

# Internet Accounting

Aiko Pras, Bert-Jan van Beijnum, Ron Sprekels, and Robert Parhonyi, University of Twente

## ABSTRACT

This article provides an introduction to Internet accounting and discusses the status of related work within the IETF and IRTF, as well as certain research projects. Internet accounting is different from accounting in POTS. To understand Internet accounting, it is important to answer questions like “what is being paid for” and “who is being paid.” With respect to the question “what is being paid for” a distinction can be made between *transport accounting* and *content accounting*. Transport accounting is interesting since techniques like DiffServ enable the provision of different quality of service classes. Each class will be charged differently to avoid all users selecting the same top-level class. The interest in content accounting finds its roots in the fast growth of commercial offerings over the Internet; examples of such offerings include remote video and software distribution. The question “who is being paid” has two possible answers: the network provider or the owner of the content. The case in which the network provider issues the bill is called *provider-based accounting*. Since this case will become more and more important, this article introduces a new architecture for provider-based accounting.

## ACCOUNTING OVERVIEW

Recent years have shown increasing interest in Internet accounting. The European Commission, for example, subsidizes various research projects on accounting; the Internet Engineering Task Force (IETF) has established an Authentication, Authorization and Accounting (AAA) Working Group; the Internet Research Task Force (IRTF) has established an AAA Architecture Research Group; and so forth.

Since the area of Internet accounting is relatively new, terminology is sometimes still confusing. In the early days of the Open Systems Interconnection (OSI) Management Framework, the term accounting was used to denote the all-embracing process of collecting, interpreting, and reporting costing and charging-oriented information on service usage. This process was divided into the following subprocesses: meter-

ing, pricing, charging, and billing. Metering is the process of measuring and collecting resource usage information, related to a single customer's service utilization. Pricing is the process of determining a cost per unit. The charging process uses this cost per unit to translate the customer's resource usage information into an amount of money the customer has to pay. This amount is then used by the billing process to inform and bill the customer.

Nowadays the term *accounting* is often used as a synonym of the more restricted process of metering. This article, however, will use the term accounting in its original and broader sense.

The purpose of this article is to provide an introduction to Internet accounting, discuss the status of accounting-related work within the IETF, IRTF, and certain research projects, and propose a new architecture for provider-based accounting. The structure of this article is as follows. This section examines the main questions behind accounting: “what is being paid for,” “who is being paid,” and “how are payments made.” The next section discusses different views concerning the question of whether Internet accounting is really needed. The subsequent section discusses the difference between Internet accounting and accounting in plain old telephony service (POTS). The article continues with a section discussing research projects and the state of the art within standardization and research bodies such as the IETF and IRTF. The final section introduces a new architecture for provider-based accounting.

## WHAT IS BEING PAID FOR?

A first thing to recognize is that Internet accounting can be divided into *transport accounting* and *content accounting*. The goal of transport accounting is to charge users for the transfer of packets or bytes over the Internet. The goal of content accounting is to charge users for the content or services that are delivered over the Internet. Whereas the IETF and IRTF focus primarily on transport accounting, many research projects have a broader focus and also investigate content accounting.

The second thing to understand is that charging can be *flat-rate* as well as *usage-based*. The merits of both approaches have been discussed

All authors are members of the Telematics Systems and Services (TSS) management group.

extensively in literature. In [1], for example, charging is structured into a subscription charge and a session charge, each in turn consisting of a setup component and a recurring or usage component. When the session charge is always zero it is called flat-rate; when the session charge depends on the session duration and/or session volume, it is called usage-based.


### WHO IS BEING PAID?

Another important question is whether the bill will be issued by the Internet provider<sup>1</sup> or by the organization who owns the content stored on servers at the remote side of the network. The term *provider-based* accounting will be used to denote the first case; the term *server-based* accounting will be used for the second. Since the transfer of information is generally charged for by the Internet provider, transport accounting usually falls into the category of provider-based accounting. It is also possible, however, that the Internet provider charges for the content, which in turn is provided by others. This case, which is similar to 900 numbers in POTS, is considered to be important for Internet providers and is the subject of various research projects. The last section of this article therefore introduces an architecture for provider-based accounting.

With provider-based accounting it is still possible to distinguish between the case in which the client pays his own Internet provider, and that in which the client pays the remote Internet provider. Corresponding to the POTS where this is an important feature, the latter case will be called *reverse charging* [2]. Reverse charging may also become interesting within the Internet. Take, for example, the case of an art college where students have to create movies. Although these movies are probably interesting to a wide audience, the art college may not be willing to invest in the exploitation of these movies and pay money to its Internet provider for the transport of these movies. If reverse charging were possible, these costs could be paid by the remote user who wants to see these movies. Another example where reverse charging may be useful is for videoconferencing, where a single participant is willing to pay all costs associated with the conference.

### HOW ARE PAYMENTS MADE?

The question "how are payments made" has two possible answers: *inband* or *outband*. Outband payment is still the dominant form of payment and can be implemented by means of credit cards, bank transfers, checks, and so on. Outband payment can be used with provider- as well as server-based accounting. Inband payment is relatively new and can be implemented in terms of cybercash. Although this form of payment will become especially useful for server-based accounting, it may also be interesting for roaming users who connect to multiple Internet providers. Figure 1 summarizes the various accounting options discussed thus far.



■ Figure 1. An overview of accounting options.

## IS INTERNET ACCOUNTING NEEDED?

To answer the question of whether accounting is needed in the Internet, it is important to distinguish between *content accounting* and *transport accounting*.

There is general agreement that content accounting is actually needed. In fact, some people even claim that content accounting will eventually make transport accounting superfluous, since the costs to transport the content can be added to the price of the content. This reasoning does not hold, however, for free content, like the movies produced by the art college in the previous section. In fact, the assumption that the price of the content can take into account all the transport costs has as its final implication that Internet providers should have complete control over all content providers. This assumption is not realistic, and it is therefore likely that transport accounting, in one form or another, will remain necessary.

Which form of transport accounting should remain is still an issue of debate. In fact there are three possibilities: flat-rate, usage-based, or a combination of both. Although flat-rate is currently the prevailing form of transport accounting, several researchers believe that usage-based charging will be introduced soon. The reasoning behind this expectation is [3, 4]:

- Usage-based charging stimulates users to use the scarce resources within the network in an efficient way. Flat-rate does not give such an incentive.
- Without usage-based charging, a small number of users will consume most of the resources: Measurements from early 1998 on the campus net of the University of Twente, which connects 2000 students via switched Ethernet, indicate that 1 percent of the students generate 40 percent of the outgoing traffic [10]. In the long run, less demanding users will not accept this form of cross-subsidy.
- The current Internet will be transformed from a best effort service into a service providing different quality of service (QoS) levels. The technique that facilitates this transition is called *differentiated services* (DiffServ); it allows for the creation of a small number of QoS classes (e.g., gold, silver, and bronze). To avoid every user choosing the gold class (which would then become again best effort), it is important to assign different tariffs to the different QoS classes. To allow users to select for each application the best QoS class, usage-based charging becomes unavoidable.

<sup>1</sup> In this article the term Internet provider will be used as the equivalent of network operator: the organization who transfers user data over the Internet.

*Because of its complexity, it can be expected that usage-based accounting in the Internet will be based on a relatively small number of parameters. Examples of possible parameters are traffic volume, traffic class, and time of day.*

Other researchers do not believe in usage-based charging and argue that transport accounting will continue to be flat-rate. An interesting study is presented in [5], where the reasoning is not based on technical arguments, but on past user behavior with similar communication facilities like ordinary mail, telegraph, and telephone service.

## INTERNET ACCOUNTING VS. TELEPHONY ACCOUNTING

It is interesting to investigate whether providers of POTS, who have many years of experience in telephony accounting, can reuse this knowledge for Internet accounting. To answer this question, it is important to distinguish between billing and metering. Billing for Internet services will be comparable to billing for POTS. As a consequence, experiences in this area can be reused.

Metering, the process of measuring the parameters within the network related to the customer's service usage, is rather different, however. Take, for example, the parameters to measure. In POTS it is common to measure call duration, time of day, and destination of the call (local, national, or international). Since the Internet is connectionless, it is principally impossible to measure call duration. Instead, some providers measure how long users are connected via their local access line to the Internet. Although this is somewhat comparable to measuring call duration, more and more users get permanent connections to the Internet (xDSL, cable, UMTS). As a result it becomes less feasible to use access duration for accounting purposes. It is also questionable whether the destination address will be a useful parameter for Internet accounting. As opposed to POTS, where subscriber numbers include a country and city code, early IP addresses do not contain any form of geographical information. Recent IP addresses that follow the rules defined by the Classless Inter Domain Routing (CIDR) standard do have some notion of location, but this information is less detailed than the geographical information contained in telephone addresses. Also, Domain Name Service (DNS) names, which as well do not really contain geographical information, are difficult to use for accounting purposes. One of the problems of using DNS names for accounting is that a single IP address may be related to multiple DNS names, each registered in a different top-level domain. If the price depends on the DNS name, the problem arises of which DNS name to choose. Although it may be difficult to charge different prices for local, national, and international traffic, it may be quite feasible to charge differently for inter-operator and intra-operator traffic. To implement this, the operator should use the information within its routing tables to determine which customers are connected to its network and which are not. Charging differently for intra- and interoperator traffic may be interesting for providers to attract customers and save on peering agreements.

Because of its complexity, it can be expected that usage-based accounting in the Internet will be based on a relatively small number of param-

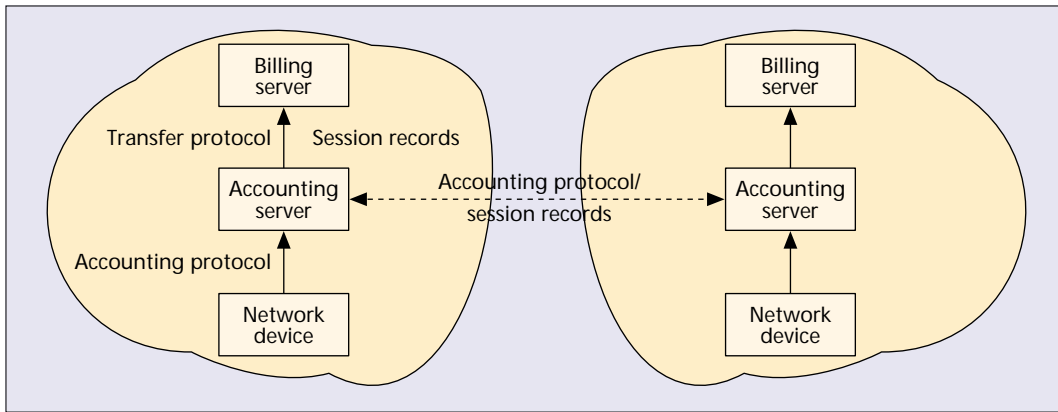
eters. Examples of possible parameters are traffic volume (transmitted, received), traffic class (in situations where DiffServ is applied), time of day, and probably the question of whether or not the destination is connected to the same provider as the sender. These parameters can already be difficult to use. Consider, for example, the case of a congested network that drops packets. The user may still have to pay for these packets, and the interesting case occurs that a provider of a congested network will charge more than a provider of a well designed network.

## RESEARCH AND STANDARDS

Although Internet accounting appears to be different from accounting in the traditional telephone world, researchers and organizations from the telephone world showed interest in Internet accounting at an earlier stage than the developers of the traditional Internet protocols. This difference in appreciation can be understood from the fact that accounting has always played an important role in the telephony world. This world is completely different from the world of traditional Internet researchers, who often come from noncommercial organizations like universities. It is not surprising, therefore, to see that the Advanced Communication Technologies and Services (ACTS) program of the European Commission [6], in which the telecom industry and operators played an important role, already subsidized accounting projects in the mid-'90s. Originally, these projects focused on asynchronous transfer mode (ATM) accounting; examples are the Contract Negotiation and Charging in ATM Networks (CANCAN) and Charging and Accounting Schemes in Multiservice ATM Networks (CASHMAN) projects. As more and more people understood that ATM would never play the role originally envisaged, the ACTS accounting projects started to investigate Internet accounting problems too. A good example of such a project is the SUSIE project, which integrated ATM, IP, and TINA concepts to study charging of premium IP services. Current projects within the European 5th framework Information Society Technologies (IST) program [7] do not pay attention anymore to ATM, but focus on Internet accounting. An example is the Market Managed Multiservice Internet (M3I) project, which investigates differential charging to provide multiple QoS levels [8].

### THE IETF

It is already nearly 10 years since the first RFC appeared on Internet accounting (RFC 1272). This first RFC, which was based on the ideas and terminology of OSI, inspired the Real-Time Traffic Flow Measurement (RTFM) group to define the so-called Meter management information base (MIB). This MIB allows the gathering of usage data from the network and may be important for accounting, performance, configuration, as well as security purposes. Other IETF groups did not pay much attention to accounting, and interest in the subject seemed to disappear. It is remarkable, however, that IETF members regained interest in accounting in the



■ Figure 2. The IETF accounting architecture.

context of ATM; this is somewhat comparable to the developments within the ACTS and IST programs of the European Commission, where there has also been a shift from ATM accounting to Internet accounting. Another development that brought accounting back on the agenda of the IETF was the work on Remote Authentication Dial In User Service (RADIUS). The RADIUS protocol defines how authentication, authorization, and configuration information should be exchanged between network access servers (NASs) and authentication servers. Because RADIUS is widely used, there is a substantial interest to extend RADIUS with new features such as accounting. This interest has led to the formation of a new Working Group (WG) within the Operations and Management Area of the IETF. The name of this new WG is Authentication, Authorization, and Accounting (AAA). It should be noted that the focus of this group is network access, and accounting is only one of their topics. One of the results of this group is the simple accounting architecture shown in Fig. 2.

#### ACCOUNTING PROTOCOLS

The IETF is primarily interested in protocols. Because of this, the AAA WG soon discussed whether or not it could adopt an existing protocol [9]. In fact there were several possible candidates to choose from, including Simple Network Management Protocol (SNMP), COPS, RADIUS/RADIUS++, TACACS/TACACS+, and Diameter. To make a choice, last summer the AAA WG followed a stepwise selection process. First the problems of roaming, mobile IP, NASs, and code-division multiple access (CDMA 2000) were investigated to find all the criteria an AAA protocol should satisfy. Then the members of the AAA WG were invited to submit proposals for their favored protocols; each proposal should clearly explain how the requirements mentioned above were satisfied. As a result, four protocols remained in the race: SNMP, COPS, RADIUS, and Diameter. The four proposals were investigated by a panel of seven persons. These persons were only allowed to judge on the basis of what was written down in the proposals; other information was not taken into account. The conclusion of the panel was to select Diameter. COPS would have been an alternative, but the impression of

the panel was that Diameter was further advanced than COPS. This is somewhat interesting, since Diameter relies on the new SCTP transport protocol, which has still a long way to go. RADIUS dropped out of the contest because it missed many features, and the result of adding these features would be something similar to Diameter. SNMP was not selected because there was disagreement on whether it could satisfy the requirements of authentication and authorization. Although there was general agreement that SNMP would be suitable for accounting, it was not selected because the intention of the AAA WG was to choose one single protocol for authentication, authorization, and accounting. Despite this formal position, it may still be expected that SNMP will continue to play a role in accounting (e.g., in relation to the METER MIB).

#### IRTF

To define a next-generation architecture for AAA, the IRTF decided by the end of 1999 to create a new research group, AAAARCH. The work of this group evolved out of the authorization activities of the IETF AAA group; the goal was to enhance the existing authorization framework to include authentication and accounting. With respect to accounting, the group focuses on policy-based accounting. The motivation for this work comes from the recognition that different providers have different accounting requirements, which change frequently. To cope with these dynamic requirements, a configurable accounting infrastructure is needed, and policies should be used to configure this infrastructure. The work of the research group aims to define the building blocks for this infrastructure and the sequence of messages between these building blocks. The language to describe accounting policies is not a subject of the research, but is expected to be defined elsewhere. Auditing is also considered to be important; for that purpose work is being performed on session and accounting identifiers.

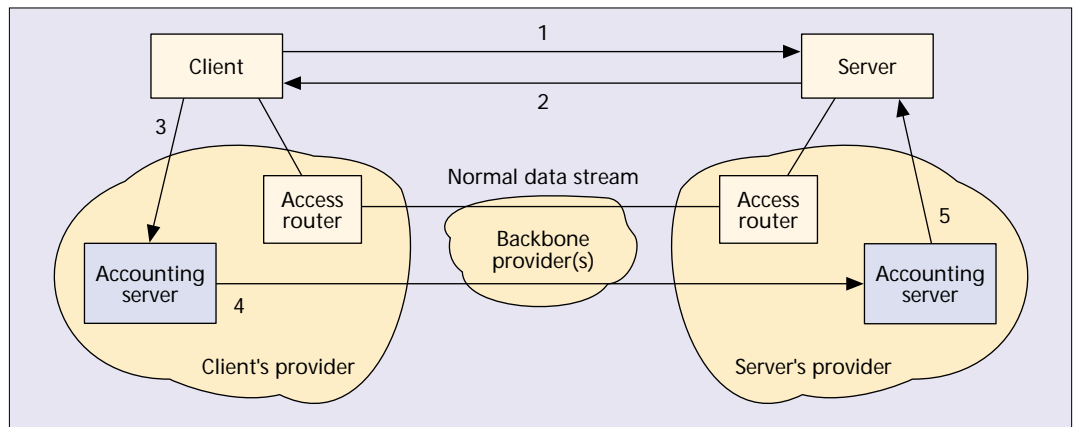
#### PROVIDER-BASED ACCOUNTING

An interesting question concerns the architecture providers should choose to account for transport and content. The issue of content accounting on behalf of others is particularly challenging, since

*Although there was general agreement that SNMP would be suitable for accounting, it was not selected because the intention of the AAA WG was to choose one single protocol for authentication, authorization, and accounting.*



In case of streaming, the client's costs can depend on the amount of data the client has received. The price is therefore actually a price per unit, and the unit may for example be 1000 packets or 1 MByte. Alternatively the price may depend on the amount of time the transmission lasted.



■ Figure 3. The basic architecture for provider-based accounting.

good solutions in this area will allow Internet providers to obtain additional income; the amount of money involved in this can easily exceed the income from the traditional transport business. Because of this importance, and because the IETF and IRTF are not (yet) working on this, the research project Internet Next Generation has developed a new architecture for provider-based accounting [10]. This section discusses the main characteristics of that architecture.

The basic architecture, shown in Fig. 3, includes a client connected to one provider, and a server connected to another Internet provider. Both providers may be connected via zero, one, or more backbone providers. The process starts with a request from the client to the server (1) to deliver content. The server answers with an *accounting request message* (2) indicating that the client should pay for the content to the client's provider.

Depending on the nature of the interaction between client and server, the accounting request message can be implemented as a special MIME type within an HTTP response message. The parameters of the accounting request message are:

- *Server information*, including a readable string identifying the owner of the server, an authentication key, and the server's IP/DNS address.
- *Price*, which may be a single value or, with multiple components, a chain of values. The currency should also be specified, as well as a timestamp to allow determination of the exact exchange rate to cover cases where the client wants to pay, for example, in euros, but the server wants to receive dollars. The timestamp is also needed to ensure that information cannot be reused (replay protection).
- *Content type*, which may take the value streaming or atomic. In case of streaming, the price is actually a price per unit, and can be expressed in, say, megabytes or minutes. Other differences between these types of content are explained later.
- *Accounting server information*, which includes the authentication key as well as the IP/DNS address of the accounting server within the server's Internet provider.

In case the client agrees to receive charged content, it forwards the information within the accounting request message to the accounting

server of its own Internet provider (3). To guarantee that no one can change the information, protection is needed via message authentication codes. Such codes can be generated using algorithms like MD5 or SHA1. If the client has paid all previous bills, the accounting server within the client's provider stores the information for future billing purposes and forwards the *accounting accepted message* to the server's provider (4). If this Internet provider is also willing to participate in the accounting process, it stores the information too and forwards the message to the content server (5).

In case of streaming, the client's costs can depend on the amount of data the client has received. The price is therefore actually a price per unit, and the unit may be, for example, 1000 packets or 1 Mbyte. Alternatively, the price may depend on the amount of time the transmission lasted; in this case a unit may be, say, 1 minute. In both cases the client should periodically send acknowledgments to indicate that it is prepared to pay for the next unit of data or time. In fact, these acknowledgments take the same route and contain the same information as the original accounting accepted message. This message and the subsequent acknowledgments in fact play a similar role to coins in a public telephone booth; if the client does not pay additional coins (acknowledgments), the service stops.

Instead of using acknowledgments, it would theoretically also be possible to introduce an *end of accounting message*. This message should be issued by the client if the client does not want to receive further content. This approach has the disadvantage, however, that the client continues to pay until the server has received the end of accounting message. If the client forgets to issue such a message, or the client's system crashes, or the message gets lost in the network, the client will still be charged. To avoid such problems, acknowledgments were introduced in the design. To keep the traffic generated by the acknowledgments at a reasonable level, it is important to choose a realistic unit size. A unit of 1 min, for example, is reasonable, but a unit of 1 ms is unreasonable, since 1000 acknowledgments/s may overload the accounting system.

Atomic content, such as a piece of software, is only useful if it is complete. If the last byte is not received, the content is useless and the client

will not be prepared to pay. The all or nothing nature of atomic content makes the use of periodic acknowledgments superfluous. Instead, the client should send an acknowledgment after reception of the last byte; only then will the client have to pay. The problem with this approach is that the client may deny reception of the last byte in an attempt to get away with not paying. For that purpose it is important to have nonrepudiation mechanisms.

The actual payments involve three steps:

- The client pays its own Internet provider. This step may be implemented via inband or outband payments. In case of outband payment, the Internet provider may combine the costs of multiple transactions onto a single bill and issue this bill on a periodic basis.
- The Internet provider of the client pays the Internet provider to which the server is connected. Because of the large number of possible trust relationships between Internet providers, it may be necessary to introduce a trusted third party (TTP) [2]. Such a TTP, which may be a bank, takes care of the financial balance between all associated Internet providers.
- Finally, the server gets paid by its own Internet provider.

It should be noted that the architecture of Fig. 3 is not only interesting for content accounting, but can easily be extended to charge for reverse traffic. In that case the accounting servers should configure the access routers such that the amount of traffic flowing between both users will be measured; details can be found in [2].

## CONCLUSIONS

There are several possibilities for structuring Internet accounting. The first possibility is to distinguish between *transport* and *content accounting*. Transport accounting is the subject of various research projects and is being standardized by the IETF. Recently the Authentication, Authorization, and Accounting WG of the IETF selected *Diameter* as the preferred protocol for transport accounting, although other protocols like SNMP may also be used. Within the research world there is no agreement on whether transport accounting will become flat-rate or usage-based.

A second possibility for structuring Internet accounting is to distinguish between *provider-based* and *server-based accounting*. Because of the large amount of money involved, provider-based accounting will become interesting for network providers. Research in the area of

provider-based accounting architectures is therefore important, especially if such architectures allow the client and server to be connected to different Internet providers. This article proposes such an architecture; an important feature of the architecture is that streaming as well as atomic content can be charged.

## ACKNOWLEDGMENTS

This article is based on the work performed within Work Unit 5 of the Internet Next Generation project [10]. This project is part of the Dutch Gigaport program, and sponsored by the Telematica Instituut (TI). The following organizations are members of the Internet Next Generation project: the University of Twente (UT), Ericsson ETM, Ericsson EMN, KPN Research, and TI. We would like to thank all members of the project for their contributions.

## REFERENCES

- [1] F. Hartano and G. Carle, "Policy-based Architecture for Internet Differentiated Services," *Proc. IFIP 5th Int'l. Conf. Broadband Commun.*, Nov. 1999.
- [2] R. A. M. Sprenkels *et al.*, "An Architecture for Reverse Charging in the Internet," *Proc. IEEE IPOM Conf.*, Cracow, Poland, Sept. 2000.
- [3] P. Varaiya, "Demand and Provisioning of Quality-Differentiated Internet Access," keynote speech, *INFOCOM'99*, New York, NY, Mar. 1999, <http://www.comm.utoronto.ca/infocom/keynote.html>.
- [4] R. Edell and P. Varaiya, "Providing Internet Access: What We Learn From INDEX," *IEEE Network*, vol. 13, no. 5, Sept.–Oct. 1999; [http://www.path.eecs.berkeley.edu/~varaiya/papers\\_ps.dir/networkpaper.pdf](http://www.path.eecs.berkeley.edu/~varaiya/papers_ps.dir/networkpaper.pdf).
- [5] A. Odlyzko, "Internet Pricing and the History of Communications," <http://www.research.att.com/~amo/doc/networks.html>.
- [6] European ACTS (Advanced Communication Technologies and Services) program: <http://www.cordis.lu/> and <http://www.uk.infowin.org/>
- [7] Euro. IST Program: <http://www.cordis.lu/ist/projects.htm>
- [8] M3I: Market Managed Multiservice Internet project: <http://www.m3i.org>
- [9] C. Metz, "AAA Protocols: Authentication, Authorization and Accounting for the Internet," *IEEE Internet Comp.*, Nov.–Dec. 1999, pp. 75–79.
- [10] Internet Next Generation project: <http://ing.ctit.utwente.nl/WU5>

## BIOGRAPHIES

AIKO PRAS (pras@ctit.utwente.nl) is a senior researcher at the Centre for Telematics and Information Technology (CTIT) of the University of Twente, the Netherlands.

BERT-JAN VAN BEIJNUM (beijnum@cs.utwente.nl) is an assistant professor at the University of Twente.

RON SPRENKELS (sprenkel@cs.utwente.nl) is a Ph.D. student at the University of Twente.

ROBERT PARHONYI (parhonyi@cs.utwente.nl) is a Ph.D. student at the University of Twente.

Research in the area of provider-based accounting architectures is important, especially if such architectures allow the client and server to be connected to different Internet providers.