

SUSIE - Charging and Accounting for QoS-enhanced IP Multicast

Georg Carle, Felix Hartanto, Michael Smirnov, Tanja Zseby

GMD FOKUS

Kaiserin-Augusta-Allee 31

D-10589 Berlin, Germany

[carle, hartanto, smirnov, zseby]@fokus.gmd.de

ACTS Project SUSIE - AC320

The SUSIE project deals with the development of a charging approach for QoS-enhanced IP services or Premium IP services provided over ATM. This includes the development of charging schemes as well as the design and implementation of the charging systems. The charging approach is evaluated and validated in the SUSIE trials. In the following, selected achievements of the SUSIE project related to Next Generation Internet are presented.

A reference model for a charging and accounting service has been defined. The model can be configured to meet requirements of both IntServ and DiffServ. An initial implementation, called VIPCAS (Value Added IP Charging and Accounting Service), has been started within the SUSIE project.

Charging and Accounting Reference Model

We developed a reference model for classifying charging, accounting and closely related processes, and for describing their interaction (see Figure 1). At the right, five layers are shown that encompass processing for charging and accounting. A configuration plane allows for providing configuration parameters for the processing layers.

The configuration can be done on-line using signalling, as for Integrated Services, or off-line using management tools, as for Differentiated Services. Configuration parameters are derived from pricing policy, accounting policy and metering policy. These policies are provided by interaction of dedicated servers with the corresponding entities of the configuration plane.

The *metering layer* performs metering of resource usage. Metering must allow to distinguish the two types of network resource: reservation of network resources, and actual usage of network resources. This distinction is useful as resources that are reserved but not used by a user may be offered to a different user, but usually under different conditions (lower price). Charging schemes may reflect this difference, e.g. by charging separately for reservation, and for actual usage. In case of multicast, depending on the cost sharing schemes, the meters can be placed at the edge routers only or also at the splitting points.

The *meter reader layer* encompasses functional entities that access data provided by metering entities and forwards it for further processing to the Accounting Processing Layer. For supporting multicast charging, this layer is also responsible for selection of appropriate meters (meter placement). Transfer of metering data to the meter reader can be initiated explicitly (the meter reader initiates transfer of metering data) or implicitly (after a triggering event such as detection of a new flow, the meter initiates transfer of metering data to the meter reader).

Entities of the *accounting processing layer* process usage data collected by meter readers, try to consolidate them based on service parameters and create accounting data sets (i.e., accounting records) which will be passed to the charging layer for pricing assignment. For supporting multicast charging, this layer is also responsible for reconstructing the multicast topology including splitting points where required by the cost sharing scheme. Additionally, the layer is also responsible for distributing collected usage data to other domains in a multi-provider environment.

The *charging layer* derives costs for accounting data sets based on service specific tariffing parameters. Different cost metrics may be applied to the same usage of resources, and may be evaluated in parallel. A detailed evaluation of the resource usage can be used for generating bills to the customer, or for internal analysis (auditing) by the service provider. A simple evaluation of current costs can be used for displaying an estimation of accumulated costs for the service user, or for control purposes by the customer organisation or by the provider. For charging of multicast services, *cost allocation* assigns costs to specific endpoints, such as sender(s) and receivers of a multicast group.

The *billing layer* translates costs calculated by the charging layer into monetary units and generates a bill for a customer. This process may combine technical considerations with economic considerations, such as volume of resources used by the customers, and marketing methods (e.g. offered discounts).

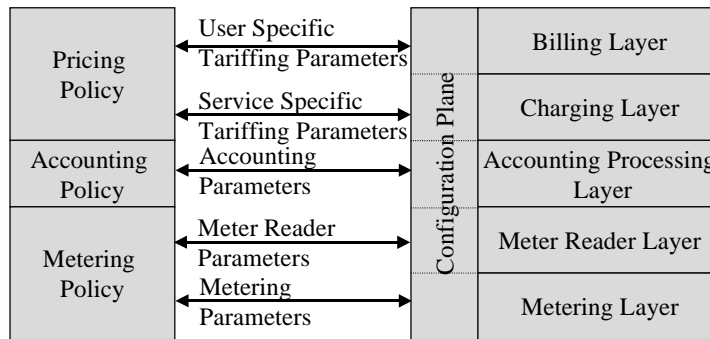


Figure 1: Charging and Accounting Reference Model

Value Added IP Multicast Charging and Accounting Service (VIPCAS)

Based on the reference model introduced in the previous section, we introduce an architecture for implementing a charging and accounting service, extending the architecture presented in [CaSZ98]. Figure 2 shows a sample architecture designed to support this service referred to as VIPCAS (Value Added IP Charging and Accounting Service). Despite its simplicity, the architecture provides generic components to build a more complex platform for charging and accounting service. In explaining the architecture we show how the layers of the CA reference model introduced in the previous section can be realised.

In this architecture, metering takes place only at the edge routers of the provider domain. This allows for core routers to be simple and fast, while more complex metering tasks can be limited to edge routers. Our platform uses a meter that is conformant to the Real Time Flow Measurement Architecture (RTFM) of IETF [RFC2063, BrMR98]. For the configuration of a meter and for the collection of data, SNMP is used with a special meter MIB [RFC2123].

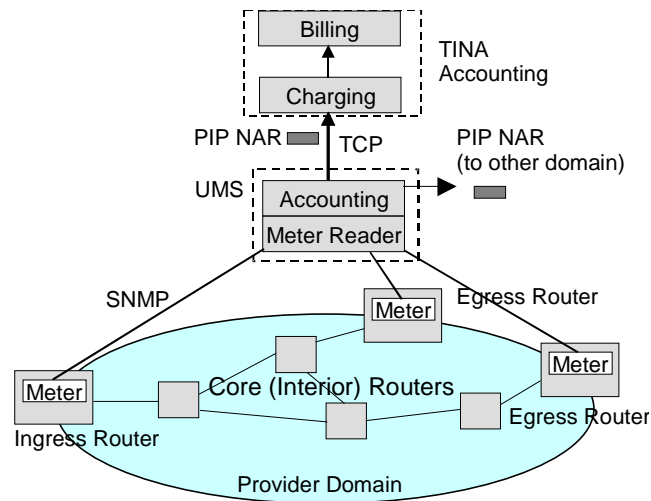


Figure 2: Generic CA Architecture

The data from the meters is collected by a meter reader. An accounting data structure is filled with the collected information by the accounting layer. For transporting the collected information to the charging layer, we have specified a data structure, referred to as PIP-NAR (Premium IP Network Accounting Record), as discussed in the next section.

Since filling of accounting data structures can be done immediately after metering data has been collected, our solution combines meter reader and accounting functions in a single entity, the usage metering server (UMS).

Filled PIP-NARs are forwarded to a charging server where a charging scheme is applied to the data. In our implementation, the charging server is implemented within a TINA/CORBA based accounting system. The UMS passes the accounting data to the TINA accounting system through an interface, called TINA Gateway, which has been defined within this SUSIE project (see section on TINA Gateway).

Within the TINA accounting system, subsequently calculated charges are passed to a billing server, where information about network related charges is used together with additional tariff information (user information, discounts, etc.) to generate bills and to prepare current billing information for users in real-time (hot billing).

Premium IP Network Accounting Record (PIP-NAR)

The PIP-NAR data structure contains reserved and used resources for an IP flow. It carries a measurement point identifier and a record description to support different styles of PIP-NARs (uni-directional/bi-directional, DiffServ/IntServ style, and others). For unique identification of the flow, a flow description is used. A detailed specification of the PIP-NAR can be found in [ETSI98].

Measurement point identification
Record Description
Flow Description
Reserved Resources
Used Resources
Data Extension

Figure 3: Basic Elements of the PIP-NAR

The semantics of the reserved resources section differ for the Integrated Services model and for the Differentiated Services model. Used resources contain the data volume and the number of transmitted packets. The PIP-NAR record format is extensible, where future metrics can be included in the Data Extension section using the TLV (type, length, value) syntax. Currently defined extensions are distance and burstiness.

TINA Gateway

The TINA Gateway provides an interface between the “Internet world”, where usage metering data is collected, and the “TINA world”, where the collected information is charged. The TINA gateway provides the following functions:

- Collect data from network nodes (UMS)
- Pre-process data if necessary
- Fill PIP-NAR
- Pass PIP-NAR to TINA accounting

Two logical entities are defined, which can either be co-located in one physical device or reside in different physical devices (Figure 4). The Usage Metering Server (UMS) takes over the meter reader functions, the pre-processing of data and the filling of the PIP-NAR structure. It then communicates the PIP-NAR to a CORBA Client (CC). This CORBA client provides an Internet style communication mechanism, like a TCP connection. Since the CORBA client is already part of the TINA/CORBA system, it can easily forward the data to the accounting system by using CORBA.

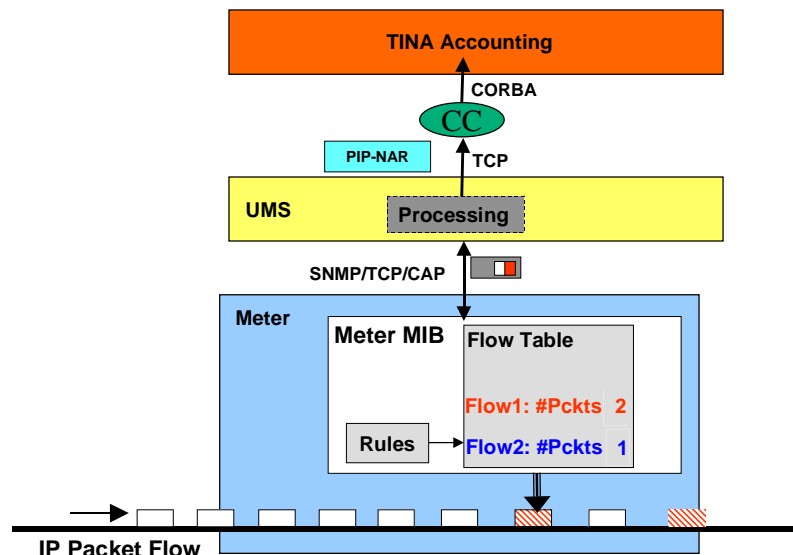


Figure 4: The TINA Gateway: Basic Architecture

The UMS fills PIP-NARs with the provided information. Since the data reported from the network are mostly not presented in the form of the PIP-NAR, some pre-processing might be necessary in order to fill the PIP-NAR structure. Furthermore, some parameters should be calculated in the UMS, because they depend on the combination of records from several locations (e.g. distance) or measurement times (e.g. burstiness). In order to calculate these parameters, it is necessary to store information from previous records in the UMS. If parameters can not be calculated, the field in the PIP-NAR is marked unknown.

As shown in Figure 4, different communication mechanisms can be used for passing the usage metering data to the UMS. One possibility is to use SNMP. In [RFC2063], the Real Time Flow Measurement (RTFM) working group, within the IETF, has defined a specialised meter MIB which can be used for the collection of accounting relevant data. An extension for the monitoring of RSVP messages is defined in [Maic98, HaBR98]. The work of the RTFM group provided a useful framework for the implementation of the UMS.

We use the Network Traffic Meter (NeTraMet) metering software [Brow97, RFC2123] as the basis for the monitoring of accounting relevant data. NeTraMet conforms to the accounting architecture defined by the IETF Real Time Flow Measurement WG [RFC 1272, RFC2063, BrMR98]. The interaction between the meter and the meter reader is based on SNMP. The NetTraMet meter is a stand-alone SNMP agent and stores the measured data in a meter MIB defined in [RFC2064, Brow98b]. Several extensions already exist, like a specialized version for using Cisco's netflowd [Brow98c] and an extension for monitoring RSVP states [Maic98]. Further extensions are planned [HaBR98]. The rules for the pattern matching process in the meter are given in special rulesets [RFC2063]. For simplification of the creation of rulesets, a Simple Ruleset Language has been developed [Brow98a].

Summary

In the first operation year of SUSIE, we have defined a reference model and implemented a charging and accounting architecture based on the reference model. The architecture allows its components to be configured for meeting charging and accounting requirements of IntServ or DiffServ in unicast or multicast scenarios. A data structure, the PIP-NAR, has been defined for the implementation. This data structure is used for transporting usage information from an accounting processing entity to a charging entity, and allows exchange of accounting information in multi-provider scenarios.

Future work is planned for extending the RTFM-conformant metering components, and for deploying and testing the architecture within a multi-provider environment. For this purpose, the PIP-NAR data structure will be used by the accounting processing layer to exchange usage information among different providers.

REFERENCES

- [BrMR98] N. Brownlee, Mills, G. Ruth: Traffic Flow Measurement: Architecture. Work in progress, IETF Internet draft <draft-ietf-rtfm-architecture-04.txt>, September 1998
- [Brow97] N. Brownlee: Reference Manual NeTraMet & NeMaC Version 4.1, Information Technology Systems & Services, The University of Auckland, New Zealand, November 1997 (<http://www.auckland.ac.nz/net/Accounting/ntm.Release.note.html>)
- [Brow98a] N. Brownlee. SRL: A Language for Describing Traffic Flows and Specifying Actions for Flow Groups. Work in progress, IETF Internet draft <draft-ietf-rtfm-ruleset-language-03.txt>, September 1998
- [Brow98b] N. Brownlee. Traffic Flow Measurement: Meter MIB, Work in progress, IETF Internet Draft <draft-ietf-rtfm-meter-mib-06.txt>, September 1998.
- [Brow98c] <http://www.auckland.ac.nz/net/Accounting/ntm.Release.note.html>
- [CaSZ98] G. Carle, M. Smirnow, and T. Zseby. Charging and Accounting Architecture for IP Multicast Integrated Services. In Proc. of the 4th International Symposium on Interworking (Interworking'98), Ottawa, Canada, July 1998.
- [ETSI98] H. Orlamünder (Editor). Parameters and Mechanisms for Charging in IP based Networks [Network Aspects]. TR/NA-080301 V1.0.4 (1998-06), ETSI Working Group NA8 Technical Document, 1998
- [HaBR98] S.W. Handelman, Nevil Brownlee, Greg Ruth, S. Stibler: RTFM Working Group - New Attributes for Traffic Flow Measurement, Work in progress, IETF Internet draft <draft-ietf-rtfm-new-traffic-flow-06.txt>, October 3, 1998

- [Maic98] A. Maiocchi. NeTraMet & NeMaC for IIS Accounting: User's Guide, CEFRIEL, Politecnico di Milan, May 1998.
- [RFC1272] C. Mills, G. Hirsch, G. Ruth: „Internet Accounting Background”, RFC1272, November 1991
- [RFC2063] N. Brownlee, C. Mills, and G. Ruth. Traffic Flow Measurement: Architecture. IETF RFC2063, January 1997.
- [RFC2064] N. Brownlee: „Traffic Flow Measurement: Meter MIB”, RFC2064, University of Auckland, New Zealand, January 1997
- [RFC2123] N. Brownlee. Traffic Flow Measurement: Experiences with NeTraMet. IETF RFC2123, March 1997.