family Identifier
AS	Autonomous System
ASBR	Autonomous System Border Router
ATM PVC	Asynchronous Transfer Mode Permanent Virtual Circuit
BGP	Border Gateway Protocol
CDF	Cumulative Distribution Function
CE	Customer Edge
CPI	Customer to Provider Interface
Diffserv	Differentiated services
DSCP	Differentiated Services Code Point
E2E	End-to-End
EF	Expedited Forwarding
EXP	Experimental (field in MPLS header that carries Class of Service information)
FCAPS	Fault-management, Configuration, Accounting, Performance, and Security
FR DLCI	Frame Relay Data Link Connection Identifier
GigE	Gigabit Ethernet
IP	Internet Protocol
IPDV	IP Delay Variation
IPLR	IP Loss Ratio
IPPM	IP Performance Metrics
IP-QoS	IP Quality of Service
IPTD	IP Transfer Delay
IPVPN	Internet Protocol Virtual Private Network
LL	Low Latency

LSP	Label-Switched Path
MPLS	MultiProtocol Label Switching
MPoP	Measurement Point of Presence
NLRI	Network Layer Reachability Information
NMS	Network Management System
NNI	Network-to-Network Interface
OAM	Operational and Management
OPSEC	Operational Security Capabilities
OWAMP	One-Way Active Measurement Protocol
PDU	Protocol Data Unit
PE	Provider Edge
PHB	Per-Hop Behaviour
PWE3	Pseudo-Wire Emulation Edge to Edge
RPSEC	Routing Protocol Security Requirements
SP	Service Provider
SAFI	Subsequence Address Family Identifier
SLA	Service Level Agreement
ToS	Type of Service
TWAMP	Two-Way Active Measurement Protocol
TWPD	Two-Way Packet Delay
TWPL	Two-Way Packet Loss
UTC	Coordinated Universal Time
VC	Virtual Circuit
VoIP	Voice over IP
VPN	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 latter 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 a single provider

The customer interface is the interface between the customer equipment and the CE router. The CE router can therefore prioritize and police traffic put onto the CE to PE link to treat differently 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 device model

In the customer-managed CE device 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 router, 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 device 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 router, rather than the PE router, 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 clause 2.5. Service definitions are discussed in clause 3. QoS measurement issues are discussed in clause 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 [IETF RFC 4364] VPNs) as well as an MPLS interconnect of both option B and option C styles defined in [IETF RFC 4364]. 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, as discussed below.

2.4.1 Provider-managed CE device model with multiple providers

The managed CE device 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 provider's network at any NNI should match the class that the packet should be carried in by the subsequent network provider(s). Whether the ingress access provider re-marks 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 bilateral agreement between two interconnecting providers; however, when three or more providers are in the end-to-end path, such practices become very challenging to be reliably supported 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 clause 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, 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 clause 4 and draws on [ITU-T Y.1543].

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 behaviour (PDB) defined in [IETF 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 potential 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 Y.1541]. These are:

- Class 5 (a "default best effort" class);
- Class 0.

3.1.1 Class 5

Traffic that has not been explicitly identified and associated with another service class will receive Class 5 ("best effort") treatment. Class 5 has no explicit guarantee with respect to latency, delay variation, or packet loss; however, providers typically do endeavour to provide for satisfactory delivery of packets in this service class, and provider service commitments for Class 5 are not uncommon.

3.1.2 Class 0

The additional service class in this Supplement is Class 0, which is intended to be used for the transmission of services that require low delay, low delay variation and low loss between customers of different service providers, and where the traffic between the customers traverses two or more disparate carriers. The class is intended to be suitable for real-time applications such as VoIP, but there is no restriction on which applications may actually use the service. Mapping of applications to service classes is left to the customers.

For many applications, Class 0 will carry traffic bidirectionally (e.g., media packets in both directions of an individual VoIP call). However, there is no requirement to provide a symmetric path for the bidirectional traffic flow between any given source and destination.

3.1.3 Other service classes

Service providers are at liberty to offer any number of service classes above and beyond those defined in this Supplement. Indeed, it is typical to offer four or more service classes to customers and also to use one or more internal classes for network control (e.g., routing protocol) traffic. It is expected that more classes will probably need to be agreed upon for inter-provider use at some point in the future, but the discussion of additional classes is for further study. As noted above, even agreeing on how to support one additional class (i.e., Class 0), beyond the default Class 5, raises many, if not most, of the hard problems that need to be addressed.

Even if there were a larger number of "standard" service classes that could be offered in an inter-provider context, it is likely that providers would continue to offer some additional classes beyond the standard set as a means of competitive differentiation. An inter-provider QoS model should allow for such flexibility.