# PATH COUPLED ACCOUNTING MECHANISMS FOR "ALL IP NETWORKS"

**Andreas Klenk*, Philipp Schlicker*, Ralph Kühneº*, Ali Fessi*, Changpeng Fanº, Falko Dressler†*, Georg Carle***

*University of Tübingen, Computer Networks and Internet, Tübingen, Germany,
†University of Erlangen, Computer Networks and Communication Systems, Erlangen, Germany
ºSiemens AG Communications, Mobile Networks, Berlin, Germany

**Keywords:** Accounting, All IP Networks, Meter Selection Algorithms, Distributed Load Balancing, NSIS

## Abstract

Accounting in Beyond 3G (B3G) networks is especially challenging, because of the rapidly increasing number of new services and tariff models introduced or to be introduced. Services are not provided statically anymore but must be deployed in short time spans to meet the volatile market demands. As the complexity of the services grows the demand for an equally flexible accounting mechanism increases as well. However nowadays' static accounting approaches with dedicated meters for certain metering tasks might scale badly in this envisioned dynamic environment. Innovative services can not be deployed if expensive accounting tasks in terms of metering resources are required and the dedicated systems cannot guarantee to handle the additional load. On-path configuration of the meters can provide for a dynamic metering configuration. The accounting configuration can be performed agnostic of the roles of the accounting entities in the B3G network architecture allowing a more flexible accounting. The distributed nature of on-path accounting allows for fault tolerance and provides load-balancing. Thus existing metering resources are utilized better and this enables the introduction of demanding accounting mechanisms for new services.

## 1 Introduction

The 3GPP currently investigates *All IP Networks (AIPN)* as perspective for the core network to provide integrated services, independent of the access systems[3]. Charging and accounting is core functionality for the network providers. In contrast to the term accounting in the 3GPP context we define accounting as "collecting data about resource usage which includes gathering, transporting, formatting and storing data about chargeable events".
Since existing accounting mechanisms cannot be adapted easily to the new requirements alternative accounting concepts must also be evaluated. On-path signalling for metering configuration[5] has introduced a new philosophy for the accounting configuration. Instead of performing accounting constantly for all traffic flows and sorting out superfluous results later on the accounting should only be active based on a configuration. The accounting configuration is independent of the particular accounting entity (meter) that performs the measurements. Hence the metering configuration is decoupled from the network architecture. Meters are configured to gather accounting data about particular data flows which pass by. In contrast to single dedicated meters for an accounting task (e.g. accounting at SGSN/GGSN/CSCF/MRF…), multiple meters can provide the same measurement functionality. The meters are configured at session setup and can share the burden of the metering dynamically. Load balancing on the data path is a technique to avoid the overload of single meters and provide for high system reliability and will be elaborated further in this paper.

### 1.1 Accounting challenges for beyond 3G networks

The vision of the introduction of a plethora of innovative services and complex service compositions that are present in future networks impose among other things challenges on the accounting and charging. The 3GET[1](3G Evolving Technologies) project proposed an architecture for service-oriented charging as a solution to adapt the configuration dynamically for any combination of services. Accounting systems face additional challenges.
Low operational costs and accuracy of the accounting are of prime importance for mobile network operators (MNO). Building the core networks with IP technology helps reducing the cost by using standard hardware from the Internet world and achieving higher bandwidth utilization. Today's growing demand for Internet content delivered by the MNO sparks the need for IP networks even more. Beyond 3G (B3G) networks are envisaged to include a packet-switched domain which provides data communication for already deployed services like GPRS and upcoming services provided over the IP Multimedia Subsystem (IMS)[2]. The IMS is based on the packet switched domain to enable, for example, multimedia services and location based services. Diverse accounting functionalities are currently standardized by the 3GPP to provide accounting and charging. These solutions are

[1] Project homepage: http://www.mobile-accounting.org/

comparably static in configuration and maintenance and are the result of a gradual development of individual accounting mechanisms for IP networks. The introduction of new services requires complex accounting and charging configurations, for example, a MMS charging should not account traffic volume and could either be implemented by specifying separate APNs for this service or by configuring the billing and charging system to ignore charging data records which contain volume information. It remains unclear if these mechanisms can cope with the expected multitude of service configurations and accounting demands.

The vision of All IP Networks must not necessarily include the familiar centralized components as currently under discussion for B3G networks for all usage scenarios but sometimes might rely solely on IP routing. This option deserves further investigation. Nowadays, meters are usually located on dedicated nodes inside the core network. All the traffic needs to pass via these points and therefore may deviate from the topological optimal path. Routing efficiency in AIPN would severely degrade if routing is forced through only a handful of points (route pinning). As an example, even if the traffic of the originator is intended for a destination in the same cell, all traffic has to be routed to the central metering entities and back.

Another challenge is how to dimension the capacity of the specialized accounting components. The degree of over provisioning might be limited by the hardware costs and must therefore be carefully calculated. However, the configured accounting components of a session must all collaborate, the failure of only a single component might result in significant monetary loss. Hence the operator is forced to employ much more accounting resources than required at each component to cope with peak usage and has to use highly reliant and normally very expensive hardware components.

### 1.2 Accounting on the data path

In contrast to 3GPP accounting configuration this paper will focus on an alternative approach. Configuration messages are sent on the data path containing only a description of the accounting tasks and *no* static assignment of the tasks to a particular meter. Thus the configuration is agnostic about which meters will perform the requested accounting. The accounting task is dynamically assigned to a meter on the data path, taking its accounting capabilities and capacity into consideration. There are potentially many meters on the data path which are capable of measuring the requested characteristics of certain user sessions. On-path signalling determines the meters which take over the metering task.

Such a configuration mechanism is usable for the configuration of the existing 3GPP accounting components. The same signalling message can establish the accounting configuration for heterogeneous charging solutions, for instance, by configuring the SGSN for packet counting, the CSCF for SIP signalling evaluation and the flow based charging component for measuring the traffic flow. Because the signalling messages are sent along the data path and are processed by each meter it can be avoided that packets are measured redundantly wasting resources.

Path-coupled accounting is more extensible than the current approach, because required accounting capabilities are dynamically identified and configured. New services with special accounting requirements can easily be integrated into the system by specifying an appropriate accounting policy.

Load balancing of the metering tasks over the available meters leads to a better utilization of the available accounting resources. If the existing meters cannot handle the additional traffic or cannot provide the desired functionalities it is feasible to deploy new meters in the network without the need to modify the existing accounting components. Additional meters can also provide for a higher accounting capacity in the system.

Path-based signalling provides also fault tolerance to provide high system reliability. If a meter fails or a path change occurs, say if the attachment to the access network changes or the routing redirects the traffic over a different path, the system autonomously reconfigures itself so as not to lose packets.

Path-coupled signalling is well suited for accounting at the data path, say for flow based accounting. However it is not feasible for off path accounting, for instance, event based accounting.

### 1.2 Structure of this document

The rest of the paper is organized as follows: section 2 discusses which algorithms can be used to select an appropriate meter and what dependencies exist between meter selection and configuration. A short description of the exchange of the measurements of the meters in form of accounting data records (ADR) is presented in section 3. Some conclusions summarize the paper.

## 2 Meter Selection

The selection of the right meter for the measurement task is of high importance for the correct functioning of the accounting system. Hence the meters on the data path must be able perform the accounting timely and accurately. To enable on-path accounting, we must make these assumptions:

a. data paths are stable over larger periods of time
b. link changes on the data path can be detected
c. at least one adequate meter is available for metering on each data path between data sender and receiver.

These assumptions are not as restrictive as they appear to be. Data paths are indeed stable over larger periods of time, changes mainly occur during changes of the access network. Path-coupled signalling protocols like NSIS possess means to detect path changes on the fly. We will elaborate the use of NSIS for on-path signalling further in this paper. The 3GPP architecture provides already many components with metering capacity on each data path.

We refer to the term *metering policy* for a definition what needs to be accounted for a user session. This policy describes which parameters are to be measured by a meter and which format to be used to report the results to the accounting. Usually many meters are situated on the data path and are potential candidates for metering the specified parameters of a user session. Note that the capabilities of the different meters may vary, for instance, some may only be equipped for counting the number of packets and utilized bandwidth, while others might be able to perform an analysis of the packets on the application level by inspecting SIP signalling packets or HTTP requests.

There might be multiple metering entities on a certain data path that are able to fulfil a specific configuration request. The selection algorithms discussed use on-path signalling directed at the destination of a flow. This configuration message will follow the same path as the data packets which must be accounted and therefore passes all potential metering candidates. Hence, the on-path signalling is well suited to discover appropriate meters. In case the traffic backwards to the sender needs to be accounted as well, a separate metering selection for the backward flow is required. Note that due to possible routing asymmetries different meters might become involved.

### 2.1 Path Coupled Accounting Configuration with the Metering NSLP

The IETF Next Steps in Signalling (NSIS) working group is currently standardizing a protocol suite for path-coupled signalling. The NSIS protocol suite consists of a generic NSIS Transport Layer Protocol (NTLP) [9], which is responsible for moving signalling messages back and forward along the data path; and different signalling applications running on top of the NTLP called the NSLPs (NSIS signalling layer protocols). One of these applications is, for instance, signalling for Quality of Service (QoS), which substitutes the Resource Reservation Protocol (RSVP, [10]). In fact, NSIS was inspired by RSVP. NSIS re-uses, where appropriate, the protocol mechanisms of RSVP while supporting a flexible signalling mechanism, such as it can be extended to support new signalling application and is not restricted to QoS signalling. One such a possible signalling application is the configuration of metering entities along the path, which is also currently proposed at the IETF for standardization [4], and which plays an important role for the accounting mechanism presented in this paper.

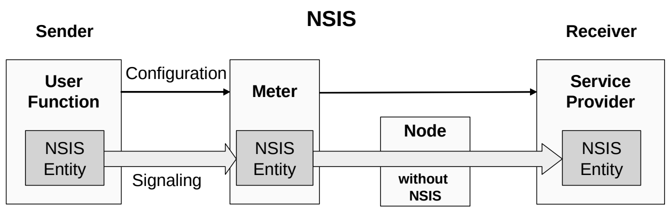

Figure 2: NSIS Signalling

As a result, the Metering NSLP is path-coupled protocol for the configuration of metering entities. Such a protocol allows the dynamic discovery and configuration of accounting entities along the data path.

A so-called "CONFIGURE" message carries the configuration information as a set of accounting tasks that need to be performed by eventually different accounting entities along the path. For example, one accounting entity should count bytes belonging to a given data flow and report the accounting records to a specified collector, and an other accounting entity should report regularly the bandwidth consumed by this data flow. The CONFIGURE message carries also policy objects, which are evaluated at each metering entity to verify that the signalling initiator is authorized to perform this configuration. Security is a very important concern for the Metering NSLP in order to avoid customers misusing the configuration capabilities for their own benefit.

Since all NSIS protocols are based on a soft state, configuration of metering entities need to be refreshed in order to avoid the metering entities stopping accounting for a running application session. This is realized by a "REFRESH" message, which refers to the previous "CONFIGURE" message (and therefore does not need to carry the configuration information again) and extends the lifetime of the configuration.

Route change is a critical issue for the Metering NSLP. A data flow might take a new path during a session, for example, due to a mobility event, or the failure of a router on the path. If the accounting entities that were gathering accounting records for this data flow are on the old path, then new accounting entities need to be discovered and configured on the new path. Route change is a general problem for other NSIS signalling application as well. Therefore, the NTLP provides mechanisms for detecting route changes, for example, based on monitoring changes in routing tables or inference from changes in signalling packet TTL. When the NTLP detects a route change, it informs the concerned NSLPs. When the Metering NSLP receives a notification from the NTLP that a route change has happened, it needs to re-initiate

### 2.2 Requirements for Meter Selection Algorithms

Path-coupled meter selection can only act as an alternative for configuration if it meets a number of requirements. The signalling for the configuration of the metering takes some time. This latency at service initialization must be as small as the time for the accounting setup in 3GPP architectures. Metering mechanisms in use today provide for a high reliability. Hence the reliability must be as good for path-coupled accounting using replication of the metering for fault tolerance and load balancing to avoid overload at meters. Sporadic traffic peaks of the service usage exacerbate the estimation of the load at each meter, thus the selection can only be made on a prediction of the future load. Therefore, an optional support for load shift in the face of overloads can increase the reliability further. Load shifts can be initiated before a meter reaches a point where it cannot account for all flows anymore.

Low administrative demand is one central requirement for the path-coupled configuration. It can be provided because of the strict separation between the specification of the metering and the configuration of a particular entity. The system must be extensible to account unforeseen innovative services and to introduce new accounting capacities seamlessly. In the face of a failure the system must solve the issues autonomously.

### 2.3 User Sessions

In a service-oriented charging approach the main focus is on user sessions being instances of some services that are provided by the network and that need to be accounted. Such services might be complex aggregates of other (sub-)services. For instance, an interactive video-conferencing service session might be a composition of a video and audio session, a session that provides for a blackboard that all participants can concurrently write on, as well as a data transfer session that allows for sending and receiving files.

In order not to account and finally bill all service sessions separately and also to prevent a need for an extensive and potentially expensive correlation on the billing level, an accounting solution must be able to differentiate between flows and events that belong to such an aggregate session and others that do not.

When accounting is configured for a user session, information must be provided that allows for such a differentiation. Normally, it will be some pattern against which an encountered event or data flow might be compared. In case of e.g. volume-based accounting, the so-called IP 5-tuple is a choice, consisting of source and destination address and port, as well as the used protocol for the data flow that has to be accounted. Another alternative is flow labelling, where the data source, or an ingress node (e.g. a video streaming server) marks all sent packets with a unique label enabling the accounting entities to recognize which packets belong to the respective user session.

These information form part of a metering policy passed along the data path and is also a basis for meter selection.

### 2.4 Static Selection Algorithms

We use this term to describe meter selection algorithms which select deterministically the same meter if the same parameters are present. The Selection is primarily based on the capabilities of a service and the accounting scenario. The semantic of the accounting task determines sometimes that no arbitrary meter can be selected but dedicated meters must be configured. The M-NSLP specifies options for such a selection, for instance, by specifying "FIRST" to trigger accounting at the first available meter on the data path which supports the requested metering task. We term an algorithm as static if it assigns the accounting task reproducibly to the same meter if the signalling message is sent again and the data path did not change. Other selection criteria could be derived from rules based classification of IP-addresses. One drawback of static selection algorithms is that dynamic conditions like load of the meters cannot be regarded.

### 2.4 Dynamic Selection Algorithms

Unlike the static algorithms dynamic algorithms can adapt to variant parameters, say system load at a meter. Hence the outcome of a dynamic selection is unknown and might be different for each signalling. These dynamic parameters are an important factor when a choice must be made which meter is responsible for a metering task. Disregarding these parameters can result in a load imbalance which can lead to an overload of certain meters while others would still have spare capacity. Another important factor are meters which are specially equipped for certain measurement tasks while other meters might also possess the same ability but are less efficient. For a static algorithm it might be a question of chance that the more adequate meter gets selected, for instance if the better equipped meter happens to be the first meter on the path. Thus dynamic selection algorithms can avoid load imbalance but introduce more complex signalling.

*Round Robin* is an example of a simple selection algorithm which can distribute the measurement tasks equally along one data path. It assigns the single metering tasks in sequential order to the meters and starts again at the beginning of the data path if the end is reached. This algorithm is simple to implement but has some major drawbacks. It operates well if the meters only receive their tasks over a single Round Robin function. If meters get assignments from different sources overload can still happen. Imagine an intersection of two Round Robin signalling paths with one meter be included in both paths. It is obvious that the meter at the intersection would get more tasks assigned then the other meters.

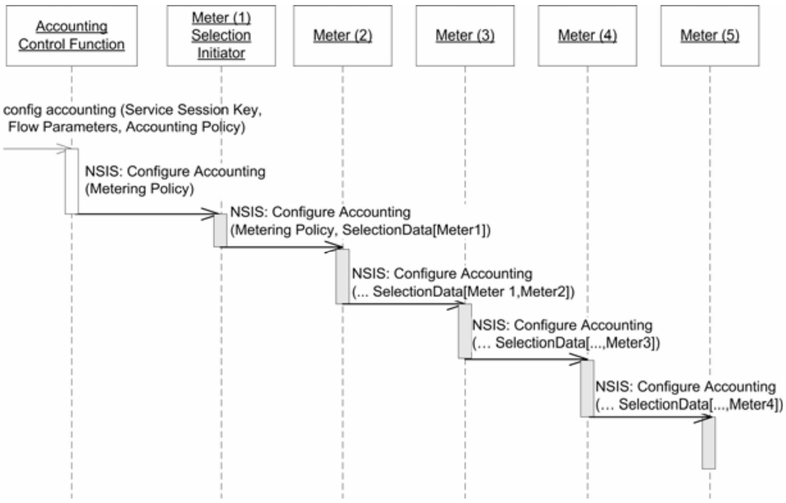


. Figure 3: Meter Selection Message Exchange

The load distribution problem requires more sophisticated algorithms which take the load and/or the meter capacity into account. Load balancing algorithms are a class of algorithms which aim for an equal distribution of the measurement load. We will discuss especially load balancing which is based on a *performance criteria* parameter.

It is important to use adequate performance criteria to characterise the current load of a meter regarding its ability to measure a certain criterion. Meters might be specialized on the efficient measurement of one characteristic, but perform worse for other metering tasks. Furthermore the hardware platforms might differ and exhibit unequal behaviour. Following are some important task dependent factors which can help to derive the performance criteria. Most of them can either just reflect the current status of the factor at the time of the request or the median of the factor value of some broader period of time to smoothen the effect of peaks in the system:

- Percentage of available CPU Load
- Prediction of number of packets a meter can handle additionally
- Prediction of number flows a meter can handle additionally
- Prediction of available bandwidth at a meter

The load balancing algorithm can use the performance criteria to base its operation on it. During the course of the load balancing, all meters on the path publish their current status expressed as performance criteria. The algorithm can conclude from the gathered information, which meter is best suited for the metering task. On-path signalling is used to send a selection message containing the metering policy along the data path (see figure 3).

An external entity, for example a AAA system, triggers the selection by providing the *metering policy* to the first node on the data path, we call this node s*election initiator*. A s*election message* is generated to signal the requirements of a metering task. This message contains the metering policy describing the session and flows and additionally a specification of what needs to be measured as well as a definition which information the resulting ADRs should contain. The NSIS transport layer protocol (NTLP) will forward this message to all NSIS enabled nodes on the data path hop by hop. All nodes with metering capabilities will evaluate the NSIS message and will process the contained *metering policy*. The meter will decide then if and to what degree it can fulfil the specified measurement task.

The capacity data is transferred to a policy decision point which can now select an appropriate meter. However the gathered capacity data can only reflect the load situation in the past and can provide for an estimation of the capacity at the meter. Hence the load at the meters might be slightly imbalanced but it can be expected that in the course of many selections this effect diminishes.

Because the load balancing is executed concurrently at different paths in the network we also refer to it as *distributed load balancing*.

## 3 Measurement Reports

Meters generate measurement reports for the assigned user sessions. The amount of the generated accounting data records(ADR) demands for efficient processing to avoid an overload at the charging system.

Hence, the applicability of an accounting system also depends on the efficient processing of ADRs. As such, the distributed nature of the metering can be complemented by the concept of distributed pre-processing of ADRs. A collector component as discussed in the context of M-NSLP can provide this functionality.

The collector is responsible for the pre-processing of the ADRs, by aggregating the reports, therewith reducing the amount of reports that must be sent to a charging system. Though it may mean a relief for the charging system a collector could not always be used to aggregate the data as far as possible. If the collector aggregates ADRs over a long period of time and then sends the results to the charging system, for example, it might cause problems for online charging where timely accounting information is required.

## 4 Conclusion

An alternative approach to accounting in B3G networks is discussed in this paper. The fundamental concept is to only actively monitor flows which require accounting and avoid superfluous measurement activities. The focus was on the motivation of different configuration algorithms which can be used on top of path coupled signalling. Central benefits were the flexible discovery and topology agnostic configuration of meters on the data path. Hence, the dynamic accounting configuration allows for an easy deployment of innovative services with low administrative costs. However, the flexible configuration approach of the accounting comes at the cost of a signalling specific configuration overhead. The impact of the required signalling and the feasibility of such an architecture is subject to ongoing research.

**Acknowledgements**

This research was support by the German Federal Ministry of Education and Research (BMBF) and Siemens AG. We'd like to thank our colleagues Uwe Föll, Uve Reimer, and Mehran Roshandel for their contribution. The helpful discussions in the Metering NSLP design team are also gratefully acknowledged.

**References**

[1] 3GPP TR 21.902: "Evolution of 3GPP system".
[2] 3GPP TR 23.228: "IP Multimedia Subsystem (IMS)"
[3] 3GPP TR 23.978: "All-IP Network (AIPN) feasibility study".
[4] F. Dressler, A. Fessi, J. Quittek, C. Kappler, and H. Tschofenig, "NSLP for Metering Configuration Signaling," Internet-Draft, draft-dressler-nsis-metering-nslp-02.txt, July 2005.
[5] A. Fessi, C. Kappler, C. Fan, F. Dressler, and A. Klenk, "Framework for Metering NSLP," draft-fessi-NSIS-m-nslp-framework-00.txt, Internet Draft, (2005).
[6] U. Foell, C. Fan, G. Carle, F. Dressler, and M. Roshandel, "Service-Oriented Accounting and Charging for 3G and B3G Mobile Environments," 9th IFIP/IEEE International Symposium on Integrated Network Management, (2005).
[7] R. Hancock et al. "Next Steps in Signalling: Framework", Internet Draft, draft-ietf-NSIS -fw-01.txt, (2002).
[8] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," Internet Draft, draft-ietf-mpls-arch-01.txt, (1998).
[9] H. Schulzrinne, R. Hancock, "GIMPS: General Internet Messaging Protocol for Signalling", Internet-Draft draft-ietf-nsis-ntlp-04, (2004).
[10] Zhang et al. "RSVP: A New Resource ReSerVation Protocol", IEEE Network Magazine, (1993).