N. Brownlee
The University of Auckland
C. Mills
BBN Systems and Technologies
G. Ruth
GTE Laboratories, Inc.

January 1997

Traffic Flow Measurement: Architecture

Status of this Memo

This memo defines an Experimental Protocol for the Internet community. This memo does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Abstract

This document describes an architecture for the measurement and reporting of network traffic flows, discusses how this relates to an overall network traffic flow architecture, and describes how it can be used within the Internet. It is intended to provide a starting point for the Realtime Traffic Flow Measurement Working Group.

Table of Contents

1 Statement of Purpose and Scope 2
2 Traffic Flow Measurement Architecture 4
   2.1 Meters and Traffic Flows 4
   2.2 Interaction Between METER and METER READER 6
   2.3 Interaction Between MANAGER and METER 6
   2.4 Interaction Between MANAGER and METER READER 7
   2.5 Multiple METERs or METER READERs 7
   2.6 Interaction Between MANAGERs (MANAGER - MANAGER) 8
   2.7 METER READERs and APPLICATIONs 8
3 Traffic Flows and Reporting Granularity 9
   3.1 Flows and their Attributes 9
   3.2 Granularity of Flow Measurements 11
   3.3 Rolling Counters, Timestamps, Report-in-One-Bucket-Only . . 13
4 Meters 15
   4.1 Meter Structure 15
   4.2 Flow Table 17
   4.3 Packet Handling, Packet Matching 17
   4.4 Rules and Rule Sets 21
   4.5 Maintaining the Flow Table 24
   4.6 Handling Increasing Traffic Levels 25

5 Meter Readers 26
5.1 Identifying Flows in Flow Records 26
5.2 Usage Records, Flow Data Files 27
5.3 Meter to Meter Reader: Usage Record Transmission. 27
6 Managers 28
6.1 Between Manager and Meter: Control Functions 28
6.2 Between Manager and Meter Reader: Control Functions . . . 29
6.3 Exception Conditions 31
6.4 Standard Rule Sets 32
7 APPENDICES 33
7.1 Appendix A: Network Characterisation 33
7.2 Appendix B: Recommended Traffic Flow Measurement Capabilities 34
7.3 Appendix C: List of Defined Flow Attributes 35
7.4 Appendix D: List of Meter Control Variables 36
8 Acknowledgments 36
9 References 37
10 Security Considerations 37
11 Authors' Addresses 37

1 Statement of Purpose and Scope

This document describes an architecture for traffic flow measurement and reporting for data networks which has the following characteristics:

- The traffic flow model can be consistently applied to any protocol/application at any network layer (e.g. network, transport, application layers).

- Traffic flow attributes are defined in such a way that they are valid for multiple networking protocol stacks, and that traffic flow measurement implementations are useful in MULTI-PROTOCOL environments.

- Users may specify their traffic flow measurement requirements in a simple manner, allowing them to collect the flow data they need while ignoring other traffic.

- The data reduction effort to produce requested traffic flow information is placed as near as possible to the network measurement point. This reduces the volume of data to be obtained (and transmitted across the network for storage), and minimises the amount of processing required in traffic flow analysis applications.

The architecture specifies common metrics for measuring traffic flows. By using the same metrics, traffic flow data can be exchanged and compared across multiple platforms. Such data is useful for:

- Understanding the behaviour of existing networks,

- Planning for network development and expansion,

- Quantification of network performance,

- Verifying the quality of network service, and

- Attribution of network usage to users.

The traffic flow measurement architecture is deliberately structured so that specific protocol implementations may extend coverage to multi-protocol environments and to other protocol layers, such as usage measurement for application-level services. Use of the same model for both network- and application-level measurement may simplify the development of generic analysis applications which process and/or correlate any or all levels of traffic and usage information. Within this docuemt the term 'usage data' is used as a generic term for the data obtained using the traffic flow measurement architecture.

This document is not a protocol specification. It specifies and structures the information that a traffic flow measurement system needs to collect, describes requirements that such a system must meet, and outlines tradeoffs which may be made by an implementor.

For performance reasons, it may be desirable to use traffic information gathered through traffic flow measurement in lieu of network statistics obtained in other ways. Although the quantification of network performance is not the primary purpose of this architecture, the measured traffic flow data may be used as an indication of network performance.

A cost recovery structure decides "who pays for what." The major issue here is how to construct a tariff (who gets billed, how much, for which things, based on what information, etc). Tariff issues include fairness, predictability (how well can subscribers forecast their network charges), practicality (of gathering the data and administering the tariff), incentives (e.g. encouraging off-peak use), and cost recovery goals (100% recovery, subsidisation, profit making). Issues such as these are not covered here.

Background information explaining why this approach was selected is provided by 'Traffic Flow Measurement: Background' RFC [1].

2 Traffic Flow Measurement Architecture

A traffic flow measurement system is used by network Operations personnel for managing and developing a network. It provides a tool for measuring and understanding the network's traffic flows. This information is useful for many purposes, as mentioned in section 1 (above).

The following sections outline a model for traffic flow measurement, which draws from working drafts of the OSI accounting model [2]. Future extensions are anticipated as the model is refined to address additional protocol layers.

2.1 Meters and Traffic Flows

At the heart of the traffic measurement model are network entities called traffic METERS. Meters count certain attributes (such as numbers of packets and bytes) and classify them as belonging to ACCOUNTABLE ENTITIES using other attributes (such as source and destination addresses). An accountable entity is someone who (or something which) is responsible for some activitiy on the network. It may be a user, a host system, a network, a group of networks, etc, depending on the granularity specified by the meter's configuration.

We assume that routers or traffic monitors throughout a network are instrumented with meters to measure traffic. Issues surrounding the choice of meter placement are discussed in the 'Traffic Flow Measurement: Background' RFC [1]. An important aspect of meters is that they provide a way of succinctly aggregating entity usage information.

For the purpose of traffic flow measurement we define the concept of a TRAFFIC FLOW, which is an artificial logical equivalent to a call or connection. A flow is a portion of traffic, delimited by a start and stop time, that was generated by a particular accountable entity. Attribute values (source/destination addresses, packet counts, byte counts, etc.) associated with a flow are aggregate quantities reflecting events which take place in the DURATION between the start and stop times. The start time of a flow is fixed for a given flow; the end time may increase with the age of the flow.

For connectionless network protocols such as IP there is by definition no way to tell whether a packet with a particular source/destination combination is part of a stream of packets or not - each packet is completely independent. A traffic meter has, as part of its configuration, a set of 'rules' which specify the flows of interest, in terms of the values of their attributes. It derives attribute values from each observed packet, and uses these to decide which flow they belong to. Classifying packets into 'flows' in this way provides an economical and practical way to measure network traffic and ascribe it to accountable entities.

Usage information which is not deriveable from traffic flows may also be of interest. For example, an application may wish to record accesses to various different information resources or a host may wish to record the username (subscriber id) for a particular network session. Provision is made in the traffic flow architecture to do this. In the future the measurement model will be extended to gather such information from applications and hosts so as to provide values for higher-layer flow attributes.

As well as FLOWS and METERS, the traffic flow measurement model includes MANAGERS, METER READERS and ANALYSIS APPLICAIONS, which are explained in following sections. The relationships between them are shown by the diagram below. Numbers on the diagram refer to sections in this document.

```
                       MANAGER
                      /       \
                 2.3 /         \ 2.4
                    /           \
                   /             \                ANALYSIS
              METER   <----->   METER READER  <----->   APPLICATION
                         2.2                     2.7
```

- MANAGER: A traffic measurement manager is an application which configures 'meter' entities and controls 'meter reader' entities. It uses the data requirements of analysis applications to determine the appropriate configurations for each meter, and the proper operation of each meter reader. It may well be convenient to combine the functions of meter reader and manager within a single network entity.

- METER: Meters are placed at measurement points determined by network Operations personnel. Each meter selectively records network activity as directed by its configuration settings. It can also aggregate, transform and further process the recorded activity before the data is stored. The processed and stored results are called the 'usage data.'

- METER READER: A meter reader reliably transports usage data from meters so that it is available to analysis applications.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.