

Network Working Group
Request for Comments: 1272

C. Mills
 BBN
D. Hirsh
Meridian Technology Corporation
 G. Ruth
 BBN
November 1991

INTERNET ACCOUNTING: BACKGROUND

Status of this Memo

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

1. Statement of Purpose

This document provides background information for the "Internet Accounting Architecture" and is the first of a three document set:

Internet Accounting Background & Status	(this document)
Internet Accounting Architecture	(under construction)
Internet Accounting Meter Service	(under construction)

The focus at this time is on defining METER SERVICES and USAGE REPORTING which provide basic semantics for measuring network utilization, a syntax, and a data reporting protocol. The intent is to produce a set of standards that is of practical use for early experimentation with usage reporting as an internet accounting mechanism.

The architecture should be expandable as additional experience is gained. The short-term Internet Accounting solution is intended to merge with OSI and Autonomous Network Research Group (ANRG) efforts and be superseded by those efforts in the long term. The OSI accounting working groups are currently defining meter syntax and reporting protocols. The ANRG research group is currently researching economic models and accounting tools for the Internet environment.

Internet Accounting as described here does not wrestle with the applications of usage reporting, such as monitoring and enforcing network policy; nor does it recommend approaches to billing or tackle such thorny issues as who pays for packet retransmission.

This document provides background and tutorial information on issues

surrounding the architecture, or in a sense, an explanation of choices made in the Internet Accounting Architecture.

2. Goals for a Usage Reporting Architecture

We have adopted the accounting framework and terminology used by OSI (ISO 7498-4 OSI Reference Model Part 4: Management Framework). This framework defines a generalized accounting management activity which includes calculations, usage reporting to users and providers and enforcing various limits on the use of resources. Our own ambitions are considerably more modest in that we are defining an architecture to be used over the short-term (until ISO and ANRG have final pronouncement and standards) that is limited to network USAGE REPORTING.

The OSI accounting model defines three basic entities:

- 1) the METER, which performs measurements and aggregates the results of those measurements;
- 2) the COLLECTOR, which is responsible for the integrity and security of METER data in short-term storage and transit; and
- 3) the APPLICATION, which processes/formats/stores METER data. APPLICATIONS implicitly manage METERS.

This working group, then, is concerned with specifying the attributes of METERS and COLLECTORS, with little concern at this time for APPLICATIONS.

3. The Usage Reporting Function

3.1. Motivation for Usage Reporting

The dominant motivations for usage reporting are:

- o Understanding/Influencing Behavior.
Usage reporting provides feedback for the subscriber on his use of network resources. The subscriber can better understand his network behavior and measure the impact of modifications made to improve performance or reduce costs.
- o Measuring Policy Compliance.
From the perspective of the network provider, usage reports might show whether or not a subscriber is in compliance with the stated policies for quantity of

network usage. Reporting alone is not sufficient to enforce compliance with policies, but reports can indicate whether it is necessary to develop additional methods of enforcement.

- o Rational Cost Allocation/Recovery.
Economic discipline can be used to penalize inefficient network configuration/utilization as well as to reward the efficient. It can be used to encourage bulk transfer at off hours. It can be used as a means to allocate operating costs in a zero-sum budget, and even be used as the basis for billing in a profit-making fee-for-service operation.

The chief deterrent to usage reporting is the cost of measuring usage, which includes:

- o Reporting/collection overhead.
This offers an additional source of computational load and network traffic due to the counting operations, managing the reporting system, collecting the reported data, and storing the resulting counts. Overhead increases with the accuracy and reliability of the accounting data.
- o Post-processing overhead.
Resources are required to maintain the post-processing tasks of maintaining the accounting database, generating reports, and, if appropriate, distributing bills, collecting revenue, servicing subscribers.
- o Security overhead.
The use of security mechanisms will increase the overall cost of accounting. Since accounting collects detailed information about subscriber behavior on the network and since these counts may also represent a flow of money, it is necessary to have mechanisms to protect accounting information from unauthorized disclosure or manipulation.

The balance between cost and benefit is regulated by the GRANULARITY of accounting information collected. This balance is policy-dependent. To minimize costs and maximize benefit, accounting detail is limited to the minimum amount to provide the necessary information for the research and implementation of a particular policy.

3.2. Network Policy and Usage Reporting

Accounting requirements are driven by policy. Conversely, policy is typically influenced by the available management/reporting tools and their cost. This section is NOT a recommendation for billing practices, but intended to provide additional background for understanding the problems involved in implementing a simple, adequate usage reporting system.

Since there are few tools adequate for any form of cost recovery and/or long-term monitoring there are few organizations that practice proactive usage reporting in the Internet. Those that do have generally invented their own. But far and away the most common approach is to treat the cost of network operations as overhead with network reports limited to short-term, diagnostic intervention. But as the population and use of the Internet increases and diversifies, the complexity of paying for that usage also increases. Subsidies and funding mechanisms appropriate to non-profit organizations often restrict commercial use or require that "for profit" use be identified and billed separately from the non-profit use. Tax regulations may require verification of network connection or usage. Some portions of the Internet are distinctly "private", whereas other Internet segments are treated as public, shared infrastructure.

The number of administrations operating in some connection with the Internet is exploding. The network "hierarchy" (backbone, regional, enterprise, stub network) is becoming deeper (more levels), increasingly enmeshed (more cross-connections) and more diversified (different charters and usage patterns). Each of these administrations has different policies and by-laws about who may use an individual network, who pays for it, and how the payment is determined. Also, each administration balances the OVERHEAD costs of accounting (metering, reporting, billing, collecting) against the benefits of identifying usage and allocating costs.

Some members of the Internet community are concerned that the introduction of usage reporting will encourage new billing policies which are detrimental to the current Internet infrastructure (though it is also reasonable to assert that the current lack of usage reporting may be detrimental as well). Caution and experimentation must be the watch words as usage reporting is introduced. Well before meters are used for active BILLING and ENFORCEMENT, they should first be used to:

- o UNDERSTAND USER BEHAVIOR
(learn to quantify and/or predict individual and aggregate traffic patterns over the long term),

- o QUANTIFY NETWORK IMPROVEMENTS,
(measure user and vendor efficiency in how network resources are consumed to provide end-user data transport service) and
- o MEASURE COMPLIANCE WITH POLICY.

Accounting policies for network traffic already exist. But they are usually based on network parameters which change seldom, if at all. Such parameters require little monitoring (the line speed of a physical connection, e.g., Ethernet, 9600 baud, FDDI). The connection to the network is then charged to the subscriber as a FLAT-FEE regardless of the amount of traffic passed across the connection and is similar to the monthly unlimited local service phone bill.

Usage-insensitive access charges are sufficient in many cases, and can be preferable to usage-based charging in Internet environments, for financial, technical, and social reasons. Sample incentives for the FLAT-FEE billing approach are:

- o FINANCIAL:
Predictable monthly charges. No overhead costs for counting packets and preparing usage-based reports.
- o TECHNICAL:
Easing the sharing of resources. Eliminating the headaches of needing another layer of accounting in proxy servers which associate their usage with their clients'. Examples of proxy servers which generate network traffic on behalf of the actual user or subscriber are mail daemons, network file servers, and print spoolers.
- o SOCIAL:
Treating the network as an unregulated public infrastructure with equal access and information sharing. Encouraging public-spirited behavior -- contributing to public mailing lists, information distribution, etc.

In other cases USAGE-SENSITIVE charges may be preferred or required by a local administration's policy. Government regulations or the wishes of subscribers with low or intermittent traffic patterns may force the issue (note: FLAT FEES are beneficial for heavy network users. USAGE SENSITIVE charges generally benefit the low-volume user). Where usage-sensitive accounting is used, cost ceilings and floors may still be established by static parameters, such as "pipe size" for fixed connections or "connection time" for dial-up connection, to satisfy the need for some predictability.

Different billing schemes may be employed depending on network measures of distance. For example, local network traffic may be flat-rate and remote internet traffic may be usage-based, analogous to the local and long distance billing policies adopted by the telephone companies.

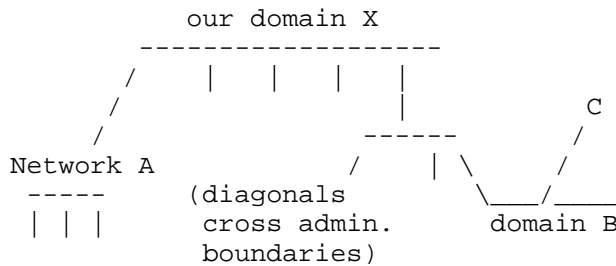
The ANRG is independently investigating policy models and infrastructure economics for billing and cost recovery.

3.3. The Nature of Usage Accounting

Although the exact requirements for internet usage accounting will vary from one network administration to the next and will depend on policies and cost trade-offs, it is possible to characterize the problem in some broad terms and thereby bound it. Rather than try to solve the problem in exhaustive generality (providing for every imaginable set of accounting requirements), some assumptions about usage accounting are posited in order to make the problem tractable and to render implementations feasible. Since these assumptions form the basis for our architectural and design work, it is important to make them explicit from the outset and hold them up to the scrutiny of the Internet community.

3.3.1. A Model for Internet Accounting

We begin with the assumption that there is a "network administrator" or "network administration" to whom internet accounting is of interest. He "owns" and operates some subset of the internet (one or more connected networks) that may be called his "administrative domain". This administrative domain has well defined boundaries.



The network administrator is interested in 1) traffic within his boundaries and 2) traffic crossing his boundaries. Within his boundaries he may be interested in end-system to end-system accounting or accounting at coarser granularities (e.g., department to department).

The network administrator is usually not interested in accounting for end-systems outside his administrative domain; his primary concern is accounting to the level of other ADJACENT (directly connected) administrative domains. Consider the viewpoint of the administrator for domain X of the internet. The idea is that he will send each adjacent administrative domain a bill (or other statement of accounting) for its use of his resources and it will send him a bill for his use of its resources. When he receives an aggregate bill from Network A, if he wishes to allocate the charges to end users or subsystems within his domain, it is HIS responsibility to collect accounting data about how they used the resources of Network A. If the "user" is in fact another administrative domain, B, (on whose behalf X was using A's resources) the administrator for X just sends his counterpart in B a bill for the part of X's bill attributable to B's usage. If B was passing traffic for C, then B bills C for the appropriate portion X's charges, and so on, until the charges percolate back to the original end user, say G. Thus, the administrator for X does not have to account for G's usage; he only has to account for the usage of the administrative domains directly adjacent to himself.

This paradigm of recursive accounting may, of course, be used WITHIN an administrative domain that is (logically) comprised of sub-administrative domains.

The discussion of the preceding paragraphs applies to a general mesh topology, in which any Internet constituent domain may act as a service provider for any connected domain. Although the Internet topology is in fact such a mesh, there is a general hierarchy to its structure and hierarchical routing (when implemented) will make it logically hierarchical as far as traffic flow is concerned. This logical hierarchy permits a simplification of the usage accounting perspective.

At the bottom of the service hierarchy a service-consuming host sits on one of many "stub" networks. These are interconnected into an enterprise-wide extended LAN, which in turn receives Internet service, typically from a single attachment to a regional backbone. Regional backbones receive national transport services from national backbones such as NSFnet, Altnet, PSInet, CERFnet, NSInet, or Nordunet. In this scheme each level in the hierarchy has a constituency, a group for which usage reporting is germane, in the level underneath it. In the case of the NSFnet the natural constituency, for accounting purposes at least, is the regional nets (MIDnet, SURAnet,...). For the regionals it will be their member institutions; for the institutions, their stub networks; and for the stubs, their individual hosts.

3.3.2. Implications of the Model

The significance of the model sketched above is that Internet accounting must be able to support accounting for adjacent (intermediate) systems, as well as end-system accounting. Adjacent system accounting information cannot be derived from end-system accounting (even if complete end-system accounting were feasible) because traffic from an end-system may reach the administrative domain of interest through different adjacent domains, and it is the adjacent domain through which it passes that is of interest.

The need to support accounting for adjacent intermediate systems means that internet accounting will require information not present in internet protocol headers (these headers contain source and destination addresses of end-systems only). This information may come from lower layer protocols (network or link layer) or from configuration information for boundary components (e.g., "what system is connected to port 5 of this IP router").

4. Meters

A METER is a process which examines a stream of packets on a communications medium or between a pair of media. The meter records aggregate counts of packets belonging to FLOWS between communicating entities (hosts/processes or aggregations of communicating hosts (domains)). The assignment of packets to flows may be done by executing a series of rules. Meters can reasonably be implemented in any of three environments -- dedicated monitors, in routers or in general-purpose systems.

Meter location is a critical decision in internet accounting. An important criterion for selecting meter location is cost, i.e., REDUCING ACCOUNTING OVERHEAD and MINIMIZING THE COST OF IMPLEMENTATION.

In the trade-off between overhead (cost of accounting) and detail, ACCURACY and RELIABILITY play a decisive role. Full accuracy and reliability for accounting purposes require that EVERY packet must be examined. However, if the requirement for accuracy and reliability is relaxed, statistical sampling may be more practical and sufficiently accurate, and DETAILED ACCOUNTING is not required at all. Accuracy and reliability requirements may be less stringent when the purpose of usage-reporting is solely to understand network behavior, for network design and performance tuning, or when usage reporting is used to approximate cost allocations to users as a percentage of total fees.

Overhead costs are minimized by accounting at the coarsest acceptable

GRANULARITY, i.e., using the greatest amount of AGGREGATION possible to limit the number of accounting records generated, their size, and the frequency with which they are transmitted across the network or otherwise stored.

The other cost factor lies in implementation. Implementation will necessitate the development and introduction of hardware and software components into the internet. It is important to design an architecture that tends to minimize the cost of these new components.

4.1. Meter Placement

In the model developed above, the Internet may be viewed as a hierarchical system of service providers and their corresponding constituencies. In this scheme the service provider accounts for the activity of the constituents or service consumers. Meters should be placed to allow for optimal data collection for the relevant constituency and technology. Meters are most needed at administrative boundaries and data collected such that service provider and consumer are able to reconcile their activities.

Routers (and/or bridges) are by definition and design placed (topologically) at these boundaries and so it follows that the most generally convenient place to position accounting meters is in or near the router. But again this depends on the underlying transport. Whenever the service-providing network is broadcast (e.g., bus-based), not extended (i.e., without bridging or routing), then meter placement is of no particular consequence. If one were generating usage reports for a stub LAN, meters could reasonably be placed in a router, a dedicated monitor, or a host at any point on the LAN. Where an enterprise-wide network is a LAN, the same observation holds. At the boundary between an enterprise and a regional network, however, in or near a router is an appropriate location for meters that will measure the enterprise's network activity.

Meters are placed in (or near) routers to count packets at the Internet Protocol Level. All traffic flows through two natural metering points: hosts and routers (Internet packet switches). Hosts are the ultimate source and sink of all traffic. Routers monitor all traffic which passes IN or OUT of each network. Motivations for selecting the routers as the metering points are:

- o Minimization of cost and overhead.
(by concentrating the accounting function). Centralize and minimize in terms of number of geographical or administrative regions, number of protocols monitored, and number of separate implementations modified. (Hosts are too diverse and numerous for easy standardization.

Routers concentrate traffic and are more homogeneous.)

- o Traffic control.
When and if usage sensitive quotas are involved, changes in meter status (e.g., exceeding a quota) would result in an active influence on network traffic (the router starts denying access). A passive measuring device cannot control network access in response to detecting state.
- o Intermediate system accounting.
As discussed above, internet accounting includes both end-system and intermediate system accounting. Hosts see only end-system traffic; routers see both the end-systems (internet source and destination) and the adjacent intermediate systems.

Therefore, meters should be placed at:

- o administrative boundaries
only for measuring inter-domain traffic;
- o stub networks
for measuring intra-domain traffic. For intra-domain traffic, the requirement for performing accounting at almost every router is a disincentive for implementing a usage-based charging policy.

4.2. Meter Types

Four possible types of metering technology are:

- o Network monitors:
These measure only traffic WITHIN a single network. They include LAN monitors, X.25 call accounting systems and traffic monitors in bridges.
- o Line monitors:
These count packets flowing across a circuit. They would be placed on inter-router trunks and on router ports.
- o Router-integral meters:
These are meters located within a router, implemented in software. They count packets flowing through the router.
- o Router spiders:
This is a set of line monitors that surround a router, measure traffic on all of its ports and coordinate the results.

4.3. Meter Structure

While topology argues in favor of meters in routers, granularity and security favor dedicated monitors. The GRANULARITY of the accountable entity (and its attributes) affects the amount of overhead incurred for accounting. Each entity/attribute/reporting interval combination is a separate meter. Each individual meter takes up local memory and requires additional memory or network resources when the meter reports to the application. Memory is a limited resource, and there are cost implications to expanding memory significantly or increasing the frequency of reporting. The number of concurrent flows open in a router is controlled by

- o the granularity of the accountable entity
- o the granularity of the attributes and sub-categories of packets
- o memory
(the number of flows that can be stored concurrently, a limit which can also be expressed as the average number of flows existing at this granularity plus some delta, e.g., peak hour average plus one standard deviation, or ...)
- o the reporting interval
(the lifetime of an individual meter)

There is a spectrum of granularity control which ranges across the following dimensions. (Most administrations will probably choose a granularity somewhere in the middle of the spectrum.)

ENTITY: Entities range across the spectrum from the coarsest granularity, PORT (a local view with a unique designation for the subscriber port through which packets enter and exit "my" network) through NETWORK and HOST to USER (not defined here). The port is the minimum granularity of accounting. HOST is the finest granularity defined here. Where verification is required, a network should be able to perform accounting at the granularity its subscribers use. Hosts are ultimately responsible for identifying the end user, since only the hosts have unambiguous access to user identification. This information could be shared with the network, but it is the host's responsibility to do so, and there is no mechanism in place at this time (e.g., an IP option, discussed in section 4.).

ATTRIBUTE: Each new attribute requires that an additional flow be maintained for each entity. The coarsest granularity is NO

categorization of packets. The finest granularity would be to maintain state information about the higher-levels protocols or type of service being used by communicating processes across the network.

VALUES: Values are the information which is recorded for each entity/attribute grouping. Usually values are counters, such as packet counts and byte counts. They may also be time stamps - start time and stop time, or reasons for starting or stopping reporting.

REPORTING INTERVAL: At the very finest level of granularity, each data packet might generate a separate accounting record. To report traffic at this level of detail would require approximately one packet of accounting information for every data packet sent. The reporting interval is then zero and no memory will be needed for flow record storage. For a non-zero reporting interval flow records must be maintained in memory. Storage for stale (old, infrequent) flows may be recycled when their data has been reported. As the reporting interval increases, more and more stale records accumulate.

The feasibility of a particular group of granularities varies with the PERFORMANCE characteristics of the network (link speed, link bandwidth, router processing speed, router memory), as well as the COST of accounting balanced against the requirement for DETAIL. Since technological advances can quickly obsolete current technical limitations, and since the policy structure and economics of the Internet are in flux, meters will be defined with VARYING GRANULARITY which is regulated according to the traffic requirements of the individual network or administration and technical limitations.

4.4. Collection Issues

There are two implicit assumptions about the nature of meters and traffic sources that they measure, both of which have substantial bearing on collectors.

1. The matrix of communicating entity pairs is large but sparse and, moreover, network traffic exhibits considerable source, destination and attribute coherence - so that lists can be quite compact.
2. Meters can be configured to generate either a static set of variables whose values are incremented, or a stream of records that must be periodically transferred and removed from the meter's memory.

Meters can generate large, unstructured amounts of information and the essential collection issue revolves around mapping collection activities into an SNMP framework (or, to the extent that this is not successful, specifying other collection paradigms).

There are three major collection concerns:

- o data confidentiality
- o data integrity
- o local and remote collection control

The prime security concern is preserving the confidentiality of usage data. (See ISO 7498 Part 2, "Security Architecture," for security terminology used herein.) Given that accounting data are sensitive, the collector should be able (or may be required) to provide confidentiality for accounting data at the point of collection, through transmission and up to the point where the data is delivered. The delivery function may also require authentication of the origin and destination and provision for connection integrity (if connections are utilized). Other security services (e.g., measures to counter denial of service attacks) are not deemed necessary for internet accounting at this time. It is assumed that security services can be provided by SNMP and its mechanisms. (This will require further investigation.)

In order to have an accurate monitoring system, reliable delivery of data should be assured through one or more of:

- o an acknowledgement retransmission scheme;
- o redundant reporting to multiple collectors;
- o having backup storage located at the meter.

There is a place for both application polling and meter traps within this scheme, but there are significant trade-offs associated with each.

Polling means that the collection point has some control over when accounting data is sent, so that not all meters flood the collector at once. However, polling messages, particularly when structured with SNMP's GET-NEXT operator, add considerable overhead to the network. Meter traps are required in any case (whether or not polling is the preferred collection method), so that a meter may rid itself of data when its cache is full.

The fundamental collection trade-off will be between primary and secondary storage at the meter, coupled with an efficient bulk-transfer protocol, versus minimal storage at the meter and a network-bandwidth-consuming collection discipline.

A final collection concern is whether packets should be counted on entry into a router or upon exit from a router. It is the nature of IP that not every packet received by a router is actually passed to an output port. The Internet Protocol allows routers to discard packets (e.g., in times of congestion when the router cannot handle the offered load); it is presumed that higher level protocols (e.g., TCP) will provide whatever reliable delivery service the user deems necessary (by detecting non-delivery and retransmitting).

The question arises, therefore, whether an internet accounting system should count all packets offered to a router (since each packet offered consumes some router resources) or just those that are finally passed by the router to a network (why should a user pay for undelivered packets?) Since there are good arguments for either position, we do not attempt to resolve this issue here. (It should be noted, however, that SMDS has chosen to count on exit only.) Rather, we require that an internet accounting should provide ability for counting packets either way -- on entry to or on exit from a router.

5. Examples


Here follows a series of examples to illustrate what data may be of interest to service providers and consumers in a number of different scenarios. In the illustrations that follow straight lines are interpreted as some sort of LAN. Diagonals are point-to-point links. Diamonds are routers. We assume that we are in a homogeneous protocol environment (IP).

5.1 A Single Segment LAN

Consumers and providers on a single LAN service can utilize the same set of data: the contribution of individual hosts to total network load. A network accounting system measures flows between individual host pairs. (On a broadcast LAN, e.g., an Ethernet, this can be accomplished by a single meter placed anywhere on the LAN.) Using this data, costs for the network management activity can be apportioned to individual hosts or the departments that own/manage the hosts.

Alternately, flows can be kept by source only, rather than source-destination pairs.

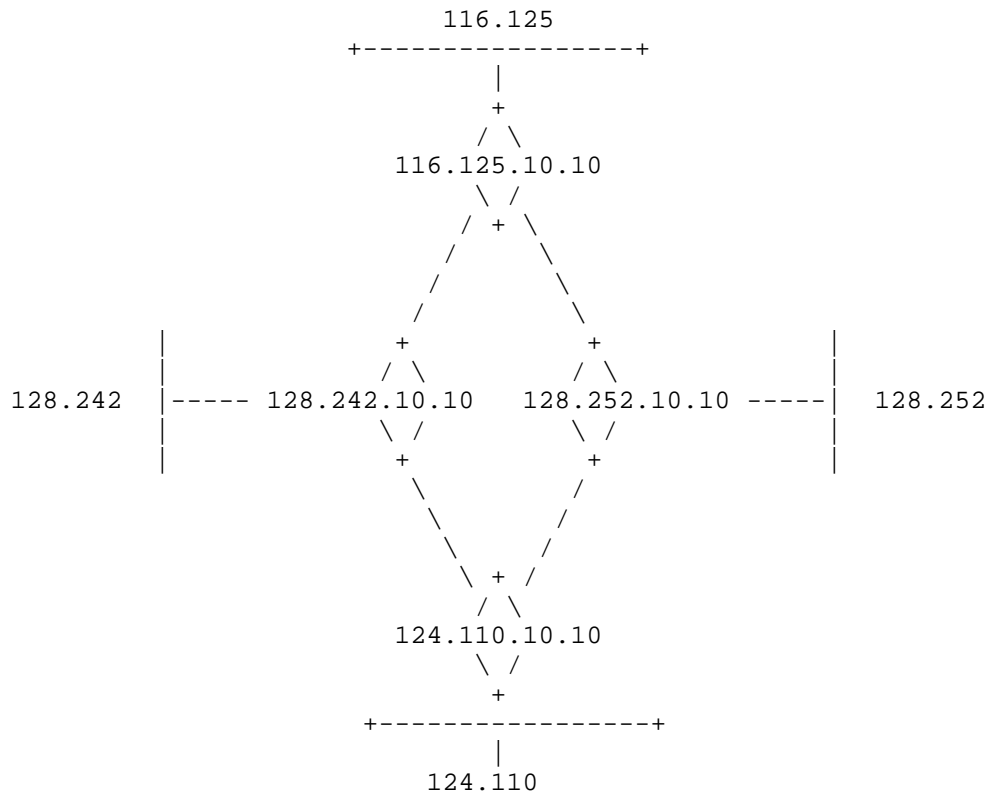
5.2 An Extended (Campus or Facility-Wide) LAN



This is the first example in which the information that is germane for service provider and consumer are not identical. The service consumers are now the individual subnets and the service provider is the facility-wide backbone. A service provider is interested in knowing the contribution of individual subnets to the total traffic of the backbone. In order to ascertain this, a meter on the backbone (the longest line in the center of the illustration) can keep track of flows between subnet pairs. Now the communications between individual hosts on adjacent subnets are aggregated into a single flow that measures activity between subnets.

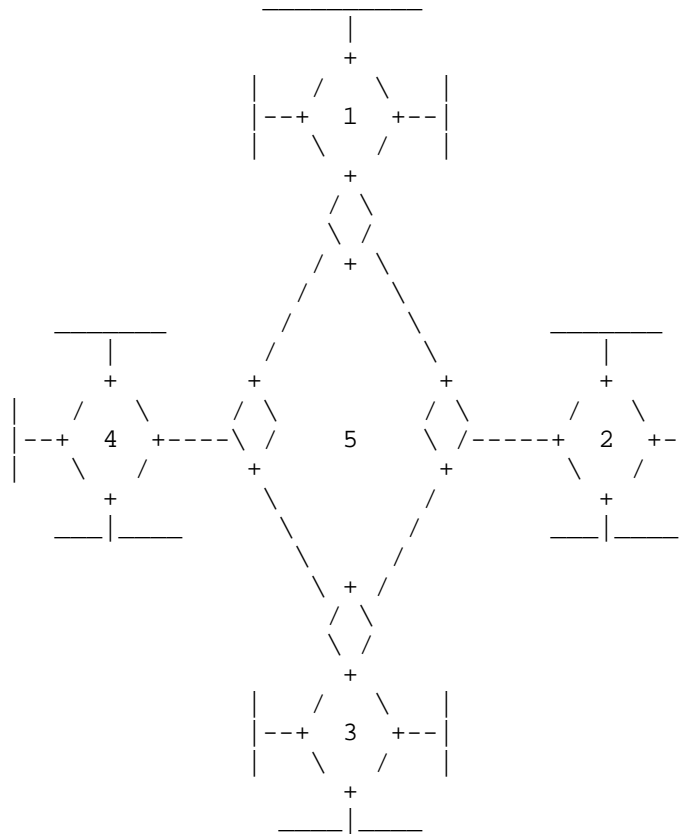
The service consumers, or subnets, might in turn want to keep track of the communications between individual hosts that use the services of the backbone. An accounting system on the backbone could be configured to monitor traffic among individual host pairs. Alternately an accounting system on each individual subnet could keep track of local and "non-local" traffic. The observed data of the two sets of meters (one for the service provider and one for the service consumers) should have reconcilable data.

5.3 A Regional Network



In this example we have a regional network consisting of a ring of point-to-point links that interconnect a collection of campus-wide LANs. Again service provider and consumer have differing interests and needs for accounting data. The service provider, the regional network, again will be interested in the contribution of each individual network to the total traffic on the regional network. This interest might extend to include measure of individual link utilization, and not just total offered load to the network as a whole. In this latter case the service provider will require that meters be placed at one end or the other on each link. For the service consumer, the individual campus, relevant measures would include the contribution of individual subnets or hosts to the total "outbound" traffic. Meter(s) placed in (or at) the router that connects the campus- network to the regional network can perform the necessary measurement.

5.4 A National Backbone



In this last case, the data that the service provider will want to collect is the traffic between regional networks. The flow that measures a regional network, or regional network pairs, is defined as the union of all member-campus network address spaces. This can be arrived at by keeping multiple individual network address flows and developing the regional network contribution as post-processing activity, or by defining a flow that is the union of all the relevant addresses. (This is a cpu cycles for memory trade-off.) Note that if the service provider measures individual network contributions, then this data is, in large measure, the data that the service consumers would require.

6. Future Issues

This last section is the collector for ancillary issues that are as yet undefined or out of current scope.

APPLICATIONS standards: Recommendations for storage, processing and reporting are left out for the moment. Storage and processing of accounting information is dependent on individual network policy. Recommendations for standardizing billing schemes would be premature.

QUOTAS are a form of closed loop feedback that represent an interesting extension of usage reporting. But they will have to wait until the basic accounting technology is reasonably defined and has been the subject of a reasonable amount of experimentation.

SESSION ACCOUNTING: Detailed auditing of individual sessions across the internet (at level four or higher) will not be addressed by internet accounting. Internet accounting deals only with measuring traffic at the IP level.

APPLICATION LEVEL ACCOUNTING: Service hosts and proxy agents have to do their own accounting for services, since the network cannot distinguish on whose behalf they are acting. Alternately, TCP/UDP port numbers could become an optional field in a meter, since the conjunction of a pair of IP addresses and port numbers occurring at a particular time uniquely identifies a pair of communicating processes.

The USER has not yet been defined, since an IP option would have to be added to the IP header to provide for this. This option would probably contain two parts - a subscriber identification and a user sub-identification - to allow for the later introduction of quota mechanisms which have both group and individual quotas. The subscriber is the fiscally responsible entity, for example the manager of a research group. In any case, routers must be able to fall back to accounting by host, since there will most certainly be hosts on the network which do not implement a new IP option in a timely fashion.

7. References

International Standards Organization (ISO), "Management Framework," Part 4 of Information Processing Systems Open Systems Interconnection Basic Reference Model, ISO 7498-4, 1984.

International Standards Organization (ISO), "Security Architecture," Part 2 of Information Processing Systems Open Systems Interconnection Basic Reference Model, ISO 7498-2, 1984.

Security Considerations

Security issues are discussed in sections 2, 3 and 4.

Authors' Addresses

Cyndi Mills
Bolt, Beranek, and Newman
150 Cambridge Park Drive
Cambridge, MA 02140

Phone: 617-873-4143
Email: cmills@bbn.com

Donald Hirsh
Meridian Technology Corporation
11 McBride Corporate Center Drive
Suite 250
Chesterfield, MO 63005

Phone: 314-532-7708
Email: hirsh@meridian.uucp

Gregory Ruth
Bolt, Beranek, and Newman
150 Cambridge Park Drive
Cambridge, MA 02140

Phone: 617-873-3150
Email: gruth@bbn.com