Design of a Charging and Accounting Architecture for QoS-differentiated VPN Services to Mobile Users*

Thanasis G. Papaioannou and George D. Stamoulis

Department of Informatics, Athens University of Economics and Business (AUEB),
76 Patision Str., 10434 Athens, Greece.
{pathan, gstamoul}@aueb.gr
telephone: +30-210-8203549, telefax: +30-210-8203686

Abstract. In the emerging context of mobile Internet, the importance of VPN services is rapidly increasing. Provision of such services was among the subjects of IST project INTERNODE. Besides the necessary technical means, charging and accounting also are key related issues, and constitute the subject of this paper. Only by dealing successfully with charging and accounting, VPN providers can recover their provision costs, increase their profits, and provide the right incentives to their users, thus leading to efficient operation of their network. In this paper, we first study the chargeable characteristics of QoS-differentiated VPN services offered to mobile users with respect to transport, security and mobility (both personal and terminal). Then, we define a complete charging scheme that is fair for the users and provides them with the incentives to use only the resources they really need. This scheme is based on the time-volume charging approach by Kelly; the adoption of this approach is justified in detail. We then show how the providers involved can share the total charge earned by each VPN service instance in a fair way, with each provider collecting the portion of charge that corresponds to the consumption of his own resources for the service. This is also a very important issue for the commercial viability of VPN services to mobile users, given that its provision spans multiple domains. Our approach also includes computation of an estimate of users' expected charge prior to using the VPN service. Finally, we specify an appropriate charging and accounting architecture pertaining to the specified charging scheme for VPNs, to the mechanism for revenue sharing, and to the technical implementation of the VPN services studied. This architecture is compliant to the relevant standards, is applicable to the current and the future Internet, was fully implemented, and can serve as a basis for applying other charging schemes as well. Our work can also serve as a methodology for designing charging and accounting architectures for a variety of Internet services.

Keywords: charging, accounting architecture, VPN, mobility and security, QoS-differentiated services

* The present work has been carried out as a part of the IST project INTERNODE (IST-1999-20117) funded by the European Union, through a subcontract with INTRACOM S.A.

A previous version of this work was presented in the workshop on *Internet Charging and QoS Technologies*, Zurich, Switzerland, October 16-17, 2002.

1 Introduction

The globalization of commerce as well as the ease in human transportation has increased the mobility of professionals and tourists. Also, in recent years the use of mobile phones has grown tremendously. The increased terminal capabilities as well as the on-going development of SMS and WAP applications have started bringing mobile Internet into reality. In this new networking environment, mobile users need to retain seamless and secure connectivity while being in a visited domain, as if being at home. Also, mobile users should be able to form private working groups independently from their respective point of attachment to the Internet. These requirements are fulfilled by the provision of VPN services to mobile users. Furthermore, VPN services should be customized in order to satisfy certain user preferences regarding levels of security and quality of service (QoS). We use the term QoS-differentiated in order to imply that there are different possible QoS levels for traffic transport together with the possibility of Best Effort. On the other hand, VPN providers should account and charge for their services, in order to recover their provision costs, increase their profits (while being competitive), and provide the right incentives to their users, thus leading to an efficient operating mode of their network. The users should be provided with the right incentives to use as many network resources, as they really need. At the same time, users should be charged in a fair way, in the sense that each of them should be charged appropriately for the usage of each of the network resources he actually uses. Moreover, the charge assigned to each VPN service instance should be shared among the providers involved in an efficient and fair way.

In this paper, we discuss and fully specify a complete, yet lightweight and thus scalable, charging and accounting architecture for QoS-differentiated VPN services offered to mobile users. These services will be referred to as M-VPN services. Our work is applicable to both cases of personal and terminal mobility. We analyze the chargeable characteristics of M-VPN services, and justify the adoption of a proper charging scheme that satisfies the requirements discussed above, and simplifies the fair sharing of the revenue resulting in a M-VPN service instance among the various providers involved. As already mentioned, the charging scheme adopted provides users with the right incentives. This helps providers to: i) set competitive tariffs that match user needs, and thus ii) maintain their position in a competitive market. Finally, we compare our work to other related articles and clarify our contribution. This includes both the specific charging and accounting architecture presented in this paper, as well as the methodology for the design of this architecture, which can be employed in cases of other services in the current and future Internet (i.e., when IPv6 will be in place). This work was part of the IST project INTERNODE [1], which studied the provision of QoS-differentiated VPN services to mobile users.

2 The INTERNODE Approach for M-VPN Service Provision

In this section, we describe the context of IST project INTERNODE for providing QoS-differentiated VPN services to mobile users (M-VPN services). INTERNODE designed, specified and implemented a platform, which can be used to create multi-domain VPN services for mobile users; e.g. users of e-business applications, of mobile Internet and intranet access, of personalized services, etc. Specifically, VPN connectivity is handled by a VPN provider that employs a VPN Service Provisioning and Support Platform (SPS). This platform is capable, among others, of automatically configuring and managing the VPNs on behalf of the VPN subscribers.

A customer subscribes to the VPN SPS and registers a number of mobile users to use the M-VPN services in a way specified in a contract. Specifically, the customer establishes a M-VPN contract with the VPN provider where security, QoS, and other characteristics of the M-VPN services for each registered user are defined. Through the SPS platform, the VPN provider performs VPN access control, charging and accounting, and management of the VPN termination points, so as to provide the security and the QoS levels declared in the M-VPN contract. Security is provided by means of the IPsec protocol [2] on a gateway-to-gateway basis. QoS is provided using the Differentiated Services (DiffServ) architecture [3], by specifying appropriate flow-level Service Level Agreements (SLAs). Implicitly assumed is the existence of inter-domain DiffServ SLAs among Connectivity Providers (CPs) that are established through Bandwidth Brokers [4] and maintained by certain mechanisms as those in [5]. Note that QoS differentiation within the IPsec tunnel can be accomplished if the Differentiated Services (DS) field of each internal packet is copied or mapped to the DS field of the external packet, as in the case of IP-within-IP tunnelling for Mobile IP [6]. [Note also that according to [2] and [6], copying or mapping of the Type of Service (TOS) field to the external packet header during tunnelling is mandatory. In most cases, this approach provides enough bits for class differentiation in a DiffServ domain.]

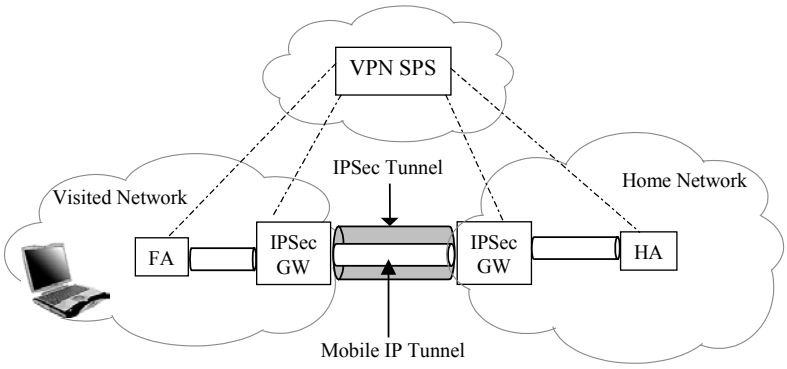

Figure 1. VPN provision to a mobile user from a visited domain to his home domain through a Foreign Agent in an IPv4 environment

As depicted in Figure 1, the VPN provider fulfils the above tasks through federated APIs with the edge CPs, i.e. the potentially visited and the home CPs. Note that all

edge CPs have to apply the security and mobility support technologies adopted by the VPN provider. Therefore, we assume that according to such federation agreements, the VPN provider leases the following equipment to each of the aforementioned CPs: i) an enhanced router that supports IPsec, which is referred to as Security Gateway (SG), and ii) an enhanced router supporting mobility, which is referred to as Home Agent (HA) [resp. Foreign Agent (FA)] if the CP serves as home (resp. visited) network. SG and HA/FA are used as mediation devices by the federated CPs, in order to associate resource usage to particular users (and thus to customers), support the mobility and security features of the M-VPN services, and monitor the conformance with the M-VPN contracts that the VPN provider has established with its customers. (Note that if DHCP is employed, then it is not necessary to employ a FA in the visited network.) The traffic generated by a mobile user is forwarded to his home domain and from there to its final destination, as specified by reverse tunneling for Mobile IP [7]. For a secure duplex end-to-end communication, the creation of two IPsec tunnels is necessary, i.e. one for each direction of the traffic between the sender and the receiver. Note that, in INTERNODE, the approach of a "trusted FA" was used, which, according to [8], does not suffer from any technical problems. Note, also, that the VPN tunnel endpoints change only in case of handover between administrative domains; they do not change in case of handover between cells (that belong to the same administrative domain).

Furthermore, due to the use of the VPN SPS platform, this approach is applicable even for the scale of Internet with either Mobile IPv4 [9] or Mobile IPv6 [10]. The only difference in the architecture applicable to the future Internet, with Mobile IPv6 in use, is that the mobile node registers itself to the HA, because there is no need for an FA in the visited network. (Recall also that the IPsec protocol supports both IPv4 and IPv6.) Next, we discuss the case of IPv4 and IPv6 coexistence, when IPv6 transition is required and configured/automatic tunneling of IPv6 traffic over IPv4 is employed. In this case, the IP-in-IP tunnel endpoints must co-exist with or appear after the IPsec tunnel endpoints, in order to attain end-to-end connectivity; see Figure 2. Otherwise, if IPv6 transition has to be employed in certain segments of the VPN tunnel, then the dual nodes that apply IPv6 to IPv4 translation for each segment should also be trusted security gateways (SGs) that decapsulate and re-encapsulate the packets according to IPsec across each segment. Finally, in case where one of the CPs employs IPv6 while the other employs IPv4, then our approach is still applicable: one of the edge CPs should employ a dual node, which would perform the transition appropriately outside the IPsec tunnel.

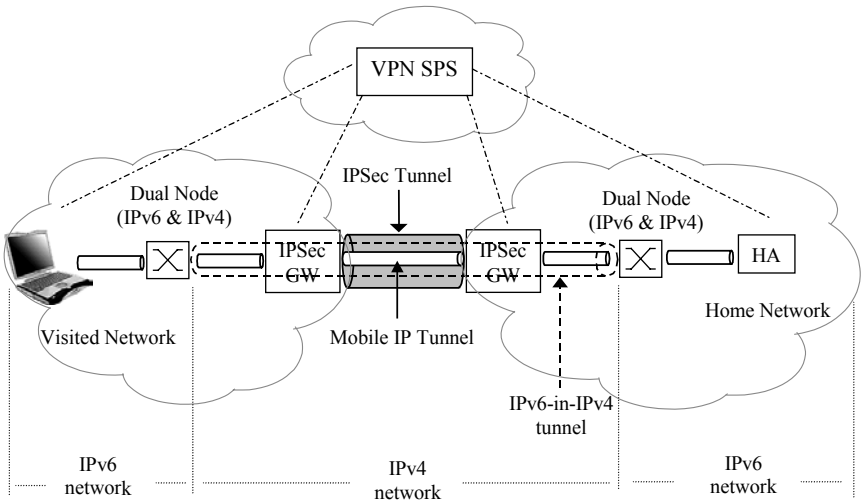

Figure 2. VPN provision to a mobile user from a visited domain to his home domain in case that IPv6 is employed by the edge CPs and IPv6 to IPv4 transition is required between them

3 Analysis of Business Roles

In this section, we define the business roles related to providing QoS-differentiated VPN services to mobile users (M-VPN services) in the way described in Section 2. We present this part prior to the charging scheme, since the structure of the latter (i.e. which role accounts and charges for what) is influenced by the business roles involved and by the interactions among them. As discussed in Section 2, the VPN provider employs a VPN Service Provisioning and Support Platform (SPS) capable (among others) of automatically configuring and managing the VPNs on behalf of its customers. The VPN tunnel may traverse through any number of intermediate CPs between the home and the visited networks, by which it is treated as specified by the corresponding inter-domain DiffServ SLAs. Moreover, each edge (home or visited) CP has the obligation, due to federation agreements with the VPN provider, to support the M-VPN services provision in terms of QoS, mobility, security and accounting/charging. Thus, the federated CPs serve as *third-party providers*, each offering a certain part of the M-VPN service instance. On the other hand, recall from Section 2 that the VPN provider is responsible for management and "packing" of the M-VPN services. Each federated CP has to send the information on the charge that arises from its participation in the service provision to the charging and accounting subsystem of the VPN SPS. Specifically, the charging and accounting subsystem of each third-party CP, in accordance with his federated agreement with the VPN SPS, has to: i) measure the usage of resources in the course of M-VPN service instances, and associate it appropriately with its users, in order to allocate individual charges, and ii) send this information to the charging and accounting subsystem of the VPN

SPS. It is a responsibility of the federated CPs to charge (or set the charging tariffs) for the resource usage in their network domain, according to the M-VPN contract. On the other hand, the VPN SPS calculates the total charge for providing the M-VPN services.

4 Specification of the Charging Scheme

In this section, we present our approach for charging M-VPN services. In particular, we adopt *additive* charging over all flows arising within a VPN; this approach is appropriate for charging VPN services, as explained in [11]. The charging scheme for individual flows should reflect the actual usage of network resources. The main issues concerning charging of individual flows within an IPsec tunnel of a VPN are analyzed below. Initially, we assume that no handover (from a visited network to another) occurs throughout the duration of the M-VPN service session. This issue is dealt with at the end of this section.

4.1 Charging for the transport of traffic

As already mentioned, the DiffServ architecture is used for QoS provision to the individual IP flows and/or the aggregates of traffic traversing the IPsec tunnels. Two charging schemes are applicable in this context: i) the scheme proposed by Courcoubetis and Siris [12] for charging DiffServ SLAs, and ii) the *time-volume* approach (also referred to as "*a*, *b*, *c* approach") proposed by Kelly [13]. The latter applies to services with quality guarantees, including ATM VBR and DiffServ; the charge equals $aT + bV + c$, where T is the time duration of a flow, V is the corresponding volume of traffic, and a, b, c are the tariff parameters. These parameters are derived from the *effective bandwidth* curve, and will be henceforth assumed to be expressed in appropriate units so that $aT + bV + c$ corresponds to a monetary amount. Both charging schemes use the on-off bound as a proxy for resource usage, which is the effective bandwidth of an on-off source with certain mean and peak rates, and serves as an upper bound to the effective bandwidth of every source with the same mean and peak rate. Also, both of the charging schemes above are applicable to paths consisting of a single link only, as well as to longer paths. They can be applied for such longer paths either additively or by focusing for charging purposes on the bottleneck link (which accounts for the largest portion of the total transport charge incurred under the additive approach) and ignoring all other links [12].

We have adopted the time-volume approach for charging IPsec flows served by DiffServ QoS classes, because this approach benefits considerably from a priori information on traffic properties. Moreover, we have adopted, for simplicity, the bottleneck link approach for the case of paths, and we assume that the bottleneck link is part of the IPsec tunnel, rather than being one of the links used within the edge networks. Indeed, information of the traffic properties can be made available for an IPsec flow through the type of user that generates this flow. In particular, we assume

that the user identity is part of the M-VPN contracts and it is indicative of the expected network usage: all users corresponding to the same identity or to a certain group of identities are taken to be of the same type, e.g. administrative employees or technical employees etc. A certain type of users corresponds to a certain type of traffic source; for example, administrative employees usually use videoconference applications, whose traffic volume and time duration can be monitored, thus leading to statistical information that can be used in order to optimize the selection of tariff parameters (see below).

As the time-volume approach is used for the computation of the charge for the traffic of an IPsec flow, it is necessary to specify the way that the proper *a*, *b*, *c* tariff is selected for charging this flow. As already mentioned, the identity of the user sending/receiving traffic over an IPsec flow determines the type of the user, and consequently the traffic source; i.e., the group of applications that the user may use together with statistics of the associated traffic. We assume that for each different type of user and for each different application there are predefined leaky bucket parameters and an estimate *m* of the mean rate value for each QoS class permissible to serve this application. Note that a M-VPN service instance is bi-directional, and thus it involves the creation of two IPsec flows. We take as mean rate of an application the mean rate of the IPsec flow that conveys the content of the application (e.g. the flow delivering video-frames), rather than of the flow conveying control signals. If both sending and receiving IPsec flows of an application convey useful content for a user type (e.g. the two flows involved in video-conference), then an estimate of the mean rate is kept for each flow. On the other hand, the M-VPN contract determines the alternative DiffServ QoS classes permissible to serve the flow of a particular application for a specific type of user. Thus, the pair <user identity, M-VPN contract> determines the eligible optimized tariffs for each IPsec flow with one such tariff being offered for each permissible QoS class. The final tariff depends on the QoS class to be actually selected by the user. As explained in [13], the pair of tariff parameters *a*, *b* that *minimizes the expected charge* for a particular QoS class (under the assumption that no handover will occur throughout the session) can be selected on the basis of the available estimate of the mean rate *m* of the application served by the IPsec flow. (The various CPs have the incentive to try to minimize the charge for their users in order to be competitive.) The optimal values of the parameters *a*, *b* also depend on the values of the leaky bucket parameters associated with the application. The tariff parameter *c* will be considered as fixed for all flows.

Recall that, initially, the charging module has an estimation of the mean rate for each application and for each type of users. The mean rate *m* for each application for a certain type of users from each particular point of attachment is constantly monitored and its "future" value is predicted, e.g. as a weighted average of past measurements. (The weights may be larger for more recent measurements.) Note that estimating the mean rate per application for each user type induces a storage and monitoring overhead. Another alternative with less overhead would be to monitor the aggregate mean rate per user type. This however may result in considerable inaccuracy in the prediction of the actual mean rate of an IPsec flow, while all such flows will have to be charged with the same *a*, *b*, *c* parameters. According to [13], this will lead to a

higher total expected charge for the entire M-VPN service instance. Thus, there is a *trade-off* between accuracy of the prediction of the mean rate and the monitoring overhead, which depends on the level of aggregation of the statistical information measured. Finally, if an IPsec flow is served Best Effort, then the time-volume formula is still applicable for computing the charge employing either the total volume or some other volume measure expressing the *burstiness* of the flow [14].

4.2 Charging for security

Each federated provider that its network domain serves as a source or sink of the M-VPN services provision is equipped with the VPN Service Provisioning and Support (SPS) platform with a Security Gateway (SG). There is a computational overhead for SGs in the establishment of a new security association (i.e., a new IPsec tunnel). Clearly, aimless or malicious creation of IPsec tunnels can be prevented by adding a fixed charge per IPsec tunnel creation or alternatively a (smaller) fixed charge per individual IPsec flow insertion to the IPsec tunnel. Note that in order for the right user incentives to be maintained, different fixed charges should be assigned to different security levels that are offered by the M-VPN services to account for the different overhead associated per such level.

Next, we discuss charging of the consumption of *computational resources* in the security procedure. There is a constraint in the service rate of the SG that encapsulates/ decapsulates packets according to IPsec, which is similar to the capacity constraint of a communications link. This constraint should be satisfied only when the traffic of the IPsec tunnels is served by DiffServ QoS classes, since in this case the traffic inserted is policed (i.e. already constrained) according to the M-VPN contracts of the customers. We make the safe assumption that the processing capacity and the buffer of the encapsulation process are larger than the capacity and the buffer of the bottleneck link. Thus, consumption of computational resources for such traffic should *not* be charged. However, in case of Best Effort traffic, an extra volume- or burstiness-based charge should be introduced so as to enforce the aforementioned capacity constraint.

Additionally, there should be a constant charge per time unit for each IPsec flow, because the identity comprising Security Parameter Index (SPI), IPsec protocol (Authentication Header or Encapsulating Security Payload), and IP destination address of an IPsec tunnel is a “scarce” resource, since there can only be finitely many IPsec tunnels between two security gateways. Thus, aimless maintenance of IPsec tunnels is prevented through charging. Finally, note that IPsec tunneling results in an increase of the traffic volume and the time duration (due to the induced computational overhead) for a M-VPN service instance, which are “automatically” included in the computation of the transport charge. Indeed, this computation employs the volume and time arising and measured after encryption.

4.3 Charging for mobility

We first consider the charging issues arising in the provision of M-VPN services in the context of terminal mobility. As explained in Section 2, the VPN SPS provides each federated CP involved with a Mobility Agent (MA, i.e., Home or Foreign Agent). As reverse tunneling for Mobile IP is used, a Mobile IP tunnel is created for each direction between the FA and the HA. There is a constraint in the service rate of the MA that encapsulates/decapsulates the packets according to Mobile IP, which is again similar to the capacity constraint of a communications link. This constraint is satisfied in case of statistically guaranteed services and no charge is required, according to the discussion above for computing the charge for security. (Again, we make the safe assumption that the processing capacity and the buffer of the encapsulation process are larger than the capacity and the buffer of the bottleneck link.) Also, in case of elastic services, an extra volume- or burstiness-based charge should be introduced.

Furthermore, each network domain has a finite set of IP addresses that a potential visited user can use interchangeably. The visitors should have the incentive to de-allocate their care-of IP addresses when not really needed. Charging for this "scarce" resource is accomplished by a fixed price per time unit for the usage time of a care-of IP address. (Note that this argument applies only to the case of Mobile IPv4, for which a care-of address is needed.) In case of personal mobility, the user leases a terminal for accessing the M-VPN services. This terminal can also be a "scarce" resource, as the number of the available terminals can be limited. Thus, leasing of this terminal should be charged by a fixed price per time unit throughout the leasing period. Note that Mobile IP tunneling also results in increased volume and time (due to the computational overhead) for a M-VPN service instance, which are included in the computation of the transport charge.

4.4 Computation of the total charge

Having dealt with all key issues on charging M-VPN services, we are now in a position to proceed with the computation of the total charge and the sharing of revenue among the players involved. As already explained, the total charge is the sum of all contributions defined in Subsections 4.1, 4.2 and 4.3. This summation spans all users of a particular M-VPN service instance that belong to the same customer. In particular, a user of a certain user type j using an application i that is served by a Diffserv QoS class q during a M-VPN service instance should be charged according to the formula below:

$$\left(p + a_{ijq}(m_{ijq})\right) \cdot T + b_{ijq}(m_{ijq}) \cdot V + c_s \qquad \textbf{(1)}$$

where T is the duration of the application and V is the corresponding volume transferred within the IPsec tunnel. p is the sum of prices for mobility and security support per time unit. $a_{ijq}(m_{ijq})$, $b_{ijq}(m_{ijq})$ are the parameters a, b of the transport tariff for the user type j with estimated mean rate m_{ijq} for the application i served by the QoS class q, which are derived from the M-VPN contract. Last, c_s represents the sum

of fixed charges (for transport, security and mobility support) associated to a new M-VPN service instance, and only depends on the security level *s*. In case the above application *i* of a user with user type *j* were served Best Effort, formula (1) should be modified according to the relevant charging schemes used for traffic transport, security and mobility; see above. A discount can be included in certain of the terms in formula (1) prior to the computation of the total charge on the basis of the identity of the customer and/or QoS/contract violations (see [15]). Last, note that the tariffs are set by the CPs in such values as to recover the costs resulting from their inter-domain SLAs for the transport of traffic.

Although the charging scheme may seem complicated, it is in fact simple: the user is offered a tariff consisting of a charge per time unit, a charge per volume unit, and a fixed charge for a new VPN connection of a certain security level. The details on how these charging parameters are derived can be presented to the user upon request. Furthermore, based on these tariffs and the estimated mean rate for the specific application that a user will use, our proposed charging mechanism is able to provide him before the start of his M-VPN session with an estimation of his expected charge during this session, assuming that no handover will occur. On the other hand, the *a*, *b* charging parameters may change during the M-VPN session in case of handover between administrative domains or in case of considerable modification of the congestion level encountered by the session (e.g. a handover between GPRS and WLAN cells). In that case, the user is notified about his new charging tariff and his corresponding expected remaining and total charges.

4.5 Sharing of Revenue

As already explained, the charge is the result of the addition of various contributions, each of which corresponds to the consumption of resources owned by a player with a certain business role. A straightforward way to share revenue among these players is for each of them to collect the revenues corresponding to the consumption of his own resources; e.g. an edge CP contributing to a M-VPN service instance would collect the revenues resulting from the transport, mobility and security support offered by his own network resources. Business agreements however can enforce different methods of revenue sharing, since different players may have different market power.

5 The Charging and Accounting Architecture

5.1 Accounting Issues

In order for the charging scheme previously described to be properly applied, accounting has to measure the traffic of each IPsec flow and associate it with a specific type of users belonging to a customer. Thus, we employ accounting per user, but charging per customer. In order to accomplish these, accounting has to:

1. Discover the identity of the user and thus specify the type of the user as well as the identity of the customer that the user belongs to. The user identity is determined by his home IP address. If the home IP address is private [16], then, according to [7], the user is uniquely identified by the pair <home IP address, Home Agent IP address>. Some link layer information is also required for user identification, in case the mobile user is registered to multiple Home Agents (HAs).
2. Measure the traffic (volume and/or burstiness measure) inserted in an IPsec tunnel by each user. This accounting procedure is performed at the ingress of the tunnel. The Security Gateway (SG) performs this task, using header information of packets at both the input and output interfaces, i.e. before and after IPsec encapsulation.
3. Be able to provide online feedback to the customer on his current charge and tariffs; this is referred to as *live accounting*. This capability may include support for debit payments and/or pre-paid services, or warnings concerning the charge accumulated.
4. Support the capability of providing the customer with one total bill for the service (i.e. opaque billing) from the VPN provider or with separate bills from each contributing provider to the service provision (i.e. transparent billing).

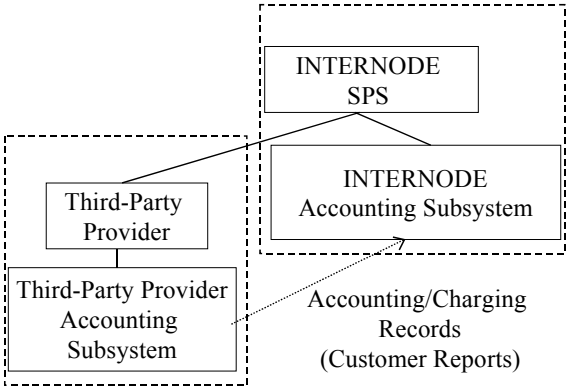

Figure 3. The overall Accounting Information Model between a CP and the VPN provider

The overall model of communication between the charging and accounting subsystem of the VPN provider and that of a third-party provider is based on the principles of the AAA architecture [17]. Thus, all charging and accounting records are generated by the charging and accounting subsystems of the third-party providers and forwarded to that of the VPN provider also in the form of Session Data Records (SDRs); see Figure 3. Recall that, according to AAA [17], a SDR contains a summary of the resources consumed by a user over an entire session. In the INTERNODE context, a SDR conveys the resource consumption of a registered user during a M-VPN session according to the M-VPN contract. Finally, both the way the charging and accounting

records are generated and the information contained therein depend on the contractual agreements between the federated CP and the VPN SPS, according to the TeleManagement Forum [15].

5.2 The Building Blocks of the Architecture

We now present the charging and accounting architecture; this is depicted in Figure 4, in a TINA hierarchy, for clarity reasons. The architecture consists of the charging and accounting subsystems of the VPN provider and those of the contributing third-party connectivity providers (CPs). Below, we describe the functionality of the building blocks of a charging and accounting subsystem. Note that we have fully implemented this architecture with all functionalities described in this section in the course of project INTERNODE. This software was integrated in the project's prototype platform for M-VPN service provision; see [1].

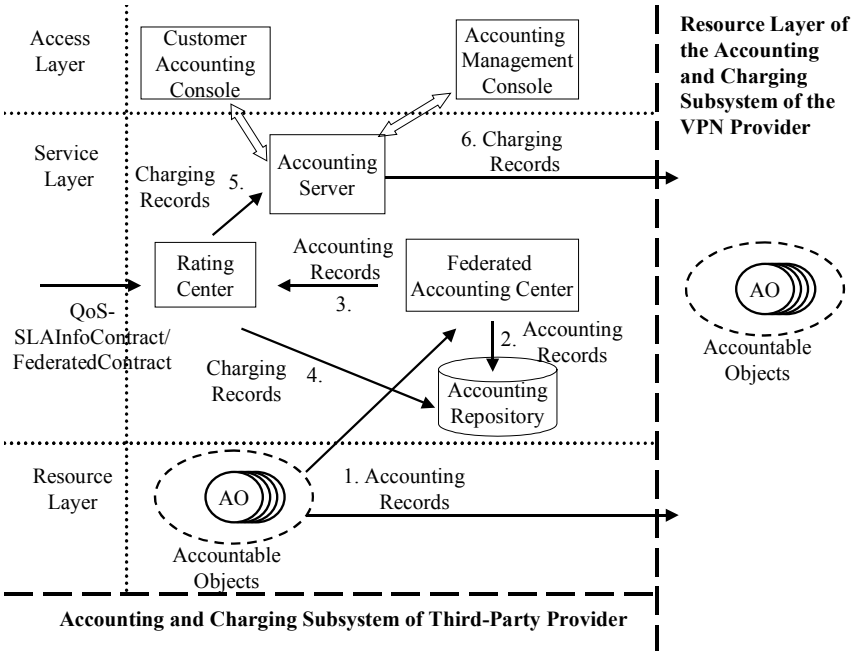

Figure 4. A high-level view of the building blocks of the accounting architecture. The charging and accounting subsystems of the VPN provider and the third-party Connectivity Providers (CPs) have the same structure

Accountable Objects (AOs): For each creation of a new instance of QoS-differentiated VPN services to mobile users (M-VPN), AOs associated with the corresponding Mobility Agents (MAs) and the Security Gateways (SGs) for this bi-directional VPN tunnel are activated, and an AO associated with the M-VPN service instance itself is created. Note that the AOs represent the hardware components used for M-VPN

service provision, and thus they are placed in TINA Resource Layer of the architecture; see Figure 4. The AOs associated with MAs and SGs belong to the charging and accounting subsystems of the third-party providers contributing to the M-VPN service instance. They collect the relevant accounting information and forward it both to the AO associated with the M-VPN service instance (belonging to the VPN provider) and to the Federated Accounting Centers (see below) of their respective charging and accounting subsystems, as depicted by arrow 1 in Figure 4. Thus, the AOs form a ladder (i.e. a hierarchy in the information flow, according to TINA accounting information model [18]) of accounting information for each M-VPN service instance, which is used for charging and accounting *auditing* by the VPN provider. This accounting ladder simplifies management of AOs and the aggregation of accounting events associated to a particular service instance. The AO associated to the entire M-VPN service instance forwards the accounting events to the Federated Accounting Center of the VPN provider. Note that the AOs remain intact for cell handovers within the same visited administrative domain. In case of handover between administrative domains, a new VPN tunnel and the AOs of its associated accounting ladder are created, while the old VPN tunnel remains alive for a timeout period in order to support handling of border movement across different administrative domains. Thus, the accounting functionality continuously measures chargeable events regardless of handovers, which thus are transparent to users.

Federated Accounting Center: The Federated Accounting Center receives all accounting information collected by the Accountable Objects (AOs) associated with the M-VPN services provided. The received accounting information is stored in a database referred to as *Accounting Repository*, as depicted by arrow 2 in Figure 4. If the Federated Accounting Center belongs to the VPN provider, it also receives Charging Records (i.e. Accounting Records and charging information) by the charging and accounting subsystem of a third-party CP. Subsequently, it categorizes usage and charging information on a per customer basis, produces Accounting Records (or Charging Records in case it received charging information), and forwards them to the Rating Center (see arrow 3 in Figure 4).

Rating Center: The Rating Center receives accounting information from the Federated Accounting Center and produces Charging Records taking also into account both the contract/QoS violations calculated according to [15] and the charging scheme (see Section 4) for the offered M-VPN contract. In case the Rating Center receives Charging Records (by the charging and accounting subsystem of a third-party CP), the Rating Center performs auditing of the received information, calculates the final charge according to the charging scheme, and produces the corresponding Charging Records. The charging information is stored to the Accounting Repository, as depicted by arrow 4 in Figure 4.

Accounting Server: The Accounting Server is the reference point of the charging and accounting subsystem to other software components, e.g. the billing system. All charging information produced by the Rating Center is forwarded to the Accounting Server (as depicted by arrow 5 in Figure 4). If the Accounting Server belongs to the federated CP, then it forwards the charging information to the Federated Accounting

Center of the VPN provider according to the federation agreements, as depicted by arrow 6 in Figure 4. Otherwise, i.e. in case the Accounting Server belongs to the VPN provider, it produces a bill and forwards it to the customer of the VPN provider according to the customer preferences (e.g. continuously/periodically, etc.).

The Federated Accounting Center, the Rating Center, and the Accounting Server are components related to the M-VPN service provision, and thus are placed in the TINA Service Layer of the architecture; see Figure 4. Finally, in the Access Layer of Figure 4, depicted are the applications enabling interaction of an administrator and a user (or customer) with the charging and accounting architecture; namely, the Accounting Management Console and the Customer Accounting Console, respectively. The information communicated within the charging and accounting subsystem of the federated CP is different than the corresponding information in that subsystem of the VPN provider. Nevertheless, the structure of their charging and accounting subsystems is the *same*.

Our architecture is a lightweight and scalable one, as it involves only the edge CPs contributing to the M-VPN services and the network domain of the VPN provider, regardless of the number of network domains intervening among the edge CPs. This also implies that the "cost" of accounting does not exceed the corresponding typical level. In order for the proposed architecture to function securely over an open network (such as the Internet), a standardized inter-domain AAA protocol should be used for the communication of the charging and accounting subsystems. In fact, for the case of Internet, the architecture is independent of the version of IP that is employed.

6 Comparison with Related Work

So far, we have presented a charging scheme and a charging and accounting architecture appropriate for QoS-differentiated VPN services to mobile users (M-VPN services). There is significant related work in the literature, which we discuss below.

The specification of an Authorization, Authentication Accounting and Charging (AAAC) architecture for QoS-enabled services offered to mobile users is presented in [19]. This architecture applies for an IPv6-based mobility-enabled end-to-end QoS environment for point-to-point communication (rather than VPN), and is based on the current IETF's QoS models, Mobile-IPv6, and AAA framework. According to [19], the accounting and charging procedures in a network domain are encapsulated in separate service equipment (i.e., software modules) and managed by an AAAC server via Application Specific Modules (ASMs). Certain AAAC servers of the network domains are traversed by the service traffic; these servers exchange authorization, authentication, accounting and charging information via a standardized inter-domain AAA protocol. Our accounting architecture can be employed in the architecture of [19], if we view our charging and accounting subsystem as a service equipment of the AAAC server. Also, our charging approach can be applied in the system of [19],

though more meaningfully for charging aggregations of flows (as opposed to individual ones) in the inter-domain level.

Furthermore, [20] presents an accounting and charging architecture and a charging scheme for QoS-enabled VPN services without mobility support. This architecture is based on currently available protocols of the Internet protocol suite and focuses on secured, reservation-based approaches. The key idea of this approach is the establishment of QoS-enabled VPN SLAs through the separate negotiation of the QoS and the VPN parts of the service by respective brokers; these are functioning in each network domain along the path that the traffic traverses. This negotiation results in a SLA establishment between the adjacent network providers along the traffic path. The accounting information is exchanged among the contracted network providers through a signaling mechanism. Our proposed charging and accounting subsystem could be used for charging and accounting within each negotiating domain. A charging scheme is also specified in [20] in an abstract way; this scheme involves a fixed part, a time-based part and a volume-based part. Our analysis of the chargeable characteristics of the M-VPN services and the charging scheme we propose, which is of the same form with the one proposed in [20], can be employed in the detailed specification of a suitable charging scheme for the services considered therein.

7 Conclusions – Future Work

In this paper, we have studied and analyzed the charging and accounting issues involved in the multiparty provision of QoS-differentiated VPN services to mobile users (M-VPN services). We have developed an appropriate charging scheme and a scalable low-overhead architecture for charging and accounting for such services that conforms to the existing standards for accounting. Charging is based on the time-volume approach by Kelly [13]. Using this charging scheme, each provider collects exactly the revenue arising due to the consumption of his resources in the service provision. On the other hand, the users are charged according to the resources they actually use. Thus, the charging scheme is fair for both the providers and the customers. It also provides users with incentives to use the M-VPN services according to their real needs, and indirectly (i.e. by means of selection among simple tariffs by users) give providers their predictions on future traffic. Furthermore, users are offered predictions of their expected charge that corresponds to their tariff selections. As explained in Section 2, our approach will still be applicable in the future Internet where Mobile IPv6 will be used, as well as in networks employing both IPv4 and IPv6. Our contribution includes both the specific charging and accounting architecture presented in this paper, as well as the methodology for the design of this architecture, which can be employed in cases of other Internet services. There are several interesting directions for further research. For example, additional functionality may be provided to customers by means of extra intelligence. This may include support for budget management in the customer side, and efficient M-VPN contract selection or re-negotiation by means of a customer agent.

References

1. INTERNODE (IST-1999-20117). “Interworking Service Architecture and Application Service Definition”. Work Package 2: Deliverable 3, December 2000. URL: http://www.internode.org
2. S. Kent and R. Atkinson. “Security Architecture for the Internet Protocol”. IETF RFC: 2401, November 1998.
3. S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang and W. Weiss. “An Architecture for Differentiated Services”. IETF RFC: 2475, December 1998.
4. K. Nichols, V. Jacobson and L. Zhang. “A Two-bit Differentiated Services Architecture for the Internet”. IETF RFC: 2638, July 1999.
5. G. Dermler, M. Günter, T. Braun and B. Stiller. “Towards a Scalable System for per-flow Charging in the Internet”. In Proceedings of *Applied Telecommunication Symposium*, Washington D.C., U.S.A., April 17-19, 2000.
6. C. Perkins. “IP Encapsulation within IP”. IETF RFC: 2003, October 1996.
7. G. Montenegro. “Reverse Tunneling for Mobile IP, revised”. IETF RFC: 3024, January 2001.
8. F. Adrangi, K. Leung, Q. Zhang and J. Lau. “Problem Statement for Mobile IPv4 Traversal of VPN Gateways”. IETF Internet Draft <mobileip-vpn-problem-statement-req-02>, April 11, 2003.
9. C. Perkins. “Mobility Support for IPv4”. IETF RFC: 3220, January 2002.
10. D. B. Johnson and C. Perkins. “Mobility Support in IPv6”. IETF Internet Draft <mobileip-ipv6-20>, January 20, 2003.
11. C. Retsas. “DSL Technology, DSL Service Models, and Charging DSL Services”. M.Sc. Thesis (in Greek), Computer Science Department, University of Crete, Heraklion, Greece, November 2000.
12. C. Courcoubetis and V. Siris. “Managing and Pricing Service Level Agreements for Differentiated Services”. In Proceedings of *IEEE/IFIP IWQoS’99*, London, United Kingdom, May 31 – June 4, 1999.
13. F.P. Kelly. “Tariffs and effective bandwidths in multiservice networks”. In: J. Labetoulle and J.W. Roberts (eds.), The Fundamental Role of Teletraffic in the Evolution of Telecommunications Networks, *14th International Teletraffic Congress*, *ITC 94*, volume 1a, pages 401-410. Elsevier Science B.V., June 1994.
14. P. Reichl, S. Leinen and B. Stiller. “A Practical Review of Pricing and Cost Recovery for Internet Services”. In Proceedings of *IEW’99:2nd Internet Economics Workshop*, Berlin, Germany, May 28-29, 1999.
15. TeleManagement Forum. TOM Application Note: “Mobile Services: Performance Management and Mobile Network Fraud and Roaming Agreement Management”. Document: GB910B, version 1.1, September 2000.
16. Y. Rekhter, B. Moskovitz, D. Karrenberg, G. J. de Groot and E. Lear. “Address Allocation for Private Internets, BCP 5”. IETF RFC: 1918, February 1996.
17. B. Aboba, J. Arkko and D. Harrington. “Introduction to Accounting Management”. IETF RFC: 2975, October 2000.
18. TINA-C. “Network Resource Architecture Version 3.0: Accounting Management”. Document No. NRA_v3.0_97_02_10, pages 137-180, February 1997.
19. Hasan, J. Jähnert, S. Zander and B. Stiller. “Authentication, Authorization, Accounting, and Charging for the Mobile Internet”. *Mobile Summit 2001*, Barcelona, Spain, September 2001.
20. B. Stiller, T. Braun, M. Günter and B. Plattner. “The CATI Project: Charging and Accounting Technology for the Internet”. *5th European Conference on Multimedia Applications, Services, and Techniques (ECMAST'99)*, Madrid, Spain, May 26-28, 1999.