International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1315

(09/2006)

SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS

Internet protocol aspects – Transport

QoS support for VPN services – Framework and characteristics

ITU-T Recommendation Y.1315

ITU-T

International Telecommunication Union

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	**Y.1300–Y.1399**
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Numbering, naming and addressing	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899

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

ITU-T Recommendation Y.1315

QoS support for VPN services – Framework and characteristics

Summary

ITU-T Recommendation Y.1315 describes the characteristics and applicability of various QoS (quality of service) service scenarios in the VPN (virtual private network) environment. The QoS service scenarios cover various QoS models in single or multiple service provider configurations. The QoS models described in this Recommendation encompass the pipe model, the hose model, and a combination and/or concatenation of both.

Source

ITU-T Recommendation Y.1315 was approved on 13 September 2006 by ITU-T Study Group 13 (2005-2008) under the ITU-T Recommendation A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. 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 Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation 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 Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation 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 Recommendation development process.

As of the date of approval of this Recommendation, ITU had received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. 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 2007

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

CONTENTS

Introduction

There are several QoS models available for VPN services. In this Recommendation, the pipe model and the hose model are studied. The objectives of such studies are to provide precise definitions for these QoS models, to construct the VPN using the toolset provided in [ITU-T Y.1291], and to outline applicability of these QoS models in the VPN environment.

The architectural components needed to build these QoS models are defined in [ITU-T Y.1291] while the mechanisms to coordinate the resources to ensure QoS between sites are defined in the related IETF and ITU-T documents.

In order to provide VPN services with QoS, the first step is to define the QoS in parametric terms. This is construed as objectives of the services to be provided. The next step is to use the architecture building blocks provided in [ITU-T Y.1291] to construct the VPN with QoS, and the final step is to evaluate the service performance provided by the network. This Recommendation refers to [ITU-T Y.1291] on the functions and features of the QoS architectural components.

This Recommendation takes input from the various VPN related ITU-T deliverables, including approved Recommendations (e.g., [ITU-T Y.1311], [ITU-T Y.1311.1], [ITU-T Y.1312], and [ITU-T Y.1314]).

ITU-T Recommendation Y.1315

QoS support for VPN services – Framework and characteristics

1 Scope

The scope of this Recommendation is to provide guidelines on how the QoS architecture tools, mostly provided in [ITU-T Y.1291], can be used to support VPN services with QoS characteristics between the network ingress points to the network egress points. In this Recommendation, different QoS service models, including the pipe model and the hose model, are studied. The characteristics and applicability of VPN services using those QoS models and concatenation of multiple QoS models in a single or multiple provider scenarios are also studied, including some application examples.

While this Recommendation discusses the QoS issues related to VPN applications, the mechanisms to ensure QoS assurance are out of the scope of this Recommendation.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[ITU-T E.800] ITU-T Recommendation E.800 (1994), *Terms and definitions related to quality of service and network performance including dependability.*

[ITU-T G.1000] ITU-T Recommendation G.1000 (2001), *Communications Quality of Service: A framework and definitions.*

[ITU-T H.360] ITU-T Recommendation H.360 (2004), *An architecture for end-to-end QoS control and signalling*.

[ITU-T Y.1291] ITU-T Recommendation Y.1291 (2004), *An architectural framework for support of Quality of Service in packet networks.*

[ITU-T Y.1311] ITU-T Recommendation Y.1311 (2002), *Network-based VPNs – Generic architecture and service requirements*.

[ITU-T Y.1311.1] ITU-T Recommendation Y.1311.1 (2001), *Network-based IP VPN over MPLS architecture*.

[ITU-T Y.1312] ITU-T Recommendation Y.1312 (2003), *Layer 1 Virtual Private Network generic requirements and architecture elements*.

[ITU-T Y.1314] ITU-T Recommendation Y.1314 (2005), *Virtual private network functional decomposition*.

3 Definitions

This Recommendation defines the following terms:

3.1 arrival curve: A mechanism to serve as a measure of traffic arriving at provider's ingress network interface.

3.2 service curve: A mechanism to serve as a measure of the packet forwarding services provided by the service provider.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms.

BE	Best Effort
CE	Customer Edge
CoS	Class of Service
CR-LDP	Constraint-based Routing – Label Distribution Protocol
Diffserv	Differentiated Services
DSCP	DiffServ Code Point
EF PHB	Expedited Forwarding Per Hop Behaviour
FIFO	First In, First Out
FILO	First In, Last Out
Intserv	Integrated Services
LAPS	Link Access Procedure – SDH
LIFO	Last In, First Out
LILO	Last In, Last Out
LSP	Label Switched Path
MPLS	Multi-Protocol Label Switching
P	Provider (node)
PE	Provider Edge
QoS	Quality of Service
RSVP	Resource reSerVation Protocol
SLA	Service Level Agreement
TE	Traffic Engineering
UNI	User Network Interface
VF	Virtual Forwarder
VPN	Virtual Private Network

5 QoS definitions

The concept of quality of service is a multi-dimensional measure of user satisfaction of the network. The various definitions of QoS have been discussed in different ITU-T Recommendations (e.g., [ITU-T E.800] and [ITU-T G.1000]). Under those Recommendations, the network performance characteristics, network operability, network support, network serviceability, and network security are all part of the quality of service.

In this Recommendation, only the parametric aspects of QoS under the VPN environment for packet-based networks are discussed. The following subclauses provide details about the QoS metrics.

5.1 Packet-network-related QoS parameters

Among various aspects of quality of service definitions, this Recommendation focuses on the performance aspects of the service provider's network in the VPN environment for the packet-based networks. The measure of the QoS for a particular service provider's packet network providing VPN services should encompass the following aspects:

- **Delay**: There are various means to measure the time delay experienced by the user traffic. It should be noted that those definitions are not equivalent (e.g., consider if the packet is fragmented inside the network or LAPS-based framing). It should be noted that the service providers are free to use any of those definitions but a consistent definition needs to be provided to the clients if multiple service providers are involved.
  - [FIFO delay] First-in, first-out: The time difference between the first bit accepted by the network and the first bit delivered to the user across the network.
  - [LILO delay] Last-in, last-out: The time difference between the last bit accepted by the network and the last bit delivered to the user across the network.
  - [FILO delay] First-in, last-out: The time difference between the first bit accepted by the network and the last bit delivered to the user across the network. It is possible to define the FILO delay but should not be used because it is packet size sensitive.
  - [LIFO delay] Last-in, first-out: The time difference between the last bit accepted by the network and the first bit delivered to the user across the network. It is possible to define the LIFO delay but should not be used because it is packet size sensitive.
- **Delay variation** (jitter): Once a specific type of delay is defined by the network, the delay variation may be further divided into the following sub-categories:
  - Maximum delay variation: The difference between the maximum delay and minimum delay experienced by the user traffic within a specific time window.
  - Mean delay variation: Within a specific time window [0,*T*], if each individual packet delay can be denoted as $d(k)$ for the kth packet in the time window, the mean delay variation can then be denoted as $D_{[0,T]} = \sum_{i=0}^{N} \left| d(i) - \overline{D} \right|$ where N denotes that there are N packets arriving in the time interval [0,*T*] and $\overline{D}$ is the average delay endured by each packet arriving in the time window [0,*T*] (i.e., $\overline{D} = \sum_{k=0}^{N} d(k)$ where N is the total number of packets arrived in the time interval [0,*T*]).
  - Medium delay variations: $(\overline{M})$: Within a specific time window [0,*T*], 50% of the packet have delay more than $\overline{M}$ and 50% packet will have delay less than $\overline{M}$.
- **Packet loss probability**: The definition of packet loss probability (as opposed to the pure packet loss counts as some service providers are considering) is also subject to the time window constraint. In this Recommendation, the "packet delivery probability" is defined as the total number of packets successfully delivered divided by the total number of the eligible packets received by the network for a given and well-defined time interval. The eligible packets are then subject to the input constraint. Hence the "packet loss probability" is one minus the "packet delivery probability".

  The particular value of delay, delay variation, and packet loss probability are not specified and are also application specific.

For delay measurement, the FIFO delay may be ensured analytically for arbitrary packets. Even the LILO delay measurement may be useful for applications; such a delay will be packet content sensitive and difficult to be ensured analytically.

The mean delay variation and medium delay variation depend on the statistics of the packet being forwarded. Given such statistics are unknown, and media characteristic sensitive, those two types of delay variations are not further pursued in this Recommendation. The focus of this Recommendation is on the maximum delay variation, which can serve as guidance for buffer design in media replay type applications.

5.2 Basic QoS mechanisms

There are essentially two mechanisms to satisfy SLA regarding QoS in the VPN environment, which can be used to construct pipes and hoses. One is based on the IntServ model and the other is based on the Diffserv model.

The IntServ mechanism, based on [b-RFC 2205] and its subsequent improvements, can be used to perform bandwidth reservation for a single packet flow or multiple packet flows. It uses the *path* message to advertise the reservation requirements and the *resv* message to actually make the reservation. Once the connection admission control allows the flow, the bandwidth will be dedicated to the flow and parametric QoS assurance is possible under the IntServ environment.

For packet-based networks, the use of Diffserv mechanisms is being deployed. When Diffserv mechanisms are used, the level of QoS a packet should receive from the network is indicated by the use of DSCP. The DSCP is being used to indicate the service class to which each packet belongs and the QoS treatment provided by the Diffserv capable packet network will be as indicated by the DSCP.

Differentiated service classes encourage the usage of the service class end-to-end, with the end-user application determining the correct service class to be used by specific application traffic flows.

For example, a traffic flow belonging to the "telephony" service class will be treated with minimum loss, jitter, and delay in the Diffserv EF PHB.

The usage of Diffserv service classes allows the communication of QoS requirements between the end user application and the packet network with minimum overhead.

6 Network input and service specification

In order to properly ensure the parametric QoS boundaries (with parameters given in clause 5.1) for the provider's network, the input traffic of the network and services performed by the service providers need to be characterized. Those characterizations are discussed in clauses 6.1 and 6.2.

6.1 Input traffic characteristics

In order to discuss the QoS in the VPN environment, the fluid model of the network will be used. In this set-up, the user input traffic will be considered as a continuous bit stream which will be received by the network. The cumulative number of bits received at the time instance t will be denoted as $R(t)$.

For a wide-sense increasing function α defined for $t \geq 0$, the function $R(t)$ is said to be constrained by α if and only if:

$$R(t) - R(s) \leq \alpha(t - s)$$

The function α is also said to be the arrival curve of the $R(t)$ or the $R(t)$ is α – *smooth*.

The popular token bucket (p,M,r,b) as defined in [b-RFC 2215], is essentially to limit the input function $R(t)$ by the arrival curve α with:

$$\alpha(t) = \min(M + pt, rt + b)$$

This arrival curve can be shown graphically as in Figure 6-1.

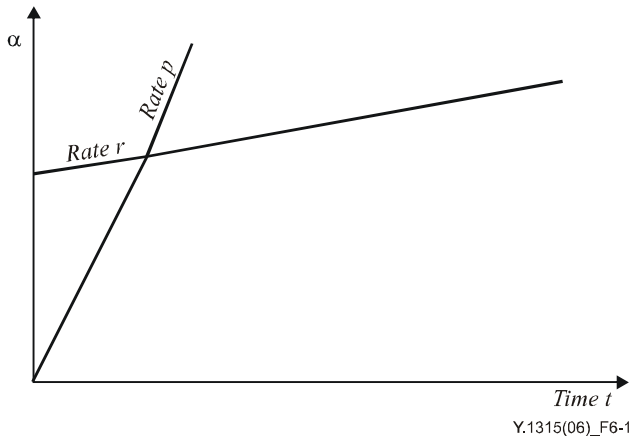

Figure 6-1 – Arrival curve for traffic limited by token bucket

6.2 Service specification

The service provided by the service provider can also be modelled via a "service curve".

For almost all the nodes, the service curve will be of the "rate-latency" in nature, which can be illustrated in Figure 6-2.

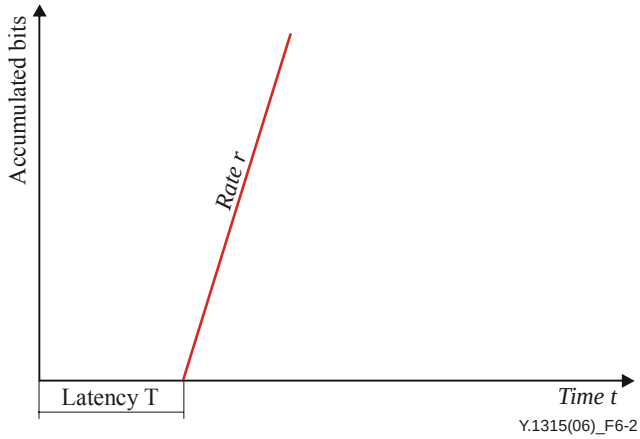

Figure 6-2 – Rate-latency service curve with service rate *r* and latency *T*

For N nodes with service rate r_i and latency with T_i, the resulting service curve will also be of the rate-latency type with rate $r = \min(r_1,...,r_N)$ and delay $T = T_1 + ... + T_N$.

In other words, as far as the service rate of the network is concerned, the service rate will be the minimum rate among all the nodes along the path while the delays are additive. It should be noted that the delay bounds obtained in considering the overall service curve are better than considering each node individually (known as "*Pay burst only once*" phenomenon).

The significance of the service curve is its ability to specify the service to be performed for a network element or a whole network, regardless of the type of buffer managements and scheduling mechanisms, and its ability to specify the concatenated behavior of a series of network elements or networks the user packet is to encounter.

The detailed definition of the service curve is introduced in Appendix II.

7 QoS architectural components

This Recommendation is based on the architectural components in [ITU-T Y.1291]. These architectural components are overviewed in this clause, the details being articulated in [ITU-T Y.1291].

This clause includes required enhancements for VPN applications to the architectural components presented in [ITU-T Y.1291].

7.1 Data plane architectural components

In this subclause, the architectural components for the data plane are considered.

7.1.1 Queue (or buffer) management

Queue or buffer management has been discussed in [ITU-T Y.1291], emphasizing various discard mechanisms when the queue is close to its capacity.

[ITU-T Y.1291] provides details for queue or buffer management. No special need has been recognized in the VPN environment regarding those components.

7.1.2 Congestion avoidance

Congestion can occur either inside the VPN, when network resources are partitioned and dedicated to the VPN, or inside the network when the resources are not partitioned. When congestion is detected, the congestion avoidance function may need to be triggered.

[ITU-T Y.1291] provides details for congestion avoidance.

7.1.3 Queuing and scheduling

Various queuing and scheduling mechanisms have been discussed in [ITU-T Y.1291], such as weighted fair queuing, priority queuing, and class-based queuing. Given the details that have been presented in [ITU-T Y.1291], no special requirements have been identified in the VPN environment.

7.1.4 Packet marking

Packet marking indicates the discard priority of the packet. It may be performed by the host, CEs, and PEs. The marking may be propagated inside the customer's network or the provider's network.

[ITU-T Y.1291] provides details for packet marking.

7.1.5 Traffic classification

Traffic classification classifies the traffic into different queuing priorities.

[ITU-T Y.1291] provides details for traffic classification.

Packets of different VPNs or even in the same VPN may have different forwarding rules. The service provider selects forwarding rules via the traffic classification functions. Forwarding rules can be controlled and negotiated by the control plane.

7.1.6 Traffic policing

Traffic policing deals with the determination of whether the traffic being presented is compliant with the pre-negotiated policies or contracts. Typically non-conformant packets are dropped. The senders may be notified of the dropped packets and cause.

The traffic is policed at the ingress of the network in accordance with the input characteristics outlined in clause 6.1. Even though a wide variety of the arrival curves may be used, only the token bucket arrival curve, shown in Figure 6-1, can be easily implemented at this time.

[ITU-T Y.1291] provides details for the traffic policing function.

7.1.7 Traffic shaping

Traffic shaping deals with controlling the VPN traffic profile. The traffic shaper at the ingress point of the network is responsible for traffic shaping, and buffers the non-compliant packets until they are in compliance with the required traffic ingress characteristics.

Traffic shaping often needs to be performed between ingress PEs and egress PEs. The traffic shapers may also need to be installed in the inter-provider boundaries. The location of the traffic shaper is at the discretion of service providers.

7.2 Control plane architectural components

In this clause, the architectural components for the control plane are considered. The control plane architectural components of [ITU-T Y.1291] are not VPN aware. In order to control the VPN properly, some enhancements to [ITU-T Y.1291] control plane components are proposed in this subclause.

7.2.1 Resource manager

The resource manager includes the "resource reservation" architectural component of [ITU-T Y.1291].

The resource manager manages the resources for a set of node(s). The set can range from a single node to the whole network.

The resource manager component may interact with many network elements involved in VPN so as to collect comprehensive network resource information.

Users inside a network may be partitioned into different resource acquisition priorities, with high priority users to acquire the potentially contentious resources at the cost of low priority users. It is the resource manager's responsibility to properly coordinate the resources and its users.

The customer network and provider network may both have their resource managers. Those resource managers control and manage the resources in customer networks and provider networks respectively, and they can communicate with each other to exchange customer network and provider network resource information, which is in order to guarantee CE-to-CE (customer edge-to-customer edge) QoS.

The resource manager is an extension of the "resource reservation" architectural component of [ITU-T Y.1291] for the VPN environment.

7.2.2 Admission control

In the established network with a fixed quantity of resources, it is impossible to admit infinite VPN service instances with strict QoS requirements, so only a limited number of these VPN service instances with QoS requirements can be admitted into the network. This function is achieved by an admission control building block. The admission control will use the resource information provided by the resource manager.

7.2.3 VPN QoS routing

QoS routing mechanisms, introduced in [ITU-T Y.1291], need to support the VPN environment, and VPN topology. QoS routing may also be resource aware and use the resource information to increase the probability of successful resource reservation.

This architectural component interacts with the resource manager to distribute the resource utilization information and to perform the path computation, with consideration of routing policies.

Inter-layering QoS routing is not a VPN specific topic. While it is still in its infancy, the inter-layering QoS routing is out of the scope of this Recommendation.

7.3 Management plane architectural components

The architectural components for the management plane are discussed in this clause.

7.3.1 Service level agreement management

A service level agreement (SLA) typically represents the agreement between a VPN customer and VPN provider that specifies the level of availability, serviceability, performance, operation or other attributes of that VPN service. It may include aspects such as pricing that are business in nature. In [b-RFC 3198] the technical part of the agreement is called the service level specification (SLS), which specifically includes a set of parameters and their values that together define the service offered to the VPN customer's traffic by the VPN provider's network.

Both QoS service models (pipe and hose) can be expressed in the SLS. SLAs should cover the single provider case and multiple provider case. In the latter scenarios, the QoS parameters may be partitioned among the service providers.

7.3.2 Traffic measurement

Traffic metering in VPN concerns monitoring of the temporal properties (e.g., rate) of VPN traffic stream against the agreed traffic profile. It involves observing traffic characteristics at a given network point and collecting and storing the traffic information for analysis and further action, such as the validation and verification of the QoS guarantee level.

It has to be noted that, in the various VPN service environments, traffic metering is not always possible. For VPNs using the hose model, the measurement on the CE-PE link can be performed. For point-to-point VPN traffic, measurement may be feasible at certain locations along the path of the flow.

7.4 Enhancements for Y.1291

There are two architecture components which require enhancements in order to provide QoS properly under the VPN environment:

1) Resource manager discussed in clause 7.2.1: the resource manager needs to be VPN aware so the collected data can be processed accordingly.
2) Traffic manager discussed in clause 7.3.2: the traffic manager needs to be VPN aware so the collected data can be processed accordingly.

8 VPN service based on the pipe model

In this clause, the pipe model used in VPN will be investigated. The pipe model is to ensure QoS on a point-to-point connectivity basis.

8.1 Definition of pipe model

The VPN customer specifies QoS requirements and/or assurances on each path between the VPN endpoints. At any given ingress point, if the customer needs to reach *n* locations, the QoS requirement for each of those *n* paths needs to be specified.

For each pair of VPN endpoints, the network provider needs to provide adequate service bandwidth (i.e., adequate service curve) along the path of each pipe to ensure that the QoS is satisfied. Resources made available to the pipe cannot be allocated to other traffic.

The customer should shape the traffic on CE delivered to the service provider on a per destination basis to avoid violation of the QoS agreement (e.g., see Figure 8-2 for details). A graphical illustration of the pipe model is shown in Figure 8-1.

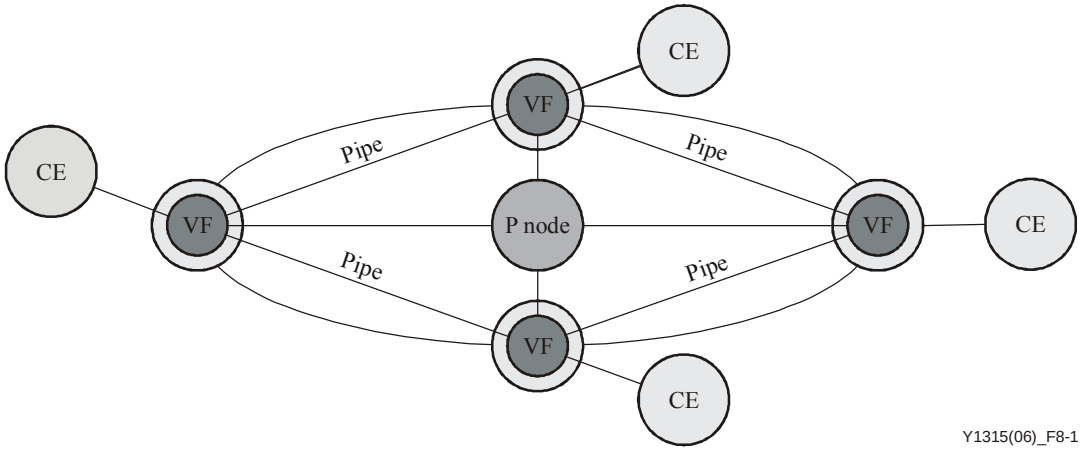

Figure 8-1 – Illustration of pipe model

In the above diagram, VF stands for virtual forwarder (VF), which contains the VPN-specific forwarding information, such as the VPN-based routing tables for IP VPNs.

A pipe may be constructed via RSVP signalling if QoS assurance is to be ensured or via relative priority if over-provisioning is practiced. There are two means to pin down the routes of the packets: one is via the connection-oriented mechanisms such as MPLS, while the other is to use source routes in the connectionless realm. The details of pipe construction, maintenance, and deletion are out of the scope of this Recommendation.

8.2 Customer edge characteristics

In order to support the pipe model, and in order to ensure QoS, for each destination, the CE may need to shape the traffic in accordance with the SLA (or input characteristics). Failing to perform such per-destination shaping may result in the violation of the agreed upon QoS assurance.

It should also be noted that, due to the dynamic nature of the packet routing, such per-destination-based shaping may not always be feasible. For example, a certain VPN user's address may change egress location on the provider's network due to routing table or routing policy modifications.

8.3 Provider equipment characteristics

In the PE, for each far-end destination, a policer (Figure 8-2) is required to ensure input traffic conformance. For the conforming traffic, with proper provisioning of the network resources, the QoS specified in the SLA can be ensured.

Based on the SLA, the non-conforming traffic may be tagged with higher discard probability or be discarded.

If the network provider has to provide the pipe model for VPN, there will be one policer per far-end destination on the ingress node to enforce the arrival curve. The network ingress point can be illustrated as in Figure 8-2.

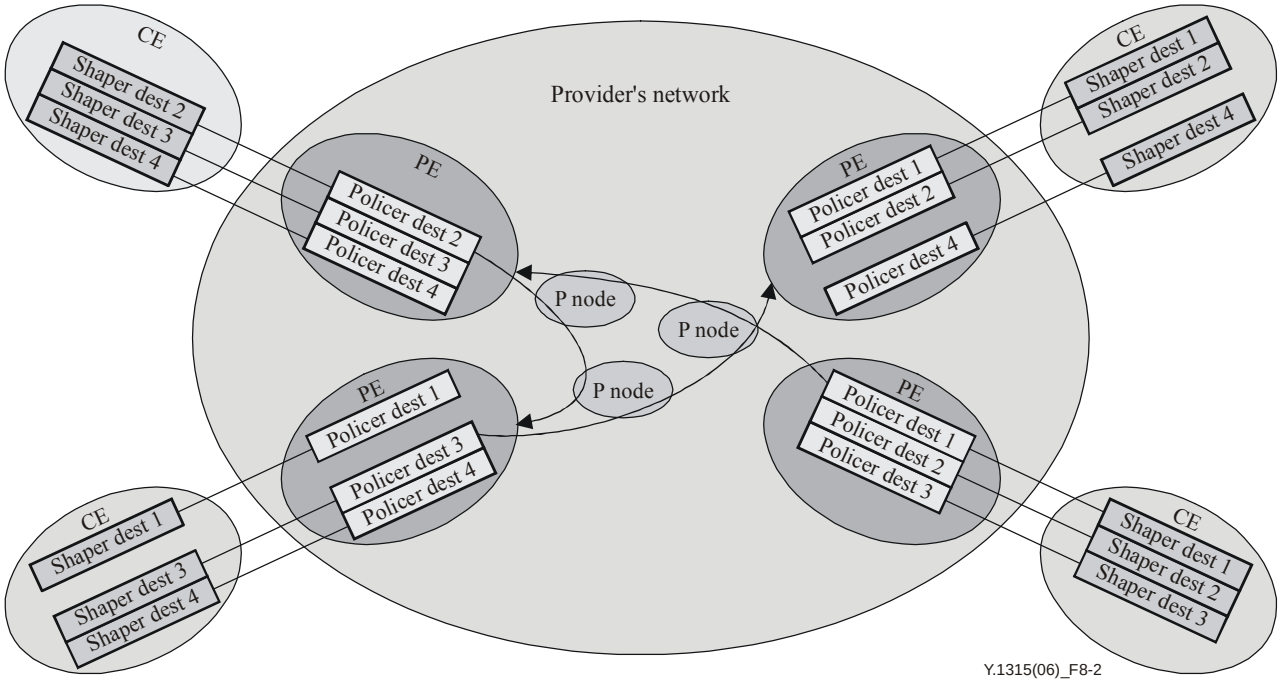

Figure 8-2 – Pipe model and traffic ingress policing

Because there is no further traffic aggregation inside the network for the conforming customer traffic, it is relatively simple to deliver the required QoS, characterized by packet delay, delay variation, and packet loss.

8.4 Applicability, complexity, and scalability

In this clause, the applicability, complexity, and scalability of the pipe VPN are addressed.

Applicability of the pipe VPN outlines the strength and limitations of this type of VPN. From the PE's perspective, each VPN's far endpoint requires a policer on the UNI (user to network interface) for non-BE traffic for traffic conformance purposes (known as arrival curve conformance). Furthermore, the policers are normally shared among a number of VPNs and the number of those policers is limited. In this case, the construction or augmentation (i.e., adding more endpoints) for a pipe VPN can only be carried out if and only if there are enough policers to enforce the traffic conformance.

Because the VPN traffic policing is performed at the edge, packet discard due to policer action only relates to the underlying UNI, and not relate to the activities on other UNIs, as compared to the hose model discussed in clause 9.4. It is sometimes possible to have the CE shape the traffic so that packet loss due to policer action can be avoided.

Network complexity discusses the issues related to VPN construction, deletion, and maintenance. For the VPNs based on the pipe model, the construction of the VPN involves construction of new pipes, resize or reuse of the existing pipes, alone with the required resource reservation, so the VPN endpoints are fully connected with the desired QoS parameters. VPN deletion triggers the release of the resources and the pipe usage, with the possibility of pipe removal. The maintenance of the pipe VPN is to manage the pipes inside the network properly so that the VPN endpoints are fully connected with the desirable QoS parameters.

The scalability of the pipe VPN is limited by the availability of policers and network processing power in managing the VPN. Because a large number of pipes are expected to exist inside the network, the size of the VPN is also limited by the resources at the physical interface and the underlying technology used to construct the VPN. It should be noted that the above statement is

universally true but, due to the special nature of the VPN based on the pipe model (n square connectivity), the network resource usage at the physical interfaces will be taxed to the maximum.

Due to the above-mentioned rationale, the pipe VPN will be applicable if the number of endpoints is small to moderate, and there is QoS assurance beyond BE traffic.

8.5 Constructing the VPN based on the pipe model

The pipe model can be constructed via Diffserv or Intserv technology. There are two means to initiate a pipe request: via the CE action, or via the network management action.

The construction procedure can be described by the following three procedures: request for creating a pipe, re-sizing a pipe, and request for destroying a pipe.

The construction of a pipe can be summarized as follows:

1) The CE or the network management system sends a request for creating a pipe service via the management plane or the control plane to the ingress PE, including QoS parameters such as bandwidth, delay, packet loss probability, etc.
2) If it is CE-initiated operation, authentication and verification procedures need to be performed but are out of the scope of this Recommendation.
3) The management plane or the control plane of the ingress PE receives the request and performs the following tasks:
    a) Ensures that the end-to-end reservations can be made, via control plane or management plane operations. The creation of the pipe will result in failure if end-to-end reservations cannot be ensured.
    b) Performs various edge functions, including but not limited to, traffic classification, marking, and policer programming.

Similar to that of the pipe construction, the pipe deletion procedure can be described as follows:

1) The CE or the network management system sends a request for destroying a pipe to the management or the control plane of the ingress PE.
2) The ingress PE ensures that the end-to-end reservation for this particular pipe is not in effect for the given VPN.
3) The ingress PE removes the required configuration, such as traffic classification, marking, and policer programming.

The resize operation for an existing pipe is similar to that of the pipe creation (with minor modifications) and is not repeated in this Recommendation. If the resize operation results in a failure, the existing VPN pipe configuration should not be impacted.

9 VPN service based on the hose model

In this clause, the hose model used in VPN is investigated. The hose model is to ensure QoS on a point-to-cloud connectivity.

9.1 Definition of hose model

The VPN customer specifies QoS requirements per VPN endpoint and not per every pair of endpoints. Specifically, associated with each endpoint, there is a pair of bandwidth values – an ingress bandwidth and an egress bandwidth. The ingress bandwidth for an endpoint specifies the incoming traffic from all the other VPN endpoints into the endpoint, while the egress bandwidth is the amount of traffic the endpoint can send to the other VPN endpoints.

The hose model (see Figure 9-1) is easy to provision for the customer and no destination sensitive customer traffic shaping is required. The challenge for the service provider is to provision adequate resources at the traffic aggregation points inside the network.

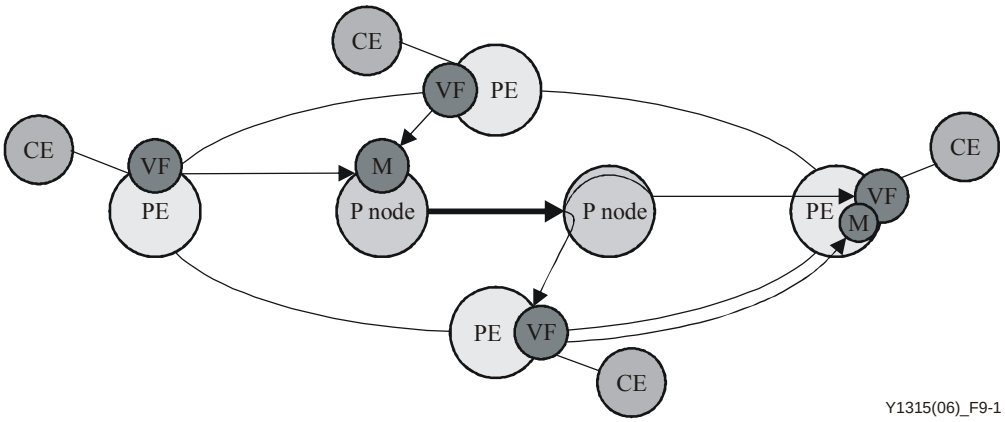

Figure 9-1 – Illustration of hose model

As in the pipe model, VF stands for virtual forwarder which contains VPN-specific forwarding information. Inside the network, on some P nodes or PEs, there can be a merge point for that VPN (denoted by M in the above illustration). The SLA enforcement and scheduling of the aggregated packet flows for that VPN need to be supported at the merging point.

The major difference between the pipe and hose model is that, in the hose model, traffic aggregation happens inside the network. In the pipe model, there is no traffic aggregation inside the network. Hence, in the pipe model, it is not possible to share the resources inside the network among different VPN endpoints.

A hose may be constructed via signalling if QoS assurance is to be ensured or via relative priority if over-provisioning is practiced. In order to construct the hose with QoS assurance, the merge point, which can be at a P node, is required to have policers so the traffic from one VPN does not have a negative impact on the network. The details of hose construction, maintenance, and deletion are out of the scope of this Recommendation.

9.2 Customer edge characteristics

In order to limit the traffic into the network within the bounds established via SLA, the CE needs to shape the traffic on the CE-to-PE link. Compared to that of the pipe model, the shaping for the hose model is not destination sensitive. The CE is then simpler as far as the traffic shaping is concerned.

9.3 Provider equipment characteristics

At the CE-to-PE link, the traffic policer on the PE is not per-destination based. Only the aggregated traffic needs to be policed based on the SLA.

If there are aggregation points inside the network, in order to ensure QoS and SLA conformance, proper resources need to be allocated.

There are a few issues related to the hose model:

- If a packet is discarded at the aggregation point inside the network based on the SLA (e.g., total VPN bandwidth is limited to a certain amount, regardless of the source of such traffic), such discards will be random and the application performance may be at risk of being poor (e.g., some applications may be more sensitive to packet loss).

- Due to the lack of per-flow reservation, and if there is no policer/shaper on the PE-to-CE link, the PE-to-CE link may be congested and hence result in packet loss. This case may be more pronounced if multiple CEs are sending traffic to a single CE. While each CE is sending traffic within its agreed-upon limit, the egress of the network, on the PE-to-CE link, will be overwhelmed.

If the network provider has to provide the hose model, the provider requires only one policer at the ingress of the network, as shown in Figure 9-2.

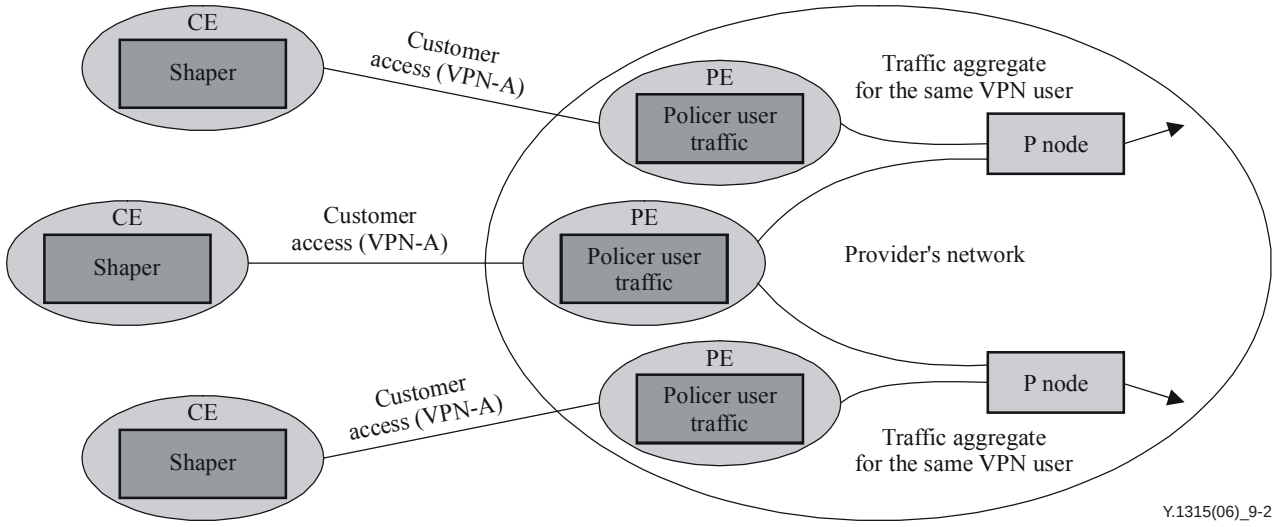

Figure 9-2 – Hose model ingress policing and network implications

The hose model is simple for the customer to configure the traffic, and simple for the network provider at the ingress node. But, in order to provide the required QoS, the provider needs to ensure that, at the traffic aggregation points, the customer traffic is properly handled.

9.4 Applicability, complexity, and scalability

From the PE's perspective, each VPN only requires one policer at the UNI. Hence the hose models are normally not limited by the number of policers at the UNI, as with the pipe model discussed in clause 8.4.

If the VPN bandwidth is to be monitored to ensure QoS and fairness, the aggregation (or merge) points inside the network (P nodes) for the hose model also require a policer. This dictates that the policer be installed not at the UNI. The policer on the P node also requires that packets carry VPN information, either implicitly or explicitly.

Due to its location not at the UNI, the policer may discard packets at that location as a result of packets received at different UNIs. Such inter-UNI interference is very undesirable since each UNI is acting independently and there is no CE action which can prevent such packet loss.

In terms of total policer usage, the hose model is less than or equal (when two flows from UNI are merged only once) to the pipe model.

In terms of network complexity, adding or deleting a VPN endpoint does not trigger network reconfiguration at the UNI, but may trigger the network-wide reconfiguration due to a possible merge-point change. The scope of network nodes affected for bandwidth reservation purposes is less than or equal to that of the pipe model, with the worst case being equal.

The scalability of the hose model is limited by the policers at the merge points (not at the UNI as that of the pipe model). The hose model may require less computing power when the bandwidth reservation needs to be changed (i.e., adding or removing an endpoint).

Due to the above rationale, the hose VPN is applicable if the number of VPN endpoints is large, and there is QoS assurance. Besides, packet loss does not result in a critical error because it is impossible for CEs to coordinate so traffic can be shaped properly at the merge points.

Even though it is possible for the hose model to support a very large number of VPN endpoints, the management of the merging points is difficult. In this sense, the hose model is best used when there is no QoS assurance in the SLA (CoS differentiation is acceptable).

10 VPN service based on both pipe and hose model

For any given provider's network, both pipe and hose model can be used to construct VPNs. A pipe model-only VPN may be constructed provided the conditions to construct the pipe model are satisfied (clause 8.4). The same statement can also apply to the hose model (with the hose model pre-requisites discussed in clause 9.4).

When both the hose and pipe models are offered, the service provider's equipment and management system needs to endure the complexity of both models.

If some segments of the VPN are constructed via the pipe model while other segments are constructed via the hose model, as shown in Figure 10-1, the general characteristics of the VPN is that of the hose model, excepted for the flows not going through the hose section of the VPN.

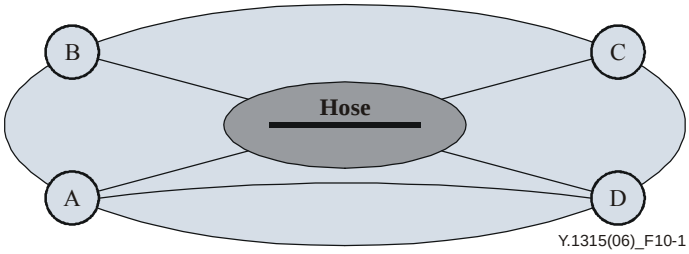

Figure 10-1 – Combining both pipe and hose model in the same VPN

The base VPN model can be a hose model, which reduces the complexity of the VPN management, especially if no QoS assurance is involved. For flows dictating QoS treatment, a pipe can be used between the flow endpoints.

In the case shown in Figure 10-1, a VPN will use the combination of both QoS models. For basic connectivity with CoS (class of service), the hose model can be deployed for simplicity (there will be no policing at the traffic merging points). For flows dictating QoS assurance, a pipe can be used. Hence the benefits of both the pipe model and the hose model can be materialized.

In particular, there is no need to have a large number of policers at the UNI points (not proportional to the number of endpoints of a given VPN), and there is no need to have policers at merge points. Of course, this is under the assumption that only a fraction of flows dictates QoS assurance.

11 Multiple service provider scenario

If a VPN service is provided via multiple VPN service providers, the overall service characteristics (i.e., service curve) can be obtained via the service curve of each provider (i.e., concatenation of service curves for each service provider).

One of the advantages of using the service curve to specify the service is its ability to work in the multiple service provider case. If each service provider's network is abstracted as an abstract node, offering service curve $\beta_i(t)$ and if the customer's packet is traversing N service providers, the effective service curve from the customer's point of view will be:

$$\beta(t) = \beta_1(t) \otimes \beta_2(t) \otimes ... \otimes \beta_N(t) \quad (1)$$

If a policer is employed between two service providers, the arrival curve for traffic delivered from service provider i to service provider $i+1$ will be:

$$\alpha^*_i(t) = \alpha(t) \otimes (\beta_1(t) \otimes ... \otimes \beta_i(t)) \quad (2)$$

There are many types of multiple service provider scenarios. In this clause, only the basic scenario with two service providers is illustrated. More complex configurations can be brought back to the basic configuration. It should be noted that there are multiple forms of inter-provider QoS arrangements and what is outlined in this clause is not meant to be exclusive.

The basic configuration of a multiple service provider scenario is the cascade of two service providers, as shown below in Figure 11-1.

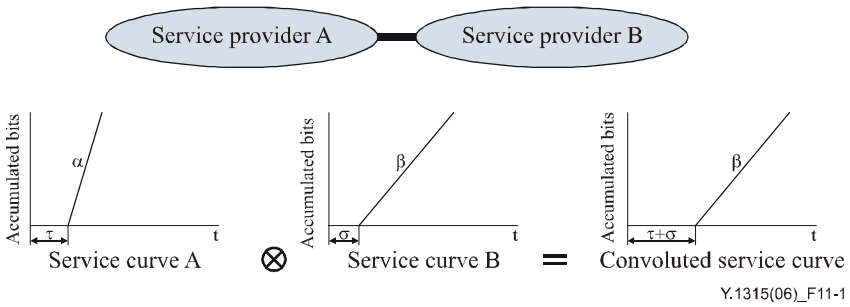

Figure 11-1 – Multiple service provider scenario

In this case, the slope introduced in the service curve will be the minimum slope of the two providers' service curves while the delay will be additive.

If a VPN service is provided by multiple service providers with a linear service curve (characterized by a delay and a slope), the overall service curve perceived by the VPN user will also be linear, with the slope to be the minimum of all these service curves, and the delay to be the sum of all delays introduced by these service curves.

Appendix I

Pipe model-based implementation example of VPN with specific QoS requirements

(This appendix does not form an integral part of this Recommendation)

I.1 Service requirements

Service providers have established their L3 VPNs based on [b-RFC 2547] and have started to provide L2 VPNs to their customers based on IETF L2 VPN WG drafts. Based on the operators' experience, a customer purchases L2 or L3 VPN services to connect his multiple sites at the corresponding layer. The service provider will ensure that traffic proper to a specific VPN does not trespass on other VPNs. Some VPNs may have topology constraints (e.g., hub-spoke).

Service providers have a practical need to build a unified network which is capable of providing L2 and L3 VPNs simultaneously. Some basic requirements of such a unified VPN network are as follows:

1) Support controllable route distribution – From the point of view of providers, VPN membership information (the sites belonging to the same VPN) and VPN reachability among sites should be controllable in flexible way and manageable in effective way.
2) Support multiple address types and overlapping address spaces (for different VPNs) – The "unified VPN" network should support both L2 and L3 VPNs, including L2 address types (e.g., FrameRelay (DLCI), ATM (VPI/VCI), and Ethernet (IEEE 802 MAC address and VLAN) address types) and L3 address types (IPv4 and IPv6).
3) Support end-to-end QoS – The MPLS technology has been widely accepted in the providers' backbone infrastructure, hence a practical strategy is to combine the MPLS QoS tools (e.g., TE, Diffserv) with the unified VPN architecture.

I.2 Unified QoS-VPN architecture

"QoS-VPN" implies a VPN with QoS requirements, including bandwidth, delay, jitter, packet loss rate among sites belonging to the VPN. The centralized QoS architecture listed in Appendix I of [ITU-T Y.1291] can be used to construct a unified QoS-VPN network. The QoS-VPN architecture can be divided into three layers: VPN service control layer, VPN transport layer, VPN transport control layer. In this architecture, MPLS is used in the transport network.

VPN transport control layer consists of VPN-CRCs, one for each VPN domain. The VPN-CRC (central router controller) corresponds to the BRM (bearer resource manager) layer depicted in [ITU-T Y.1291]. The typical VPN domain is an autonomous system in the current IP network. A VPN-CRC performs the following functions:

– Manages the network resources (including bandwidth, processor, and buffer) of a VPN-LBN (VPN-logical transport network);
– Maintains the VPN-LBN network topology;
– Performs path selection and then path information instruction to the ingress PE;
– Performs resource allocation and admission control within a VPN-LBN;
– Maintains membership and reachability information at a per QoS-VPN level;
– Performs related signalling to achieve membership auto-discovery and single-sided provisioning.

Diffserv-aware MPLS technology is used to process the different VPN flows belonging to the same PE-to-PE pipe.

The generic LSP label stack with three layers is used: TE tunnel outer label, VPN label, and flow label. The top label is the MPLS TE tunnel label with all the micro-flows inside which have the same QoS requirement. The VPN label, together with the flow label, identifies the source and the destination of the packet.

Figure I.1 shows an implementation example of the unified QoS-VPN architecture in one single AS (autonomous system).

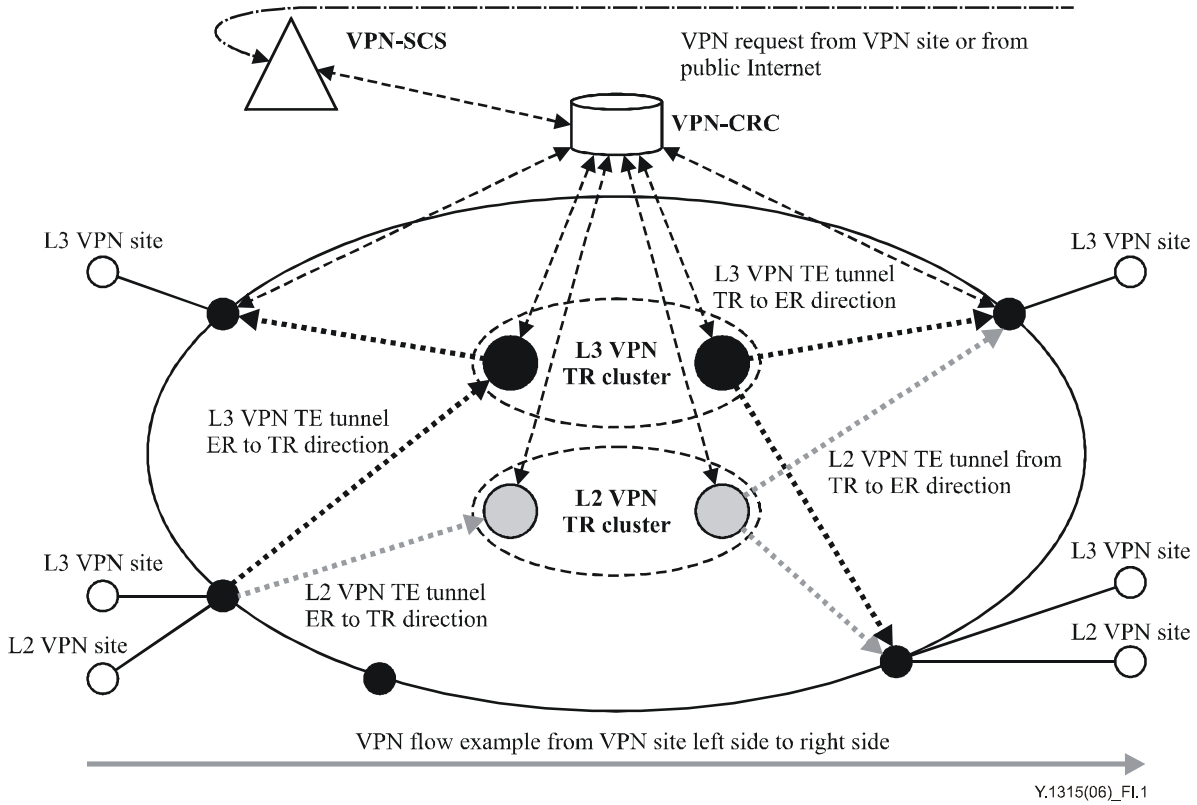

Figure I.1 – Implementation example of the unified QoS-VPN in one single AS

In order to improve the scalability of practical deployments, the concept of VPN central control cluster (CCC) system may be introduced (not indicated in Figure I.1). The main function of VPN central control cluster (CCC) is to achieve VPN resource isolation with the use of trunk TE tunnels between ingress trunk routers and egress trunk routers. These trunk ingress and egress routers form the trunk router cluster for a particular VPN service.

In the inter-AS scenario, VPN-LBN resources can be planned and managed according to the different AS. In this case, TE tunnels are initiated and terminated within the AS. In order to ensure the end-to-end QoS of a VPN service flow, there should be a mechanism to link the different TE tunnel segments.

I.3 Considerations

In this architecture, the PE supports upstream and downstream TE tunnel set-up through RSVP-TE/CR-LDP and traffic classification. The number of upstream and downstream TE tunnels is not of the order of n^2 where n is the number of PEs inside the network due to the traffic tunnel construction.

The key objective of this implementation example is to reduce the features required to be supported inside the PE and to improve the scalability and upgrade ability of the PE.

With this architecture, the PE can be manually configured and performs the functions of VPN flow identification, label binding, and VPN admission control. The network resource allocation and reservation is carried out by the logical central trunk clusters. This implementation example can provide both L3 VPN and L2 VPN services and with stringent QoS requirements. Further details on the protocols between PE to CRC, TR to CRC, and CRC to CRC are out of the scope of this Recommendation.

Appendix II

Service curve

(This appendix does not form an integral part of this Recommendation)

Consider the input function $R(t)$ (cumulative bits input into the network) and the output function $R*(t)$ (cumulative bits output from the network), the "service curve" $\beta(t)$ offered by the service provider will be defined as:

$$R*(t) \geq R(t) \otimes \beta(t) \tag{II-1}$$

where the symbol $\otimes$ denotes the min-plus convolution and can be defined as:

$$R(t) \otimes \beta(t) = \inf_{s \leq t} \{R(s) + \beta(t-s)\} \tag{II-2}$$

For a constant bit rate service, the service curve will then be defined as $\beta(t) = rt + t_0$ where t_0 denotes the initial delay in serving the packets.

For a flow, constrained by arrival curve $\alpha(t)$, traversing a system that offers a service curve of $\beta(t)$, the output flow will be constrained by the arrival curve $\alpha*(t)$, with $\alpha*(t)$ expressed as (min-plus deconvolution):

$$\alpha*(t) = \alpha(t) \otimes \beta(t) = \sup_{u \geq 0} (\alpha(t+u) - \beta(u)) \tag{II-3}$$

For a flow traversing N nodes (networks), while each node (network) offering a service curve $\beta_{n_i}(t)$, the service curve, after the nodal concatenation, will be:

$$\beta_{C_N}(t) = \beta_{n_1}(t) \otimes \beta_{n_2}(t) \otimes \ldots \otimes \beta_{n_N}(t) \tag{II-4}$$

The arrival curve after traversing those N nodes will be:

$$\alpha*(t) = \alpha(t) \otimes \beta_{C_N}(t) \tag{II-5}$$

Bibliography

LE BOUDEC (J.Y.), THIRAN Patrick, *"Network Calculus"*, Springer, 2001.

CHUANG (J.-F.), CHANG (C.-M.), *"Deterministic loss ratio quality of service guarantees for high speed networks"*, IEEE Communications Letters 4, pp. 236-238, 2000.

CRUZ (R.), CHANG (C.S.), LE BOUDEC (J.-Y.), and THIRAN (P.), *"A min-plus system theory for constrained traffic regulation and dynamic service guarantees"*, Technical Report SSC/1999/024, EPFL, 1999.

CRUZ (R.), and TANEJA (M.), *"An analysis of traffic clipping"*, In Proc. 1998 Conf. on Information Science & Systems, Princeton University, 1998.

CRUZ (R.L.), *"A calculus for network delay, Part i: Network elements in isolation"*, IEEE Trans. Inform. Theory, Vol. 27-I pp. 114-131, 1991.

CRUZ (R.L.), *"A calculus for network delay, part ii: Network analysis"*, IEEE Trans. Inform. Theory, Vol. 37-I, pp. 132-141, 1991.

[b-ITU-T X.85] ITU-T Recommendation X.85/Y.1321 (2001), *IP over SDH using LAPS*.

[b-IETF RFC 1661] IETF RFC 1661 (1994), *"The Point-to-Point Protocol (PPP)"*.

[b-IETF RFC 2205] IETF RFC 2205 (1997), *"Resource Reservation Protocol (RSVP)"*.

[b-IETF RFC 2215] IETF RFC 2215 (1997), *"General Characterization Parameters for Integrated Service Network Elements"*.

[b-IETF RFC 2547] IETF RFC 2547 (1999), *"BGP/MPLS VPN"*.

[b-IETF RFC 3198] IETF RFC 3198 (2001), *"Terminology for Policy-Based Management"*.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	**Global information infrastructure, Internet protocol aspects and next-generation networks**
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2007