

## Accounting Requirements for IPng

### Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Abstract

This document was submitted to the IETF IPng area in response to RFC 1550. Publication of this document does not imply acceptance by the IPng area of any ideas expressed within. Comments should be submitted to the [big-internet@munari.oz.au](mailto:big-internet@munari.oz.au) mailing list.

### Summary

This white paper discusses accounting requirements for IPng. It recommends that all IPng packets carry accounting tags, which would vary in size. In the simplest case a tag would simply be a voucher identifying the party responsible for the packet. At other times tags should also carry other higher-level accounting information.

### Background

The Internet Accounting Model - described in RFC 1272 - specifies how accounting information is structured, and how it is collected for use by accounting applications. The model is very general, with accounting variables being defined for various layers of a protocol stack. The group's work has so far concentrated on the lower layers, but the model can be extended simply by defining the variables required, e.g., for session and application layers.

Brian Carpenter [1] suggests that IPng packets should carry authenticated (source, destination, transaction) triplets, which could be used for policy-based routing and accounting. The following sections explain how the transaction field - hereafter called an 'accounting tag' - could be used.

### Lower-layer (Transport) Accounting

At the lower (network) layers the tag would simply be a voucher. This means it is an arbitrary string which identifies the party

responsible, i.e., willing to pay for, a packet. It would initially be set by the host which originates the packet, hence at that stage the tag would identify the user who sent it.

A tag could be changed at various points along a packet's path. This could be done as part of the routing policy processing so as to reflect changes of the party responsible over each section of the path. For example:

user - provider	tag identifies user
provider A - provider B	tag identifies provider A

The tag could be used by accounting meters to identify the party responsible for a traffic flow, without having to deduce this using tables of rules. This should considerably simplify accounting for transit traffic across intermediate networks.

#### Higher-layer (Session and Application) Accounting

At higher layers there is a clear need to measure accounting variables and communicate them to various points along a packet's path, for example an application server may wish to inform a client about its usage of resources. A tag containing this information could be read by meters at any point along the packet's path for charging purposes, and could also be used by the client to inform the user of charges incurred.

It would make the collection of accounting data much simpler if this information was carried in a standard tag within each packet, rather than having different protocols provide this service in differing ways.

For 'old' applications which remain unaware of the tag field, a meter could be placed at a gateway for the application's host. This 'gateway' meter could determine what the application is by watching its streams of packets, then set an appropriate value in their tag fields.

#### Structure of the accounting tag

The two uses of tags outlined above must be able to coexist. Since many - indeed most - of the packets will only carry a voucher, it seems simplest to keep this as part of the routing tuple (see below).

For the application variables, a separate tag seems sensible. This would simply contain a list of the variables. Having two tags in this way would keep separate the management of routers and meters.

If the encryption/digital signature overhead of the second tag proves to be too high, it should be possible to combine this with the voucher.

The fine detail of this, or at least the way variables are packed into the tags, could be standardised by the Accounting Working Group in due course. For the purpose of IPng all that is required is the ability to carry one or two variable-size objects in every packet.

#### References

- [1] Carpenter, B., "IPng White Paper on Transition and Other Considerations", RFC 1671, CERN, August 1994.

#### Security Considerations

For IPng to provide reliable transport in a hostile environment, routing and accounting information, i.e., the (source, dest, network-tag) and (application-tag) tuples, must be tamper-proof. Routers and meters which need to use the tuples will need to hold appropriate keys for them. Network operators will have to plan for this, for example by determining which routers need which sets of keys. This will be necessary in any case for reliable policy-based routing, so the extra work required to set up accounting meters should be acceptable.

#### Author's Address

Nevil Brownlee  
Deputy Director  
Computer Centre, The University of Auckland  
Private Bag 92019, Auckland, New Zealand

Phone: +64 9 373 7599  
Fax: +64 9 373 7425  
EMail: n.brownlee@auckland.ac.nz

