

Domain Based Metering

*Róbert Párhonyi*¹

*Bert-Jan van Beijnum*¹

¹Faculty of Computer Science, University of Twente

P.O. Box 217, 7500 AE Enschede, The Netherlands

E-mail: {parhonyi, beijnum}@cs.utwente.nl

ABSTRACT

In the last two years, there is a gaining interest in usage based accounting for the Internet. One of the driving forces for this growing interest is the progress being made in providing some form of quality assurance in IP packet forwarding. One of the key processes in an accounting system is metering, gathering of network usage parameters on which the accounting is or can be based.

In this paper are addressed issues that are potentially of interest in Internet accounting. One of these issues is the possibility to base accounting on traffic or flows that crosses the borders of a domain to different neighboring domains; that is, is it possible to meter these different flows? In this paper, the backgrounds and rationales of this, so-called, domain based metering are discussed and two preliminary experiments are reported.

1. INTRODUCTION

The Authentication, Authorization and Accounting working group (AAA WG) of the IETF is currently developing an accounting management architecture [1]. This architecture defines a set of network entities like network devices that collect accounting data, accounting servers and billing servers, and the kind of interactions that are foreseen between the devices. The actual protocols for these interactions are currently under debate. The network device collects resource consumption data in the form of accounting metrics. This information is then transferred to an accounting server. The processed accounting data is then submitted to a billing server, which typically handles rating and invoice generation, however, other functions can be realized using the same data, such as auditing, cost allocation, trend analysis and capacity planning functions.

The Realtime Traffic Flow Measurement working group (RTFM WG) of the IETF has between its objectives the production of an improved Traffic Flow Model and the development the RTFM Architecture and Meter MIB as 'standards track' documents with the IETF. The RTFM Architecture [2] provides a general method for metering network traffic flows. Flow definitions are very flexible and may be based on aggregates (for instance groups of source and destination IP addresses). The RTFM Architecture fits in the general Internet Management Architecture (SNMP). Metering information is stored in the so-called Meter-MIB [3].

In the Internet as it is used today, provides a best-effort service to its customers. For many applications this is sufficient. However, for high quality real-time applications like video on demand, or more generally a multi-media retrieval service, and shared cooperative working applications, the best-effort service is not sufficient. Having learned from QoS provisioning in ATM networks, the IETF community is working on Class of Service (CoS) provisioning, one of the technologies brought forward is known as Differentiated Services (DiffServ).

To create an incentive for users to contend for the CoS truly needed for there application, an accounting strategy that is based on the CoS used is reasonable. Furthermore, instead of basing accounting on flat rates per CoS it is more reasonable to base it on the actual traffic volume being transport by the network conform a specific class of service. It is with these arguments in mind that usage based accounting will be deployed in the CoS aware Internet.

One of the major issues in an accounting system is determining the set of parameters on which to base the accounting. The choices made here highly influence the metrics to be measured and the content of the session records to be produced. A first

classification of parameters, that seems to be useful for accounting, is:

- *Content/Service accounting*: In content and service accounting, parameters are considered that are relevant for services or the content that is delivered. Based on the business model and objectives used, an accounting and charging strategy can be applied.
- *Transport accounting*: ultimately, services are to be provided and content is to be delivered through the network: distance has to be bridged by the network. Usage of network resources and transport of packets is another category of parameters on which accounting may be based.

The starting point for accounting is collecting (raw) accounting data; this process is usually called metering. There are different methods, techniques and products for metering.

In the Telecom world there is a lot of experience in accounting. One of the features is the accounting based on the destination a user connects to. The Internet is organized differently; however, one of the questions is, if it is possible to account on basis of different destinations. A criterion may be the domain that traffic is sent to.

This research has been performed as part of the Internet Next Generation Project¹ and focuses, for now, on transport accounting. The scope of the paper is to present the results of preliminary experiments during the study of domain based metering. When neighboring domains are referred one may think about transient domains as well. For these experiments only strictly neighboring domains were considered, the metering of traffic received from/sent to transient domains is not subject of investigation.

In this paper is addressed the issue of differentiating accounting on basis of intra-domain and inter-domain traffic, and the feasibility to account on basis of different destination domains. In particular, was investigated whether an existing metering technique, the Meter-MIB, allows this domain based metering.

This paper is organized as follows. In Section 2 is presented an overview of various metering techniques. Domain Based Metering is explained and discussed in more detail in Section 3. In Section 4 two Domain Based Metering experiments and their results are described and discussed. Conclusions and further research are presented in Section 5.

¹ <http://ing.ctit.utwente.nl/>

2. METERING TECHNIQUES

Several vendors have developed and implemented proprietary metering tools usually incorporated in their routers.

Cisco has a metering product called NetFlow [4]. NetFlow technology efficiently provides the metering base for a key set of applications including network traffic accounting, usage-based network billing, network planning, network monitoring, and data mining capabilities for both service provider and enterprise customers. With NetFlow network traffic flow statistics can be collected and analyzed. It allows for flows to be defined on basis of different classes of service.

CableTron's Lightweight Flow Admission Control Protocol (LFAP) [5] provides Layer 3/4 flow accounting. The LFAP client forwards information on per-flow basis to the Spectrum Flow Admission Center (FAS). LFAP coupled with the FAS provides application level accounting services. Flows may also be defined on basis of different classes of service.

Both NetFlow and LFAP base their flows on a single source and destination IP address.

A publicly available implementation of the Meter-MIB is NeTraMet [6]. This release includes a Meter-Reader and Manager as well. The flows to be metered are defined by a so-called ruleset. Rulesets can be developed off-line and downloaded to the Meter-MIB by a manager. Metered data can be retrieved from the Meter-MIB by a Meter-Reader. The communications between these entities are using the SNMP protocol.

3. DOMAIN BASED METERING

The research questions related to transport accounting of the Internet Next Generation project are relevant in many ways, they help in understanding what is possible and what is not, as such the answers contribute in an accounting architecture and the feasibility of building accounting systems. Through this research and associated experiments, the aim is to answer these types of questions and contribute in developing an Internet accounting architecture. To list a few:

- What are the parameters to base transport accounting on: apart from what is theoretically feasible, there is also a large installed management base (MIBs).
- Transport accounting can be performed at different granularities, for instance: user/customer based, per DSCP or on aggregated level (e.g. in case of inter-domain accounting). One of the questions how to

instrument (meter) this, and are the existing metering techniques (standardized or proprietary) adequate to do this?

- How to collect and process transport accounting data (hence, protocol and architectural issues)?
- Nowadays a customer pays for the time he or she is connected to the Internet, or the amount of traffic he or she sends and receives. Is it possible to charge only the receiver for the traffic (we call this reversed charging)?
- Is it possible to base accounting on local and non-local traffic? Or, to bring this even one step further: is it possible to account differently for traffic that is sent to and received from different domains?

A simplified model of how the Internet is organized is as follows. There are the users or customers who, via an Internet Service Provider (ISP) have access to the Internet; these users may be either individuals or organization. The ISP has a network infrastructure, a *domain*, to provide Internet services. This network is connected to the networks of other ISPs. Interconnections of networks of different ISPs are the result of bilateral agreements between ISPs. Neighbors are the ISPs connected to one particular ISP. These bilateral agreements, usually referred to as Service Level Agreements (SLAs), specify for instance the amount of traffic (per time unit) that is maximally exchanged between the two domains. In the context of DiffServ networks, it may be foreseen that such capacity arrangements are made on an aggregated traffic level per Class of Service type. Also, depending on the expected balance between incoming and outgoing traffic, accounting arrangements may be part of the SLA.

Thus, for the ISP it makes a difference if traffic from and to its customers crosses the borders of the ISP's network (or domain) or not. As a result of this there exists an incentive for an ISP to account its users differently for local and non-local traffic.

Given an ISP, having several SLAs with other ISPs, it may have different accounting arrangements in each SLA. Hence, there may be a difference in terms of costs for traffic flowing to different other domains. Simply said, traffic over a link to or from one other domain might be more or less expensive than traffic over another link to another domain.

Due to that there is a motivation for an ISP to account its users differently for traffic to and from neighboring domains. Then the question is whether the traffic flowing to and from different neighboring

domains can be metered? This type of traffic or flow metering is called *domain based metering*.

There are different simple metering techniques to meter and record individual flows based on source and destination IP addresses (see section 2), but the challenge is to be able to summarize the aggregation of these individual flows between well-defined subdomains without any further processing.

Consider the following simple scenario: it is given an ISP (ISP A) which has between its end-customers several companies as customers (1, 2, 3), each company has its own IP address range, forming a subdomain within the ISPs domain (see Figure 1). ISP A has SLAs with the neighboring domains (ISP B, ISP C, and ISP D). Assume that data traffic between the customers of the provider A is free of charge, and the each link towards the neighboring ISPs has a different cost. In this case to charge its customers the ISP A must perform local and non-local metering, and inter-domain metering.




Figure 1: Scenario

Hence, in domain based metering it is not the intention to meter flows to and from a particular destination domain, because this would imply to have knowledge about the chain of SLAs to reach the destination. The intention of domain based metering is to meter flows to and from the ISP's neighboring domains on an aggregated level.

4. EXPERIMENTS WITH NETRAMET

To investigate the feasibility of local/non-local flow metering and domain based metering, two experiments have been carried out. The metering tool selected for these experiments was NeTraMet.

4.1. About NeTraMet

The Meter-MIB [3] is a sub-tree of the MIB-2, it consists of several groups that allow the management of the meter and the storage of flow data.

NeTraMet [6] is a public domain implementation of the Meter-MIB of the accounting architecture of the AAA WG. NeTraMet is a program which works as a meter, it stores the measured flow data in the Meter-MIB, and provides an SNMP agent to make it available to meter readers. It observes passing packets and builds flow data records for the flows of interest. It also can be used as well for real-time network monitoring and trouble-shooting. The performance and scalability of the metering tool [9] were not subjects of the investigations. NeMaC is included in the distribution of NeTraMet and it is used for remote configuration and management of meters and to read the collected flow data using the Simple Network Management Protocol (SNMP).

The header of each packet received by the meter is processed. All header fields are extracted from the header and a Match Key is created for the packet, then this key is matched against the current ruleset. The ruleset specifies if the packet is to be counted or ignored.

The Packet Matching Engine (PME) [2] does the matching process. The PME is a virtual machine that uses a set of instructions defined by the ruleset. A ruleset is a sequence of low-level machine-based instructions to be executed by the PME.

To specify rulesets at a higher level of abstraction, a special language was defined. The Simple Ruleset Language (SRL) [7] is a procedural language for creating rulesets. SRL programs are compiled into rulesets that can then be downloaded to meters.

The network manager configures the Meter-MIB. The involved management and configuration actions are creating a ruleset, downloading it to the Meter-MIB and activating the ruleset.

Flows are bi-directional and they can be specified at many different levels of aggregation. In general, flows are defined between endpoints: a source endpoint and destination endpoint. These endpoints can be specified at virtually any desired layer, e.g. MAC addresses, IP-addresses or Port-Numbers (see also [8]). In this study only IP addresses were considered. One single endpoint may consist of a single IP-address or a set of IP-addresses; the latter allows the definition of aggregated flows. If needed, a flow may be specified on basis of a single endpoint, in practice this means that the second endpoint is the entire (IP) address space.

Apart from address attributes to define a flow, there are many other flow attributes associated, for instance: times of first and last packets, counts for packets and bytes in each direction, Differentiated Services Code Point (DSCP).

4.2. The Experimental Network

In the Figure 2 the network configuration for the domain based metering experiments is given. It consists of four domains (D1, D2, D3, and D4), and each domain has its address mask to identify the network devices incorporated. A router interconnects these four domains.




Figure 2: The network configuration

The subnet masks for the domains are:

- **D1:** **130.89.17.64/29** and there are 3 hosts in this domain (130.89.17.67, 130.89.17.68 and 130.89.17.69);
- **D2:** **130.89.17.72/29** and there are 2 hosts in this domain (130.89.17.74, 130.89.17.75);
- **D3:** **130.89.17.80/29** and there are 2 hosts in this domain (130.89.17.82, 130.89.17.83);
- **D4:** **130.89.17.88/29** and there are 2 hosts in this domain (130.89.17.90, 130.89.17.91).

Using the network setup presented in the Figure 2 two kind of metering experiments were performed using NeTraMet. During the experiments no use has been made NeTraMet's capability to specify flows on basis of DS-Code Points, this is scheduled for future experiments. The goal in the current experiment was to investigate whether domain based metering is feasible at all.

In the first experiment two flows are considered: one flows stays within the local domain and the second flow has the local domain as a source endpoint and all other domains as destination endpoint.

In the second experiment another scenario is considered. Suppose that there are bilateral agreements between ISPs concerning the accounting of traffic, metering is necessary to determine the amount of data transmitted from one domain to each of the other domains. If there are n neighboring domains, then the metering process should result in $n-1$ inter-domain flows and an intra-domain flow.

4.3. Local and Non-Local Flows

Using the network topology presented in Figure 2 consider D1 the local domain, and D2 + D3 + D4 the non-local domain.

An SRL program was written for the classification and aggregation of local and non-local flows. There were defined two sets of address masks: one for the local domain (130.89.17.74/29) and another describing the non-local domain (130.89.17.72/29, 130.89.17.80/29, and 130.89.17.88/29). The meter and its manager were running on different hosts in the D1 domain.

The source code of the ruleset used in this experiment is the following:

```

DEFINE local = 130.89.17.64/29;
DEFINE nlocal = (130.89.17.72/29,
130.89.17.80/29, 130.89.17.88/29);
IF (SourcePeerAddress == local &&
    DestPeerAddress == local) {
    STORE FlowKind := 'L';
    COUNT;
} ELSE
IF (SourcePeerAddress == local&&
    DestPeerAddress == nlocal) {
    STORE FlowKind := 'N';
    COUNT;
} ELSE IGNORE;
SET 3;
FORMAT ToPDUs " " ToOctets " "
FlowKind;

```

NeTraMet's manager, NeMaC, was used to download the compiled ruleset to the meter, to retrieve the flow data from the Meter MIB and to present it in an easy readable format. The output file has three columns: number of transmitted PDUs and bytes, and the FlowKind value of the flow. The measured flow data can be viewed in Table 1.

Traffic	Transmitted PDUs	Transmitted bytes
Time: 13:39:30 Fri 19 May 2000		
Local	64	7498
Non-local	100	9800
Time: 13:40:00 Fri 19 May 2000		
Local	136	16135
Non-Local	222	21756
Time: 13:45:30 Fri 19 May 2000		
Local	858	99696
Non-Local	1864	182672

Table 1: Local and non-local traffic flow meter data

The ruleset constructed and the metering result show that it is possible for this network configuration to define local and non-local traffic flows, and therefore an accounting scheme based on different ratings for local and non-local traffic is feasible. Extending the scope further: an Internet Service Provider normally has a block of assigned IP-addresses, therefore the ISP can easily configure meters to allow for a local vs. non-local accounting scheme. Also, the ruleset may be refined to differentiate between local and non-local traffic for individual users. In this case two flows are to be specified for each user. Because meters can be placed close, for instance to access routers (hence, close to where users access the network), it is possible to monitor all traffic for each user at a single place.

4.4. Domain Based Metering Experiment

For the second metering experiment the question to answer is: is it possible to differentiate flows on basis of different neighboring domains?

For this experiment the same network configuration is used (see Figure 2). Each domain can be considered an ISP, and suppose that different SLAs between these ISPs exists. Through metering these SLAs can be checked, and accounting scheme may be based on this metered data as well. Domain D1 is considered the local domain, all other domains are non-local domains.

The SRL program specifying the ruleset could be as follows:

```

DEFINE D1 = 130.89.17.64/29;
DEFINE D2 = 130.89.17.72/29;
DEFINE D3 = 130.89.17.80/29;
DEFINE D4 = 130.89.17.88/29;

```

```

IF (SourcePeerAddress == D1 &&
    DestPeerAddress == D1) {
    SAVE SourcePeerAddress/29;
    SAVE DestPeerAddress/29;
    STORE FlowKind := '1'; #49
    COUNT;
} ELSE
IF (SourcePeerAddress == D1 &&
    DestPeerAddress == D2) {
    SAVE SourcePeerAddress/29;
    SAVE DestPeerAddress/29;
    STORE FlowKind := '2'; #50
    COUNT;
} ELSE
IF (SourcePeerAddress == D1 &&
    DestPeerAddress == D3) {
    SAVE SourcePeerAddress/29;
    SAVE DestPeerAddress/29;
    STORE FlowKind := '3'; #51
    COUNT;
} ELSE
IF (SourcePeerAddress == D1 &&
    DestPeerAddress == D4) {
    SAVE SourcePeerAddress/29;
    SAVE DestPeerAddress/29;
    STORE FlowKind := '4'; #52
    COUNT;
} ELSE
    NOMATCH;
ELSE
    IGNORE;
SET 4;
FORMAT SourcePeerAddress " "
    DestPeerAddress " "
    ToPDUs " " ToOctets " "
    FromPDUs " " FromOctets " "
    FlowKind;

```

The measured flow data is presented in the Table 2. NeMaC read the Meter MIB and presented it in a readable format. The “To ...” columns represent the amount of PDUs and bytes that were sent from the D1 domain, and the “From ...” columns represent the received amount of PDUs and bytes in the domain D1. For local traffic there is only “To ...” data because both source- and destination addresses are in the same domain, so the sent and received number of packets within D1 are the same.

In case that in domain D1 an accounting policy is applied in which each individual user is charged for traffic to and from the other domains, the ruleset can easily be adjusted to reach this higher level of granularity. The network configuration used is very simple, only ‘end domains’ were considered. In case one of the neighboring domains is a transient domain, the destination endpoint might become more complicated. In fact, depending on the IP numbering it might even be the case that the flow intended to meter cannot be captured by a single flow. Another issue that might complicate this even further if there are two or more transient domains and within D1 some kind of dynamic load balancing is applied. In this case there are no static routes to the destination.

A possible solution can possibly be found with a proper selection of the metering points. Issues like these are for further investigation.

Traffic	To PDUs	To bytes	From PDUs	From bytes
Time: 12:06:00 Fri 19 May 2000				
Local (D1)	27	7394	0	0
D1 → D2	4	192	34	8976
D1 → D3	12	1176	22	2342
D1 → D4	59	16335	50	14186
Time: 12:06:30 Fri 19 May 2000				
Local (D1)	93	15191	0	0
D1 → D2	64	6272	260	25480
D1 → D3	72	7056	68	19246
D1 → D4	136	19135	100	18500
Time: 12:12:00 Fri 19 May 2000				
Local (D1)	1287	177783	0	0
D1 → D2	1234	120932	32	3136
D1 → D3	1240	121520	505	76401
D1 → D4	823	118041	1300	127400

Table 2. Domain based traffic flow meter data

5. CONCLUSIONS AND FUTURE WORK

In this paper were discussed a number of accounting issues, the focus of the work is on transport accounting. A set of accounting issues relevant for a CoS aware Internet, in particular a DiffServ network, have been identified. The idea and rationale of Domain Based Metering has been discussed in detail and its feasibility have been studied by two preliminary experiments.

It can be concluded that for each ISP a domain and user based metering scheme based on local and non-local traffic is possible, also in the case when further flow differentiation on basis of DS Code Points is implemented. The main issue that remains to be considered here is that the used accounting architecture must scale in terms of number of users. Solutions for the scalability problem can be found in the area of distributed management architectures, a scalable solution is feasible with, for instance, the Script-MIB [9].

The metering experiment, in which different flows are considered on basis of different neighboring

domains, shows that this is feasible in some cases. However, further study and experimentation is needed to be conclusive about this. In case the neighboring domains are 'end-domains', therefore these domains are not transient domains, the problem can be solved rather easily as the experiments demonstrate. In case neighboring domains are transient domains, the definition of end-points of flows may become quite difficult.

In future experiments flows with different DS Code Points are going to be taken into account. Further study in depth and experimentation on domain based metering are foreseen. Issues to be addressed are, for instance, transient domains, and routing policies with a domain. One of the main criteria for the feasibility is the complexity of the rulesets needed to allow the intended metering.

REFERENCES

- [1] B. Aboba, J. Arkko, D. Harrington, "*Introduction to Accounting Management*", May 2000 (Internet draft draft-ietf-aaa-acct-03.txt)
- [2] N. Brownlee, C. Mills, G. Ruth, "*Traffic Flow Measurement: Architecture*", October 1999 (RFC 2722)
- [3] N. Brownlee, "*Traffic Flow Measurement: METER MIB*", October 1999 (RFC 2720)
- [4] Overview of the NetFlow FlowAnalyzer, Cisco's on-line documentation
<http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/nfa/>
- [5] P. Amsden, J. Amweg, P. Calato, S. Bensley, G. Lyons, "*Cabletron's Light-weight Flow Admission Protocol Specification v.1.0*", March 1997 (RFC 2124)
- [6] N. Brownlee, "*NeTraMet & NeMaC v. 4.3 reference manual*", June 1999
<http://www.auckland.ac.nz/net/Accounting/ntm.Release.note.html>
- [7] N. Brownlee, "*SRL: A Language for Describing Traffic Flows and Specifying Actions for Flow Groups*", October 1999 (RFC 2723)
- [8] S. Löffler, "Using Flows for Analysis and Measurement of Internet Traffic", August 1997
- [9] R. Marsman, "Development of Prototype-scripts for the Script-MIB", MSc. Thesis, University of Twente, the Netherlands, May 2000.
<http://wwwsnmp.cs.utwente.nl/nm/education/assignments/>