QoS Provisioning Framework in IP-Based VPN

Mirjana D. Stojanovic
University of Belgrade, Serbia

Vladanka S. Acimovic-Raspopovic
University of Belgrade, Serbia

INTRODUCTION

A virtual private network (VPN) can be broadly defined as a "restricted communication between a set of sites, making use of a backbone that is shared with other traffic not belonging to that communication" (Carugi & De Clercq, 2004, p.116). Since the late nineties, with pervasive deployment of the Internet protocol (IP) technology in corporate networks, IP-based VPNs, in several forms and based on different network technologies, have become a promising solution for a wide range of corporate network services.

An IP VPN represents a "VPN implementation that uses public or shared IP network resources to emulate the characteristics of an IP-based private network" (European Institute for Research and Strategic Studies in Telecommunications (EURESCOM), 2002, p.14). IP VPNs complement the infrastructure of corporate wide-area networks (WANs) with solutions for heterogeneous requirements of virtual organizations including remote access (e.g., telecommuters, mobile users), intranet (interconnecting company offices and branches) and extranet (e.g., restricted access of business partners to a corporate WAN).

IP VPN approaches can be classified according to management responsibility to customer-provisioned VPNs and provider-provisioned VPNs (PPVPNs). Depending on the location of the specific VPN equipment, PPVPNs can further be classified to CE (customer edge) based VPNs and PE (provider edge) based VPNs (also called "network-based VPNs"). PE-based IP VPNs are distinguished depending on the offered service as: (1) PE-based L2 VPNs, which offer layer 2 services, and (2) PE-based L3 VPNs, which offer layer 3 services.

Vendor-specific IP VPN solutions at the end of nineties preceded the standardized ones. ITU-T Study Group 13 started with standardization activities for IP-based VPNs in 2000, while the PPVPN working group has been founded within the Internet Engineering Task Force (IETF), in 2001. Comprehensive overviews of standardization efforts and solutions for IP VPNs have been published recently (Carugi & De Clercq, 2004; Knight & Lewis, 2004).

Solutions for IP VPN quality of service (QoS) have been studied since the late nineties. Braun, Guenter, and Khalil (2001) identified the following three components as a foundation for QoS support in IP VPNs: (1) differentiated services—DiffServ (Blake, Black, Carlson et al., 1998); (2) traffic engineering (Awduche, Chiu, Elwalidet al., 2002), and (3) multi-protocol label switching—MPLS (Rosen, Viswanathan, & Callon, 2001). They also proposed management architecture for QoS-enabled VPNs, including business, service and network management models.

EURESCOM (2002) systematized generic QoS management capabilities needed for provisioning different types of IP VPNs. Zeng and Ansari (2003) addressed QoS issues for PE-based L3 VPNs from the provider's perspective, including IP QoS architectures and their associated mechanisms, as well as management aspects. Szigeti and Hattingh (2004) provided a practical approach for designing end-to-end QoS considering MPLS VPNs and IPsec (IP security) VPNs.

In this chapter, we address QoS aspects in PE-based L3 VPNs, and propose a framework for end-to-end QoS provisioning in a DiffServ-aware provider's network. The framework specifies VPN service level agreements, QoS mechanisms, adaptation of routing through traffic engineering and the QoS management system.

REFERENCE MODEL OF PE-BASED L3 IP VPNs

In PE-based L3 IP VPNs, all network routers are capable of forwarding VPN traffic to appropriate destinations, that is the customer and provider need to cooperate at the routing layer. Reference model of PE-based L3 IP VPNs is presented in Figure 1 (Callon & Suzuki, 2005).

Figure 1. A reference model of PE-based L3 VPNs (Adapted from Callon & Suzuki, 2005)

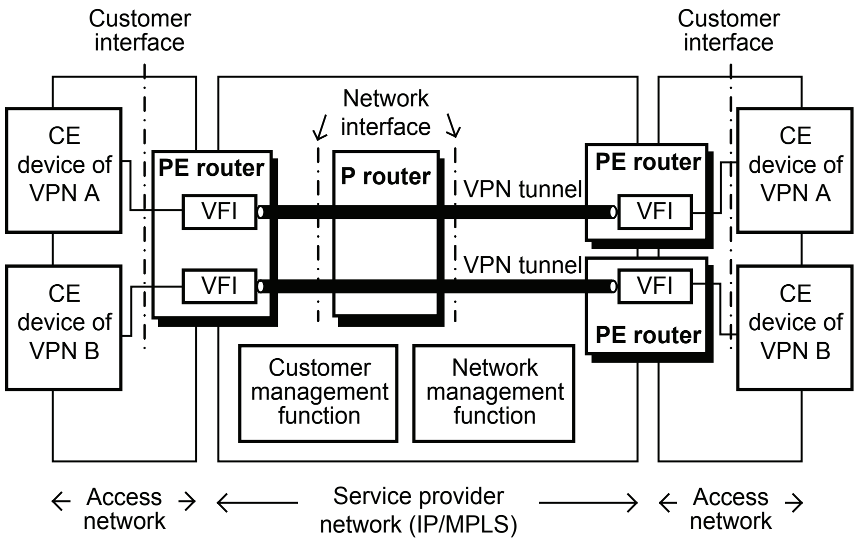

The backbone of provider's network consists of a single or multiple cooperating domains, relying on the IP/MPLS technology. A set of P (P—provider core) routers corresponds to core IP/MPLS routers, which are transparent to VPN services. A set of PE (PE—provider edge) routers corresponds to network edge routers and implements VPN functionality by means of a set of VPN forwarding instances (VFIs). VFI entities implemented in different PE routers of the single provider's network are mutually interconnected by means of VPN tunnels. *VPN tunnel* is a logical connection which is created by encapsulating each IP packet at the ingress PE router, that is by adding a proper header in such way that the encapsulated traffic is forwarded through network based on the destination address, and optionally other fields in that header. Depending on the VPN architecture, the VFI concept may be implemented as a virtual router or VPN routing and forwarding (Andersson & Madsen, 2005; Carugi & De Clercq, 2004).

A customer edge (CE) device corresponds to a host, router or switch at the customer's site and is connected to one of the provider's PE routers by an access connection realized through the access network. Each access connection is associated with the appropriate configuration of the VFI entity in PE router. In general, CE device can be simultaneously connected to a number of PE routers or a number of CE devices belonging to the same VPN may be connected to the same VFI entity in PE router.

Customer management function encompasses provisioning of the customer-specific attributes like identifiers, personal information, subscribed services, access control policy, accounting and billing information, statistical data, and so on. Network management function is responsible for monitoring and provisioning the attributes of PE and CE devices and their mutual relationships.

A FRAMEWORK FOR QOS PROVISIONING IN PE-BASED L3 VPNs

Considering network scalability requirements, our proposed framework for QoS provisioning relies on standardized (Blake et al., 1998) or proprietary architectures with differentiated services, that is DiffServ-aware architectures. They have a common property that incoming packets belonging to different traffic flows, but with similar QoS requirements, may be associated to the same traffic class and processed in the same manner at network nodes. Complex processing operations, including packet classification and traffic conditioning, are performed at the edge routers, while core routers perform simple forwarding operations based on the traffic class code in packet header.

Providing end-to-end IP QoS requires extension of the basic DiffServ model with a variety of mechanisms dealing with the call and packet handling, as well as