post-NSFNET statistics collection[1]

Hans-Werner Braun <hwb@sdsc.edu>
K. Claffy <kc@sdsc.edu>

Abstract

As the NSFNET backbone service migrates into a commercial environment, so will access to the only set of publicly available statistics for a large national U.S. backbone. The transition to the new NSFNET program, with commercially operated services providing both regional service as well as cross-service provider switching points, or Network Access Points (NAPs), will render statistics collection a much more difficult endeavor. In this paper we discuss issues and complexities of statistics collection at recently deployed global network access points such as the U.S. federal NAPs.

1 background

The United States National Science Foundation (NSF) is making a transition away from supporting Internet backbone services for the general research and education community. One component of these services, the NSFNET backbone, was dismantled as of the end of April 1995. To facilitate a smooth transition to a multi-provider environment, and hopefully forestall the potential for network partitioning in the process, the NSF is sponsoring several official Network Access Points (NAPs) and providing regional service providers with incentives to connect to all three points.[2] [3]

NAPs are envisioned to provide a neutral, AUP-free meeting point for network service providers to exchange routing and traffic. The three NSF-sponsored priority NAPs are:

- Sprint NAP, in Pennsauken, NJ
- Pacific Bell NAP, in Oakland/Palo Alto, CA
- Ameritech NAP, in Chicago, IL

NSF also sponsors a non-priority NAP, in Washington, DC, operated by Metropolitan Fiber Systems. The Sprint NAP was operational as of November 1994, but the other NAPs were not yet ready until mid-1995, mainly because Sprint was the only NAP who did not try to start off with switched ATM technology, but rather begin with an FDDI implementation. In addition, NSF is sponsoring a very high speed backbone service (vBNS), based on ATM technology, to support meritorious research requiring high bandwidth network resources. The vBNS represents a testbed for the emerging broadband Internet service infrastructure in which all parts of the network will be experimented with: switches, protocols, software, etc., as well as applications. It will be a unique resource for network and application researchers nationwide to explore performance issues with the new technologies (e.g., how host systems and interfaces interact with ATM components of the wide area network) [1, 2].

As the NSFNET era comes to a close, we will no longer be able to rely on what was the only set of publicly available statistics for a large national U.S. backbone. The transition to the new NSFNET program, with commercially operated services providing both regional service as well as cross-service provider switching points (NAPs), will render statistics collection a much more difficult task. There are several dimensions of the problem, each with a cost-benefit tradeoff. We examine them in turn.

2 dimensions of problem

2.1 contractual and logistical

In the cooperative agreement with the NAPs, the National Science Foundation has made a fairly vague request for statistics reporting. As with the NSFNET program, NSF was not in a position to specify in detail what statistics the network manager should collect since NSF did not know so much about the technology as the providers themselves did. The situation is similar with other emerging network service providers,

[1] This research is supported by a grant from the National Science Foundation (NCR-9119473). Reprinted with permission from National Academy of Sciences.

[2] See *http://www.merit.edu* for more information on the NSFNET transition.

[3] Specifically, NSF-sponsored regional providers, i.e., those who received funding from the NSF throughout the life of the NSFNET, will only continue to receive funding if they connect to all three NSF NAPs. Even so, the funding will gradually diminish within 4 years, an interval of time in which regional providers will have to become fully self-sustaining within the marketplace.

Figure 1: Merit notification of the cessation of NSFNET statistics collection

```
NIC.MERIT.EDU /nsfnet/statistics/READ.ME                                9 January 1995

NSFNET performance statistics have been collected, processed, stored, and reported by the Merit
Network since 1988, in the early stages of the NSFNET project.  During December 1994, the
numbers contained in Merit's statistical reports began to decrease, as NSFNET traffic began to
migrate to the new NSF network architecture.  In the new architecture, traffic is exchanged at
interconnection points called NAPs (Network Access Points.)  Each NAP provides a neutral
interconnection point for U.S.-based and international network service providers.
Once the new architecture is in place, Merit will be unable to collect the data needed to
continue these traffic-based reports.  The reports will be discontinued by spring 1995.
```

whose understanding of the technology and what statistics collection is possible is likely to exceed that of NSF. The NAPs and NSPs, however, at least in early 1995, were having enough of a challenge in getting and keeping their infrastructure operational; statistics have not been a top priority. Nor do the NAPs really have a good sense of what to collect, as all of the new technology involved is quite new to them as well.

A concern is that the NAPs will likely wait for more specific requirements from NSF, while NSF waits for them to develop models on their own. Scheduled meetings of community interest groups (e.g., NANOG, IEPG, FEPG, EOWG, Farnet)[4] that might develop statistics standards have hardly enough time for more critical items on the agenda, e.g., switch testing and instability of routing. The issue is not whether traffic analysis would help, even with equipment and routing problems, but that traffic analysis is perceived as a secondary issue, and there is no real mechanism (or spare time) for collaborative development of an acceptable model.

cost-benefit tradeoff: fulfill the deliverables of the cooperative agreement with NSF, at the least cost in terms of time and effort taken away from more critical engineering and customer service activities.

2.2 academic and fiscal

Many emerging Internet services are offered by companies whose primary business has thus far been telecommunications rather than IP. The NAP providers (as well as the vBNS provider) are good examples. Traditionally phone companies, they find themselves accustomed to having reasonable tools to model telephony workload and performance (e.g., Erlang distributions). Unfortunately, the literature in Internet traffic characterization, both in the analytical and performance measurement domains, indicate that wide area networking technology has advanced at a far faster rate than has the analytical and theoretical understanding of Internet traffic behavior.

The slower and more containable realms of years ago were amenable to characterization with closed-form mathematical expressions, which allowed reasonably accurate prediction of performance metrics such as queue lengths and network delays. But traditional mathematical modeling techniques, e.g., queueing theory, have met with little success in today's Internet environments. For example, the assumption of Poisson arrivals was acceptable for the purposes of characterizing small LANs years ago. As a theory of network behavior, however, the tenacity of the use of Poisson arrivals, whether in terms of packet arrivals within a connection[5], connection arrivals within an aggregated stream of traffic, or packet arrivals across multiple connections, has been quite remarkable in the face of its egregious inconsistency with any collected data [3, 4]. Leland, *et al.* [5] and Paxson and Floyd [6] investigate alternatives to Poisson modeling, specifically the use of self-similarity (fractal) mathematics to model IP traffic.

There is still no clear consensus on how statistics can support research in IP traffic modeling, and there is skepticism within the community regarding the utility of empirical studies that rely on collecting real data from the Internet, i.e., some claim that since the environment is changing so quickly, any collected data is only of historical interest within weeks[6]. There are those whose research is better served by tractable mathematical models than by empirical data that represent at most only one stage in network traffic evolution.

[4] North American Network Operators Group; International Engineering and Planning Group, Federal Engineering and Planning Group, Engineering and Operations Working Group; Federal Association of Regional Networks.

[5] For various definitions of connection.

[6] review cycles for research journals being months...

A further contributing factor to the lag of Internet traffic modeling behind that of telephony traffic is the early financial structure of the Internet. A few U.S. government agencies assumed the financial burden of building and maintaining the transit network infrastructure, leaving little need to trace network usage for the purposes of cost allocation. As a result Internet customers did not have much leverage with their service provider regarding the quality of service.

Many of the studies for modeling telephony traffic came largely out of Bell Labs, who had several advantages: no competition to force profit margins slim, and therefore the financial resources to devote to research, and strong incentive to fund research that could ensure the integrity of the networks for which they charge. The result is a situation today where telephone company tables of "acceptable blocking probability" (e.g., inability to get a dial tone when you pick up the phone) reveal standards that are significantly higher than our typical expectations of the Internet.

We do not have the same situation for the developing Internet marketplace. Instead we have dozens of Internet providers, many on shoestring budgets in low margin competition, who view statistics collection as a luxury that has never proven its utility in Internet operations. How will statistics really help keep the NAP alive and robust since traffic seems to change as fast as they could analyze it anyway?

We are not implying that the monopoly provider paradigm is better, only observing aspects of the situation that got us to where we are today: we have *no way to predict, verify, or in some cases even measure* Internet service quality in real time.

There is some hope that some of the larger telecommunication companies entering the marketplace with will eventually devote more attention to this area. The pressure to do so may not occur until the system breaks, at which point billed customers will demand, and be willing to pay for, better guarantees and data integrity.

cost-benefit tradeoff: undertake enough network research to secure a better understanding of the product the NAPs sell, without draining operations of resources to keep the network alive.

failing that, fund enough research to be able to show that the NAPs are being good community members, contributing to the 'advancement of Internet technology and research' with respect to understanding traffic behavior, at the least cost in terms of time and effort taken away from more critical engineering and customer service activities.

2.3 technical

With deployment of the vBNS and NAPs, the situation grows even more disturbing. The National Information Infrastructure continues to drive funding into hardware, pipes, and multimedia-capable tools, with very little attention to any kind of underlying infrastructural sanity checks. And until now, the primary obstacles to accessing traffic data in order to investigate such models has been political, legal (privacy), logistic, or proprietary.

With the transition to ATM and high speed switches, it will *no longer even be technically feasible* to access IP layer data in order to do traffic flow profiling, certainly not at switches within commercial ATM clouds. The NAPs were chartered as *layer 2 entities*, that is, to provide a service at the link layer without regard for higher layers. Because most of the NSFNET statistics reflected information at and above layer 3 (i.e., the IP layer), the NAPs cannot use the NSFNET statistics collection architecture [7] as a model upon which to base their own operational collection. Many newer layer 2 switches, e.g., DEC gigaswitch, ATM switches, have little if any capability for performing layer 3 statistics collection, or even looking at traffic in the manner allowed on a broadcast medium (e.g., FDDI, Ethernet), where a dedicated machine can collect statistics without interfering with switching. Statistics collection functionality in newer switches takes resources directly away from forwarding of frames/cells, driving customers toward switches from competing vendors who sacrifice such functionality in exchange for speed.

cost-benefit tradeoff: (minimalist approach) collect what is necessary at least cost to switching performance

2.4 ethical

Privacy has always been a serious issue in network traffic analysis, and the gravity of the situation only increases at NAPs. Regardless of whether a NAP is a layer 2 entity or a broadcast medium, Internet traffic is a private matter among Internet clients. Most NAP providers have agreements with their customers not to reveal information about individual customer traffic. Collecting and using more than basic aggregate traffic counts will require precise agreements with customers regarding what may be collected and how it will be used. Behaving out of line with customers' expectations or ethical standards, even for the most noble of research intentions, does not bode well for the longevity of a service provider.

An extreme position is to not look at network header data (which incidentally is very different from user data, which we do not propose examining) at all because it violates privacy. Analogous to unplugging one's machine from the Ethernet in order to make it secure, this approach is effective in the short term, but has some undesirable side effects. We need to find ways rather to minimize exposure rather than surrendering the ability to understand network behavior. It seems that no one has determined an 'optimal' operating point in terms of what to collect, along any of the dimensions we discuss. Indeed, the optimal choice often depends on the service provider, and changes with time and new technologies.

We acknowledge the difficulty for the NAPs, as well as any Internet service provider, to deal with statistics collection at an already very turbulent period of Internet evolution. However, it is just such a time, marked ironically with the cessation of the NSFNET statistics, that a baseline architectural model for statistics collection would be most critical, so that customers can trust the performance and integrity of the services they procure from their network service providers, and yet so service providers do not tie their hands behind their backs in terms of being able to preserve robustness, or forfend demise, of their own clouds.

cost-benefit tradeoff: accurately characterize the workload on the network so NAPs and NSPs can optimize (read, maintain) their networks, but at least cost to these privacy ethics we hold so dear.

3 utility of analysis: examples

There is no centralized control over all the providers in the Internet. The providers do not always coordinate their efforts with each other, and quite often are in competition with each other.

Despite all the diversity among the providers, the Internet-wide IP connectivity is realized via Internet-wide distributed routing, which involves multiple providers, and thus implies certain degree of cooperation and coordination. Therefore, there is a need to balance the providers' goals and objectives against the public interest of Internet-wide connectivity and subscribers' choices. Further work is needed to understand how to reach the balance.

Y. Rekhter, Routing in a multi-provider Internet, *[8]*

As the underlying network becomes commoditized, many in the community begin to focus on higher layer issues, such as the facilities, information, and opportunities for collaboration available on the network. In this context, facilities include supercomputers and workstations, information includes world wide web and gopher servers, and opportunities for collaboration include e-mail and multi-user domains (MUDs). One could think of these three categories as corresponding to three types of communication: machine-to-machine, people-to-machine, and people-to-people.

Within each category, multiple dimensions emerge:

- **context**: static topical, such as cognitive neuroscience research or geograhically topical, such as local news
- **temporal**: from permanent resources to dynamically created communication resources used for brief periods, e.g, distributed classroom lecture or seminar
- **geographic distribution**: which may require transparency at times and boundary visibility at another

As we build this higher layer infrastructure taking the lower layers increasingly for granted, the need for statistics collection and analysis is not diminished. On the contrary, it is even more critical to maintain close track of traffic growth and behavior in order secure the integrity of the network. In this section we highlight several examples of the benefits of traffic characterization to the higher layers that end users care about.

3.1 long range Internet growth tracking

Few studies on national backbone traffic characteristics [7, 9, 10, 11, 12, 13] exist, limiting our insight into the nature of wide area Internet traffic. We must rely on WAN traffic characterization studies that focus on a single or a few attachment points to transit networks to investigate shorter-term aspects of certain kinds of Internet traffic, e.g, TCP [3], TCP and UDP [14], *dns* [15].

The authors have devoted much attention in the last two years to investigating the usefulness, relevance, and practicality of a wide variety of operationally collected statistics for wide area backbone networks. In particular, we have undertaken several studies on to what extent much of the statistics that the NSFNET project has collected over the life of the NSFNET backbone are useful for a variety of workload characterization efforts. As the NSFNET service agreement ends, leaving the R&E portion of Internet connectivity in the hands of the commercial marketplace, we are without a common set of statistics, accessible to the entire community, that can allow even a rough gauge of Internet growth. We consider it important to establish a community-wide effort to support the aggregation of such network statistics data from multiple service providers, such as that being developed in [16]. We view a consensus on some baseline statistics collection architecture as critical to Internet long-term stability.

3.2 service guarantees and resource reservation

Another ramification of the transition of the the R&E portion of Internet connectivity from the NSFNET service into a commercial marketplace is the need for a mechanism to compare the quality of service providers. Rather than procured via collaborative undertakings between the federal government and academia and industry, services in today's internetworking environments will be market commodities. Effective provision of those commodities will require the ability to describe Internet workload using metrics that will enable customers and service providers to agree on a definition of a given grade of service. Furthermore, metrics for describing the quality of connectivity will be important to market efficiency since they will allow customers to compare the quality of service providers when making procurement decisions.

In addition, users will want to be able to reserve bandwidth resources, which will require that the network provider have an understanding of the current traffic behavior in order to efficiently allocate reservations without leaving unused reserved bandwidth unnecessarily idle.

However the research community has yet to determine such metrics, and the issue requires immediate attention. A precursor to developing common metrics of traffic workload is a greater understanding of network phenomena and characteristics. Insight into workload requires tools for effective visualization, such as representing the flow of traffic between service providers, or across national boundaries. Without such tools for monitoring traffic behavior, it is difficult for an Internet service provider to do capacity planning, much less service guarantees.

In addition to our studies of operational statistics, we have also undertaken several studies that collect more comprehensive Internet traffic flow statistics and we have developed a methodology for describing those flows in terms of their impact on an aggregate Internet workload. We have developed a methodology for profiling Internet traffic flows which draws on previous flow models [3, 17], and developed a variety of tools to analyze traffic based on this methodology. We think that NAP and other network service providers would gain great insight and engineering advantage by using similar tools to assess and track their own workloads.

3.3 accounting

The ability to specify or reserve the services one needs from the network will in turn require mechanisms for accounting and pricing (else there is no incentive not to reserve all one can or not to use the highest priority). Many fear pricing will stifle the open, vibrant nature of the Internet community; we suggest that it may rather motivate the constructive exploration of more efficient network implementation of innovative networking applications over the Internet. Understanding the possibilities for resource accounting will also aid commercial network providers in the process of setting cost functions for network services.

3.4 web traffic and caching

The world wide web is another domain in which operational collection and analysis of statistics is vital to support of services. Similar to our NSFNET analysis work, we have explored the utility of operationally collected web statistics, generally in the form of *http* logs. We analyzed two days of queries to the popular mosaic server at NCSA to assess the geographic distribution of transaction requests. The wide geographic diversity of query sources and popularity of a relatively small portion of the web server file set present a strong case for deployment of geographically distributed caching mechanisms to improve server and network efficiency.

We analyzed the impact of caching the results of queries within the geographic zone from which the request was sourced, in terms of reduction of transactions with and bandwidth volume from the main server [18]. We found that a cache document timeout even as low as 1024 seconds (about 17 minutes) during the two days that we analyzed would have saved between 40% and 70% of the bytes transferred from the central server. We investigated a range of timeouts for flushing documents from the cache, outlining the tradeoff between bandwidth savings and memory/cache management costs. Further exploration is needed of the implications of this tradeoff in the face of possible future usage-based pricing of backbone services that may connect several cache sites.

Other issues that caching inevitably poses include how to redirect queries initially destined for a central server to a preferred cache site. The preference of a cache site may be a function of not only geographic proximity, but also current load on nearby servers or network links. Such refinements in the web architecture will be essential to the stability of the network as the web continues to grow, and operational geographic analysis of queries to archive and library servers will be fundamental to its effective evolution.

For very heavily accessed servers, one must evaluate the relative benefit of establishing mirror sites, which could provide easier access but at the cost of extra (and distributed) maintenance of equipment and software. However, arbitrarily scattered mirror sites will not be sufficient. The Internet's sustained explosive growth calls for an architected solution to the problem of scalable wide area information dissemination. While increasing network bandwidths help, the rapidly growing populace will continue to outstrip network and server capacity as they attempt to access widely popular pools of data throughout the network. The need for more efficient bandwidth and server utilization transcends any single protocol such as *ftp*, *http*, or whatever protocol next becomes popular.

We have proposed to develop and prototype wide area information provisioning mechanisms that support both caching and replication, using the NSF supercomputer centers as 'root' caches. The goal is to facilitate the evolution of U.S. information provisioning with an efficient national architecture for handling highly popular information. A nationally sanctioned and sponsored hierarchical caching and replicating architecture would be ideally aligned with NSF's mission, serving the community by offering a basic support structure and setting an example that would encourage other service providers to maintain such operations. Analysis of web traffic patterns is critical to effective cache and mirror placement, and ongoing measurement of how these tools affect, or are affected by, web traffic behavior is an integral part of making them an effective Internet resource.

3.5 regulation

Although in previous years we have investigated telecommunications regulation issues in both the federal and state arenas, we prefer instead to invest, and see others invest, in research that might diminish the need for regulatory attention to the Internet industry. For example, Bohn *et al.* [19] have proposed a policy for IP traffic precedence that can enable a graceful transition into a more competitively supplied Internet market that might reduce the need for federal interest in regulation. The paper discusses how taking advantage of existing IP functionality to use multiple levels of service precedence can begin to address disparities between the requirements of conventional and more new and highly demanding applications.

Our approach thus far has assumed the belief that at least so far, the less regulation of the Internet, and in fact the more progress in removing regulatory barriers for existing telecommunications companies so they can participate more effectively in the Internet market, the better. However, regulation may be necessary in order to foster a healthy competitive network environment where consumers can make educated decisions regarding which network service providers to patronize. A key requirement in analyzing the relative performance of and customer satisfaction with network service providers is public availability of statistics that measure their capacity, reliability, security, integrity, and performance.

4 recommendations

In this section we highlight several examples of recommended statistics collection objects and tools.

Reaching consensus on the definition of a community-wide statistics collection architecture will require cooperation between the private and public sectors. We hope that key federally-sponsored networks such as the NSF very high speed backbone network service (vBNS)[1] and the federally sponsored NAPs can serve as a role model in providing an initial set of statistics to the community and interacting with the research community to refine metrics as research reveals the relative utility of various ones.

4.1 existing work

The last community wide document of which we are aware was Stockman's RFC 1404 [20], "A model for common operational statistics". Since that time the Internet environment and has changed considerably, as have the underlying technologies for many service providers such as the NAPs. As a result these specific metrics are not wholly applicable to every service provider but they serve as a valuable starting point. We emphasize that the exact metrics used are not a critical decision at this point, since refinements are inevitable as we benefit from experience with engineering the technologies; what is essential is that we start with some baseline and create a community facility for access and development in the future.

From Stockman [20]:

The metrics used in evaluating network traffic could be classified into (at least) four major categories:

- *utilization metrics: (input and output packets and bytes, peak metrics, per protocol and per application volume metrics)*
- *performance metrics (round trip time (rtt) on different protocol layers, collision count on a bus network, ICMP Source Quench message count, count of packets dropped)*
- *availability metrics (longer term) (line, route, or application availability)*
- *stability metrics (short term fluctuations that degrade performance) (line status transitions, fast route changes (flapping), routes per interface in the tables, next hop count stability, short term ICMP anomalous behavior)*

Some of these objects are part of standard SNMP[7] MIBs; others of private MIBs; others are not possible to retrieve at all due to technical limitations, i.e., measurement of a short term problematic network situation only exacerbates it or takes longer to perform than the problem persists. For example, counts of packets and bytes, for non-unicast and unicast, for both input and output are fairly standard SNMP variables. Less standard but still often supported in private MIBs are counts of packet discards, congestion events, interface resets or other errors. Technically difficult variables to collect, due to the high resolution polling required, include queue lengths and route changes. Although such variables would be useful for many research topics in Internet traffic characterization, operationally collected statistics will likely not be able to support them. For example, one characteristic of network workload is 'burstiness', which reflects variance in traffic rate. Network behavioral patterns of burstiness are important for defining, evaluating, and verifying service specifications, but there is not yet agreement in the Internet community on the best metrics to define burstiness. Several researchers [21, 22, 6] have explored the failure of Poisson models to adequately characterize the burstiness of both local and wide area Internet traffic. This task relies critically on accurate packet arrival timestamps, and thus on tools adequate for packet tracing of the arrivals of packets at high rates with accurate (microsecond) time granularities. Vendors may still find incentive in providing products that can perform such statistics collection, for customers that need fine-grained examination of workloads.

The minimal set of metrics recommended for IP providers in Stockman [20] were: packets and bytes in and out (unicast and non-unicast) of each interface, discards in and out of each interface, interface status, IP forwards per node, IP input discards per node, and system uptime. All of the recommended metrics were available in the Internet Standard MIB. The suggested polling frequency was 60 seconds for unicast packet and byte counters, and an unspecified multiple of 60 seconds for the others. Stockman also suggested aggregation periods for presenting the data by interval: over 24-hour, 1 month, and 1 year periods, aggregate by 15 minuets, 1 hour, and 1 day, respectively. Aggregation includes calculating and storing the average and maximum values for each period.

4.2 switched environments

In switched environments, where there is no IP layer information, the above statistics are not completely applicable. Without demand from their customer base, many switch vendors have put statistics at a second priority since it tends to detract from forwarding performance anyway. Some ATM switches can collect per-VC (virtual circuit) statistics such as those described in the ATM MIB [23].

One alternative for providers that support IP over an internal ATM infrastructure is to collect the IP statistics described above, and in fact use objects such as host-host matrices to plan what number and quality of ATM virtual circuits might be necessary to support the IP traffic workload. For switched FDDI environments, the provider could collect statistics on each LAN coming into the switch, and collapse it during analysis to determine possible compound effects within the switch. in addition segmenting traffic by interface, customer, and perhaps by protocol/application. If there is no access to network layer information, such as at NSF NAPs or certain ATM switches, the network service provider will still have an interest in these statistics, since sorting the resulting arrays would give the NSP an indication what fraction of traffic comes from what number of users, which may be critical for planning switched or permanent virtual circuit configuration, and by the same token for accounting and capacity planning purposes. However converting virtual circuits to end customers, most likely IP customers for the near future, requires maintaining mappings to higher layers.

[7] Simple Network Management Protocol

4.3 dedicated studies of IP workloads

Even with a solid set of operational statistics there are times when one wants dedicated collection to gain greater insight into short-term dynamics of workloads. For example, there are limitations of the operationally collected statistics for the NSFNET for describing flows in terms of their impact on an aggregate Internet workload [24]. We have developed tools for supporting operational flow assessment, to gain insight into both individual traffic signatures as well as heavy aggregations of end users. We have tested our methodology using packet header traces from a variety of Internet locations, yielding insight far beyond a simple aggregated utilization assessments, into details of the composition of traffic aggregation, e.g., what components of the traffic dominate in terms of flow frequency, durations, or volumes. We have shown that shifts in traffic signatures as a result of evolving technologies, e.g., toward multimedia applications, will require a different approach in network architectures and operational procedures. In particular, the much higher demands of some of these new applications will interfere with the ability of the network to aggregate the thousands of simultaneous but relatively short and low volume flows that we observe in current environments.

The methodology [24] defines a flow based on actual traffic activity from, to, or between entities, rather than using the explicit setup and teardown mechanisms of transport protocols such as TCP. Our flow metrics fall into two categories: metrics of individual flows and metrics of the aggregate traffic flow. Metrics of individual flows include: flow type, packet and byte volume, and duration. Metrics of the aggregate flow, or workload characteristics seen from the network perspective, include: counts of the number of active, new, and timed out flows per time interval; flow interarrival and arrival processes; and flow locality metrics. Understanding how individual flows and the aggregate flow profile influence each other, is essential to securing Internet stability, and requires ongoing flow assessment to track changes in Internet workload in a given environment.

Because it requires comprehensive and detailed statistics collection, we recognize that the NAPs and other service providers may not be able to afford to continuously monitor flow characteristics on an operational basis. Nonetheless we imagine that NAP operators will find it useful to undertake traffic flow assessment at least periodically to obtain a more accurate picture of the workload their infrastructure must support. The methodology and tools that implement it[8] will be increasingly applicable, even on a continuous basis, for NAP tasks such as ATM circuit management, usage-based accounting, routing table management, benchmarks by which to shop for equipment from vendors, and load balancing in future Internet components.

The methodology can form a complementary component to other existing operational statistics collection, yielding insights into larger issues of Internet evolution, i.e., how environments of different aggregation can cope with resource contention by an ever-changing composition and volume of flows. Internet traffic cross-section and flow characteristics are a moving target, and we intend that our methodology serve as a tool for those who wish to track and keep pace with its trajectory. For example, as video and audio flows, and even single streams combining voice and audio, become more popular, Internet service providers will need to parametrize them to determine how many such end user streams they will be able to support, and how many more resources each new such stream would require. Multicast flows will also likely constitute a increasingly significant component of Internet traffic, and applying our methodology to multicast flows would be an important step toward coping with their impact on the infrastructure.

4.4 community exchange

A key requirement in analyzing the relative performance of and customer satisfaction with network service providers is public availability of statistics that measure their capacity, reliability, security, integrity, and performance. One vehicle conducive to the development of an effective marketplace is a client/server based architecture for providing access to statistics of various NSPs. Each NSP would support a server that provides query support for collected statistics for its clients or potential customers. MCI provides this functionality for statistics collected on the NSF-sponsored vBNS network. Community wide participation in such an open forum would foster a healthy competitive network environment where consumers could make educated decisions regarding which network service providers to patronize.

[8] The software we used is available at *ftp://ftp.sdsc.edu/pub/sdsc/anr/software/Flows.tar.Z.*

5 conclusion

"when we get through we won't be done."

Steve Wolff, director, DNCRI, NSF on the NSFNET transition (one among many since 1983) at 1994 Farnet meeting

We conclude by mentioning an issue of recent popular interest as well as symbolic of the larger problem: Internet security and prevention of criminal behavior. Ironically, workload and performance characterization issues are inextricably intertwined with security and privacy. Much of the talk about the Internet's inherent insecurity due to the inability to track traffic at the required granularity is misleading. It is not an *inability* to examine traffic operationally that has prevented it thus far, whether for security or workload characterization (and the same tools could do both), but rather its priority relative to the rest of the community research agenda.

As a result, the gap has grown large between the unambiguous results of confined experiments which target isolated environments, and the largely unknown characteristics of the extensive Internet infrastructure that is heading toward global ubiquity.

Empirical investigation and improved methodology for doing so can improve current operational statistics collection architectures, allowing service providers to prepare for more demanding use of the infrastructure and allowing network analysts to develop more accurate Internet models. In short, we can contribute to a greater understanding of real computer networks of pervasive scale. by reducing the gaps among: (1) what network service providers need; (2) what statistics service providers can provide; and (3) what in-depth network analysis requires.

We encourage discussion as soon as possible within the community on developing a policy on statistics collection/exchange/posting of available NAP/ISP statistics, with supporting tools to allow greater understanding of customer requirements and service models, equitable cost allocation models for Internet services, verification that a given level of service was actually rendered, and evolution toward a level of Internet performance that matches or surpasses that of most telecommunication systems today.

author information

Hans-Werner Braun is a Principal Scientist at the San Diego Supercomputer Center. Current research interests include network performance and traffic characterization, and working with NSF on NREN engineering issues. He also participates in activities fostering the evolution of the national and international networking agenda. San Diego Supercomputer Center, P.O. Box 85608, San Diego, CA 92186-9784; *hwb@sdsc.edu*.

Kimberly Claffy received her doctoral degree from the Department of Computer Science and Engineering at the University of California, San Diego in June, 1994, and is currently a Associate Staff Scientist at the San Diego Supercomputer Center. Her research focuses on establishing and improving the efficacy of traffic and performance characterization methodologies on wide-area communication networks, in particular to cope with the changing traffic workload, financial structure, and underlying technologies of the Internet. San Diego Supercomputer Center, P.O. Box 85608, San Diego, CA 92186-9784; *kc@sdsc.edu*.

References

[1] H.-W. Braun, C. E. Catlett, and K. Claffy, "http://vedana.sdsc.edu/," tech. rep., SDSC and NCSA, March 1995. San Diego Supercomputer Center.

[2] H.-W. Braun, C. E. Catlett, and K. Claffy, "National Laboratory for Applied Network Research," tech. rep., SDSC and NCSA, March 1995. San Diego Supercomputer Center.

[3] R. Caceres, P. Danzig, S. Jamin, and D. Mitzel, "Characteristics of wide-area TCP/IP conversations," in *Proceedings of ACM SIGCOMM '91*, pp. 101–112, September 1991.

[4] V. Paxson, "Empirically-Derived Analytic Models of Wide Area TCP Connections." *IEEE/ACM Transactions on Networking, vol. 2 no. 4*, August 1994.

[5] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," in *IEEE/ACM Transactions on Networking*, vol. 2, February 1994.

[6] V. Paxson and S. Floyd, "Wide-area traffic: the failure of Poisson modeling," in *Proceedings of ACM SIGCOMM '94*, February 1994.

[7] K. Claffy, H.-W. Braun, and G. C. Polyzos, "Long-term traffic aspects of the NSFNET," in *Proc. of INET '93*, pp. CBA–1:10, August

1993.

[8] Y. Rekhter, "Routing in a multi-provider Internet." Internet Request for Comments Series RFC 1787, April 1995.

[9] M. Davis, "Analysis and optimization of computer network routing," Master's thesis, University of Delaware, 1988.

[10] S. Heimlich, "Traffic Characterization of the NSFNET National Backbone," in *Proceedings of the 1990 Winter USENIX Conference*, December 1988.

[11] K. Claffy, G. C. Polyzos, and H.-W. Braun, "Traffic characteristics of the T1 NSFNET backbone," in *Proceedings of IEEE Infocom 93*, pp. 885–892, March 1993.

[12] K. Claffy, *Internet workload characterization*. PhD thesis, UC, San Diego, June 1994.

[13] K. Claffy, H.-W. Braun, and G. C. Polyzos, "Tracking long-term growth of the NSFNET backbone," *Communications of the ACM*, vol. 37, pp. 34–45, August 1994.

[14] A. Schmidt and R. Campbell, "Internet protocol traffic analysis with applications for ATM switch design," *Computer Communiations Review*, vol. 23, pp. 39–52, April 1993.

[15] P. B. Danzig, K. Obraczka, and A. Kumar, "An analysis of wide-area name server traffic," in *Proceedings of ACM SIGCOMM '92*, August 1992.

[16] H. Wolff and the IETF Opstat working group, "The opstat client-server model for statistics retrieval," April 1995. draft-ietf--opstat-client-server-03.txt.

[17] R. Jain and S. A. Routhier, "Packet trains—measurement and a new model for computer network traffic," *IEEE Journal on Selected Areas in Communications*, vol. 4, pp. 986–995, September 1986.

[18] H.-W. Braun and K. Claffy, "Web traffic characterization: an assessment of the impact of cach ing documents from NCSA's web server," in *Second international world wide web conference*, October 1994.

[19] R. Bohn, H.-W. Braun, K. Claffy, and S. Wolff, "Mitigating the coming Internet crunch: multiple service levels via precedence," *Journal of High Speed Networks*, forthcoming 1994. available for anon-ftp at `ftp.sdsc.edu:pub/sdsc/-anr/papers/`, and `http://www.sdsc.edu/-0/SDSC/Research/ANR/kc/precedence/-precedence.html`.

[20] B. Stockman, "A model for common operational statistics." RFC 1404, January 1993.

[21] W. Willinger, "Self-similarity in high-speed packet traffic: Analysis and modeling of ethernet traffic measurements," *Statistical Science*, 1994.

[22] M. W. Garrett and W. Willinger, "Analysis, modeling and generation of self-similar VBR video traffic," *Proceedings of ACM SIGCOMM '94*, September 1994.

[23] M. Ahmed and K. Tesink, "Definitions of managed objects for atm management, version 7.0." Internet Request for Comments Series RFC 1695, August 1994.

[24] K. Claffy, G. C. Polyzos, and H.-W. Braun, "Internet traffic flow profiling," *IEEE Journal on Selected Areas in Communications*, to appear 1995.