Measurement and Analysis of IP Network Usage and Behavior

R. Cáceres, N.G. Duffield, A. Feldmann, J. Friedmann, A. Greenberg, R. Greer, T. Johnson, C. Kalmanek, B. Krishnamurthy, D. Lavelle, P.P. Mishra, K.K. Ramakrishnan, J. Rexford, F. True, J.E. van der Merwe

AT&T *

Abstract

Traffic, usage, and performance measurements are crucial to the design, operation and control of Internet Protocol (IP) networks. This paper describes a prototype infrastructure for the measurement, storage and correlation of network data of different types and origins from AT&T's commercial IP network. We focus first on some novel aspects of the measurement infrastructure, then describe analyses that illustrate the power of joining different measured data sets for network planning and design.

1 Introduction

The Internet is extremely diverse, both in terms of the technologies that form the network and the applications that use it. It includes many types of communication links, from slow dialup lines to fast optical fibers. It also carries many types of traffic, from electronic commerce transactions to images from faraway spacecraft. Furthermore, new technologies and applications emerge frequently. For example, a number of wireless subnetworks have come online in recent years. Similarly, the World Wide Web appeared in the mid-1990's and quickly grew to be the dominant source of traffic. Because of all these factors, measurement remains the only effective way to determine how the Internet is being used and how it is behaving.

In early 1997 we began an effort to measure and analyze the usage and behavior of WorldNet, AT&T's Internet Service Provider (ISP) business. We have deployed a prototype measurement infrastructure with both passive and active components. Passive components include custom-built packet sniffers we call PacketScopes, flow-level traffic statistics gathered at all the border routers, and an otherwise idle top-level router that we query for routing information. Active components include a mesh of measurement machines located at major router centers. These machines exchange measurement traffic and collect loss, delay, and connectivity statistics throughout the day. We continuously feed all these measurements into a high-performance data repository we call the WorldNet Data Warehouse. We have used the data both to drive research studies and to help run the network.

The variety of measurement data at our disposal, together with the ability to manipulate data from multiple sources in a single storage and analysis platform, enable new and important characterizations of many aspects of the Internet. As a motivating example, consider the millions of Web users who connect to the Internet daily over dialup lines. What workload does each user place on the network? How is their traffic treated by the network? How much of the user-perceived latency is due to the network and how much is due to the server? Any one form of measurement is insufficient to answer these questions thoroughly.

Let's start with the problem of characterizing Web

*Address for correspondence: N.G. Duffield, AT&T Labs, 180 Park Avenue, Rm. B139; Florham Park, NJ 07932, USA; E-mail: duffield@research.att.com; Tel: +1 973 360 8726

accesses by dialup users. We have gathered traces of all traffic flowing between a large dialup modem pool inside WorldNet and the rest of the Internet. However, dialup users are typically assigned a different IP address each time they call in. Although our traffic traces contain very detailed information, including IP addresses and timing information for each packet, they are not enough to determine when a dialup session begins and ends. We also maintain a separate log of all the calls made into the WorldNet dial platform. This log includes the begin and end times of each dialup session along with the IP address assigned to that session. By correlating this information in the call log with our packet traces, we can isolate the traffic streams for individual dialup sessions and thus proceed to answer the question of how a typical dialup user accesses the Web.

Next, let's turn to the problem of determining how Web traffic to and from dialup users is shaped by the network. We have deployed PacketScopes next to a dialup modem pool in WorldNet and next to a bank of Web servers that form part of Easy World Wide Web (EW3), AT&T's Web hosting business. By gathering simultaneous traces of the traffic flowing through these two points, we can determine, for example, which packets have been lost in between. However, there is no guarantee that the users represented in the client-side traces will access the Web pages represented in the server-side traces. We can arrange to capture traffic at both ends by dialing up to the modem pool and accessing the Web servers ourselves, at the same time measuring performance as seen by our own Web clients. We can also make use of the access logs routinely maintained by the EW3 Web servers. By joining all these active and passive measurements, we can obtain a complete picture of a Web transfer, including the request by a Web client, the progress of the related Web traffic through the network, and the handling of the request by a Web server. This analysis addresses the questions on network and server performance posed earlier.

The rest of this article is organized as follows. In Section 2 we describe briefly AT&T's IP infrastructure in the United States, the types and provenance of measurements collected from it, the repository where the measurement data is stored and analyzed, and the measures we've taken to preserve privacy. In Section 3 we describe two measurement systems in more detail: the AT&T Labs PacketScope and an active measurement system currently deployed by AT&T. In Section 4 we describe some sample analyses that we've carried out on the measurement data.

2 Data Warehouse and its Sources

2.1 Networking Infrastructure

AT&T's IP services in the United States are supported by the following broad networking infrastructure. The IP backbone comprises a number of major hub sites distributed in the continental US, connected by a mesh of high-speed connections. In this paper we consider customer access to the network by one of two means: dedicated or dial access. Dedicated access is via access routers at one or more of the major hub sites; this is the usual mode of access for business customers. Dial users gain access over the public switched telephone network to modem groups distributed across the US; this mode of access is used mainly by domestic customers. The dial service is supported by multiple data centers containing servers handling customer registration and billing functions, while other servers provide standard services such as email and netnews. There are also servers supporting Web hosting for both dial and dedicated users. Connectivity to other segments of the Internet is achieved through gateways at major hub nodes. These gateways connect to private peers and to public interconnection points.

2.2 Data Sources

The WorldNet Data Warehouse currently maintains a large number of different data sets. In this section we summarize those measured data sets that are updated regularly; the origins of these data in the AT&T IP network are represented in Figure 1.

Packet Headers: These are collected by three AT&T PacketScopes installed at representative locations: on a T3 link between WorldNet and one of its peers; on a FDDI ring carrying traffic to and from a group of roughly 450 modems shared by approximately

18,000 dialup customers; and, on a FDDI ring carrying traffic to and from a commercial Web hosting service. Each PacketScope can produce approximately 90 GB of data per weekly sample. The PacketScope architecture is described in more detail in Section 3.1.

Flow Statistics: Cisco NetFlow [3] is enabled in Internet Gateway Routers (IGRs). These routers export per-flow summaries including: start time; flow duration; byte and packet count; source and destination IP addresses, port numbers and autonomous systems.

Routing Tables: BGP routing tables and forwarding tables for key backbone and access routers. An otherwise idle router peers with a gateway router, and collects routing updates which are then exported to the warehouse.

Router Statistics: Obtained by SNMP polling. They include: link/router utilization, fault events for all routing elements of backbone, access and gateway routers.

Router Configuration: Includes information on topology, security, access lists, BGP/OSPF routing metrics, etc.

Registration Records: Authentication, authorization and accounting (AAA) systems supply customer account records, including, e.g., customers' choice of pricing plan and date of registration.

Call Record Journal: Per-session summaries of dialup customer sessions, including modem dialed, session start and end times, total bytes and packets in each direction, and connect speed.

Web Server Logs: Collected at AT&T's Web hosting service. Fields include: IP address or name of the remote client; date and time of the request; first line of the request including the HTTP method and URL; HTTP response status code; number of bytes in the response; the referrer field.

Email Server Logs: SMTP and POP3 transaction summaries.

Active Measurements: The Worldnet In-Process Metric (WIPM) system conducts active measurements of packet loss, delay and throughput between each pair of major routing hubs, as described further in Section 3.2.

2.3 Data Warehouse

The WorldNet Data Warehouse is a prototype high-performance repository for the storage and analysis of measurements and operational data from the AT&T IP network described above. It provides extensive facilities for mining, correlating, analyzing, and archiving this information. Many research projects have grown up around the warehouse, both in the construction of experimental measurement infrastructure and in the analysis of measured data. Research collaborations have spanned a user group comprising members of the research, engineering and management communities within AT&T. These have led to the construction of prototype decision-support systems for the engineering, marketing, and strategic planning of AT&T's Internet services.

The warehouse is organized around two computing platforms. First, a general-use server employs Oracle to store pre-computed aggregate and desensitized detail information. Second, a large data storage server contains detailed measurements; these are stored in either raw form or using AT&T's Daytona database [9] described below. This second platform comprises a Sun Enterprise 6000 server with twenty-two 333 MHz processors, 13 GB of main memory, 1.8 TB of disk storage and a 14-TB attached tape library

The two-platform architecture has several advantages. Data stored on the Oracle-based platform is available to a large number of AT&T users, since Oracle client software is widely used within AT&T. As another benefit, Oracle has a wide variety of its own and third-party sophisticated GUI applications, including various data warehouse aggregation engines.

The Daytona data management system [9] offers several definitive advantages recommending its use in the data storage server. The primary advantage is speed. Daytona translates its Cymbal high-level query language completely into C, which is then compiled into a native machine executable; most other systems interpret their SQL. The Cymbal query language is a very powerful and expressive 4GL that synthesizes a procedural language with several declarative languages including SQL. Its ability to express complex queries, together with its parallelism, enable its effective use of the data storage

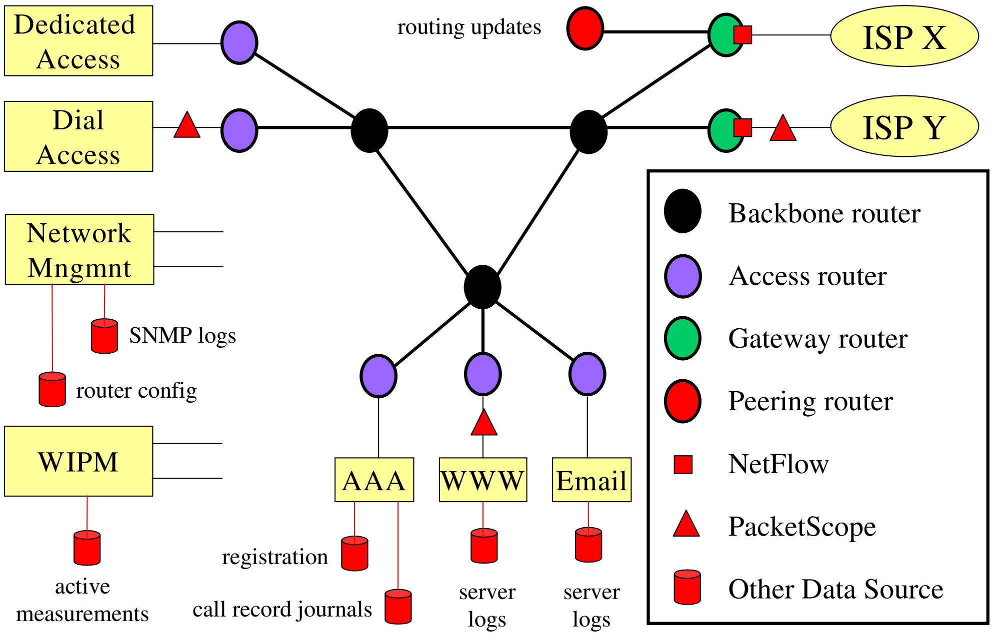

Figure 1: Measurement Infrastructure and Warehouse Data Sources

server's considerable resources. In contrast to the server processes acting as operating systems used by other data management systems, Daytona uses Unix as its server, thus saving on code and complexity.

Daytona reduces disk usage as compared with other data management systems through a combination of partitioning and compression techniques. This is a great advantage for the WorldNet Data Warehouse since the amount of data to be stored is truly enormous. By compressing each record individually, compressed tables remain fully indexable with decompression happening quickly only for the records of interest. If the compression features are not used, then Daytona stores its data in flat ASCII file format. Users have the option to work with these directly using, e.g., `awk` or `Perl`.

2.4 Privacy Measures

Privacy is an important concern any time network usage data is collected. We have taken explicit measures to safeguard user privacy in every aspect of the work described in this article. For example, our PacketScopes encrypt IP addresses as soon as they read packets from the network, before they write traces to stable storage. This procedure anonymizes the traces and helps prevents their misuse. We also discard the user data portion of these packets and work only with protocol header information.

To further protect privacy, we have implemented numerous security procedures to prevent unauthorized access to our measurement data. For instance, our PacketScopes allow only authenticated and encrypted access through private lines inside the AT&T firewall, not through the external Internet. In addition, the Warehouse is composed of specially administered machines inside the AT&T firewall. Access to raw measurement data is thus restricted while desensitized data is made available on a separate machine.

Finally, throughout our work we report only aggregate statistics that do not identify individual users. The primary purpose of the data is to support network engineering and to better understand the evolution of various protocols and applications. Through all the above privacy measures we ensure conformance with the AT&T Online Privacy Policy described in `http://www.att.com/privacy`.

3 Instrumentation

3.1 PacketScope

The AT&T Labs PacketScope is a high-performance system for the collection of IP packet headers. Its principal components are a Compaq 500 MHz Unix workstation for packet capture, a 10GB striped disk array, and a 140GB tape robot. Each has dedicated T1 access to the Data Warehouse that can be used to, e.g., remotely control data collection, retrieve data and install software.

The main design criteria for the PacketScope were as follows. *Passivity:* the monitors cannot disturb the monitored network. *Security:* control of the PacketScopes is authenticated and encrypted, with no control possible through the interface to the monitored network. *Versatility:* Packet header capture is performed by the Unix workstations using a modified version of the `tcpdump` tool [10]. This platform enabled rapid development, including customizations for security and high-level protocol header tracing. Per-unit hardware cost was *not* a major issue for our limited deployment, in distinction with the approach of [1]. We now explain the detailed PacketScope configuration; see Figure 2.

Passive Link Access At the workstation, the device driver for the monitored network interface was modified to allow reads but not writes. In addition, no IP address is assigned to the monitored interface. To monitor a multi-access channel, such as a FDDI ring, attaching a station to the channel in the usual manner is all that is required. To monitor each direction of a point-to-point T3 link, a copy of the line signal is extracted from the monitor jack of a DSX-3 panel through which the link passes in a Central Office. The signal is passed through an isolating amplifier and thence on a dedicated monitoring T3 to the PacketScope.

T3 Monitoring At the PacketScope, the monitoring T3 is terminated at a Cisco 7505 router, which forwards all packets to the monitor for capture. For security purposes it is desirable not to have an IP path from the router to the monitor. Instead, the router is configured with a static default route to a fictitious IP address in a subnet associated with the router interface to which the workstation attaches. Forwarding is accomplished by setting a static ARP entry in the router that associates the MAC address of the monitor's interface with this IP address. `tcpdump` captures packets in promiscuous mode on this interface.

Packet Capture In the workstation, the path from the device driver to the packet filter utility has been enhanced to provide good packet-capture performance under heavy receive load; overload behavior is graceful. `tcpdump` saves the captured packet headers to disk. Optionally, these headers are then copied to tape. By a combination of buffering to disk and exchanging tape cartridges we have taken continuous header traces over a period of over 2 weeks. This ability to capture long, uninterrupted traces sets our approach apart from others, e.g., [1].

Multimedia Traffic Monitoring We have deployed in the PacketScopes a new tool for monitoring Internet multimedia traffic. The `mmdump` tool [2] parses traffic from H.323, RTSP, SIP, and similar multimedia session-control protocols to set up and tear down packet filters as needed to capture multimedia sessions. Dynamic packet filters are necessary because these protocols dynamically negotiate TCP and UDP port numbers to carry the media content. Thus, we cannot rely on well known port numbers as we have for more traditional traffic types. We are using `mmdump` to track the emergence of streaming media, voice over IP, and other multimedia applications in the Internet. We are also evaluating their impact on the network, for example, we would like to determine what forms of congestion control, if any, these new traffic types use.

HTTP Protocol Tracing A significant portion of packets observed are due to Web traffic. In addition to understanding the packet-level characteristics of such traffic, it is useful to characterize behavior at larger scales such as entire HTTP transactions. We have deployed in our PacketScopes `blt` [5], a tool for extracting HTTP-level traces from raw packet traces. A single HTTP transaction may be contained in the payload of several IP packets; the

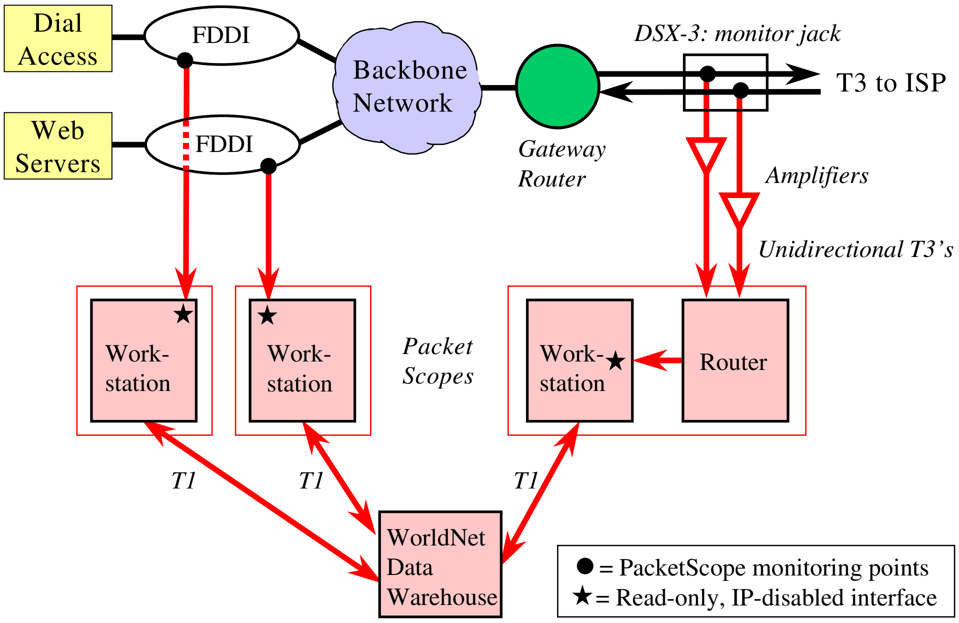

Figure 2: PacketScope Configuration: measurement infrastructure shown in red. Tape robots, disk arrays, and some other PacketScope elements not shown.

challenge is to demultiplex IP packets into TCP connections and then into individual HTTP transactions. By processing `tcpdump` traces online, we construct a full log of HTTP request and response headers and relevant timestamp information. Offline post-processing matches HTTP request and response information where applicable.

3.2 Active Measurements

The Worldnet In-Process Metric (WIPM) system provides an infrastructure for measuring edge-to-edge performance of the IP backbone using a suite of active measurements. The system currently measures and reports round-trip packet delay, round-trip packet loss, throughput, and HTTP transfer delay (including DNS resolution time) across the fully-connected mesh of the IP backbone, as well as to a set of remote access and external Internet sites, in real time. While passive measurements provide views of the performance of individual links or nodes, active measurements provide complementary views of the performance of a path consisting of several links and nodes. This measurement data is used for traffic engineering, performance debugging, network operations, and to measure compliance with performance targets. Such targets are used as a basis for defining service level agreements with customers.

The WIPM system is composed of a series of measurement servers deployed in all backbone routing complexes, which function as probe delivery and measurement points. Probe results are collected and correlated by a secure, distributed harvesting application, and are then ingested into a centralized data warehouse where they are available for reporting, alarming, and correlation with other data sources. Historical data is retained indefinitely, and is available on-line in the data warehouse for a period of one year. WIPM is similar to other active measurement infrastructures that have been deployed in recent years, including AMP, Felix, IPMA, NIMI, Surveyor, and Test Traffic; see [4].

4 Analysis Studies

4.1 NetScope: A Unified Toolkit for Managing IP Networks

Managing large IP networks requires understanding the current traffic flows, routing policies, and network configuration. Yet, the state of the art for managing IP networks involves manual configuration of each IP router, and traffic engineering based on limited measurements. The networking industry is sorely lacking in software systems that a large ISP can use to support traffic measurement and network modeling, the underpinnings of effective traffic engineering.

The AT&T Labs *NetScope* [8], a unified toolkit to manage IP networks, draws on the many feeds of configuration and usage data. Router configuration files and forwarding tables are joined together to form a global view of the network topology. The topology model includes layer-two and layer-three connectivity, link capacities, and the configuration of the routing protocols, as well as the homing locations for customer and external IP addresses. The tool also associates the routers and links with statistics derived from network measurements. For example, the feed of SNMP data provides information about CPU load, as well as link utilization and loss statistics, which can be associated with the appropriate router and its incident links. Similarly, active measurements of delay, throughput, or loss, can be associated with the appropriate link(s). Finally, by combining the NetFlow data with the topology model, we can compute a traffic matrix of load between pairs of routers in the backbone.

Joining these many feeds of configuration and usage data in a common information model facilitates a wide range of activities for the design and operation of IP networks. In addition, NetScope includes a model for how packets are routed through the network, based on the configuration of the routing protocols. Using NetScope, a network provider can experiment with changes in network configuration in a simulated environment, rather than the operational network.

Figure 3 depicts a NetScope view in a traffic engineering application. The visualization shows the flow of traffic from a single ISP peer through the AT&T common IP backbone. NetScope draws on the traffic matrix and the routing to display the sink tree for the traffic from the ISP peer. Nodes correspond to router centers and edges to layer-3 links between router centers. Color and size are used to show the amount of traffic flowing to each node and link. By combining configuration and usage data with a routing model, NetScope can answer important questions about the behavior of the network, such as "Why is this link congested?" and "How would a change in the configuration of intradomain routing help alleviate congestion?"

4.2 Web Flow Management

We used our PacketScope traces to evaluate policies for carrying Web traffic over a flow-switched network infrastructure. The explosive growth of Internet traffic has motivated the search for more efficient packet-switching techniques. To exploit recent advances in switching hardware, several proposals call for grouping sequences of related packets into flows and sending these packets through fast switching paths, or shortcuts. Shortcuts improve the performance experienced by network traffic but consume extra network resources to create and maintain.

We based our study on continuous one-week traces of Web traffic gathered by our PacketScope in Bridgeton, MO, and by a similar system within our research laboratory. We assessed the effects of three network parameters (address aggregation, idle timeout, and initial trigger) on three cost and performance metrics (shortcut setup rate, number of simultaneous shortcuts, and percent of traffic carried on shortcuts). We focused on the full probability distributions of Web flow characteristics and the metrics of interest. We found that moderate levels of aggregation and triggering yield significant reductions in overhead with a negligible reduction in performance. Our results also suggested schemes for further limiting overhead by temporarily delaying the creation of shortcuts during peak load, and by aggregating related packets that share a portion of their routes through the network. Further information is available from our earlier paper on this topic [6].

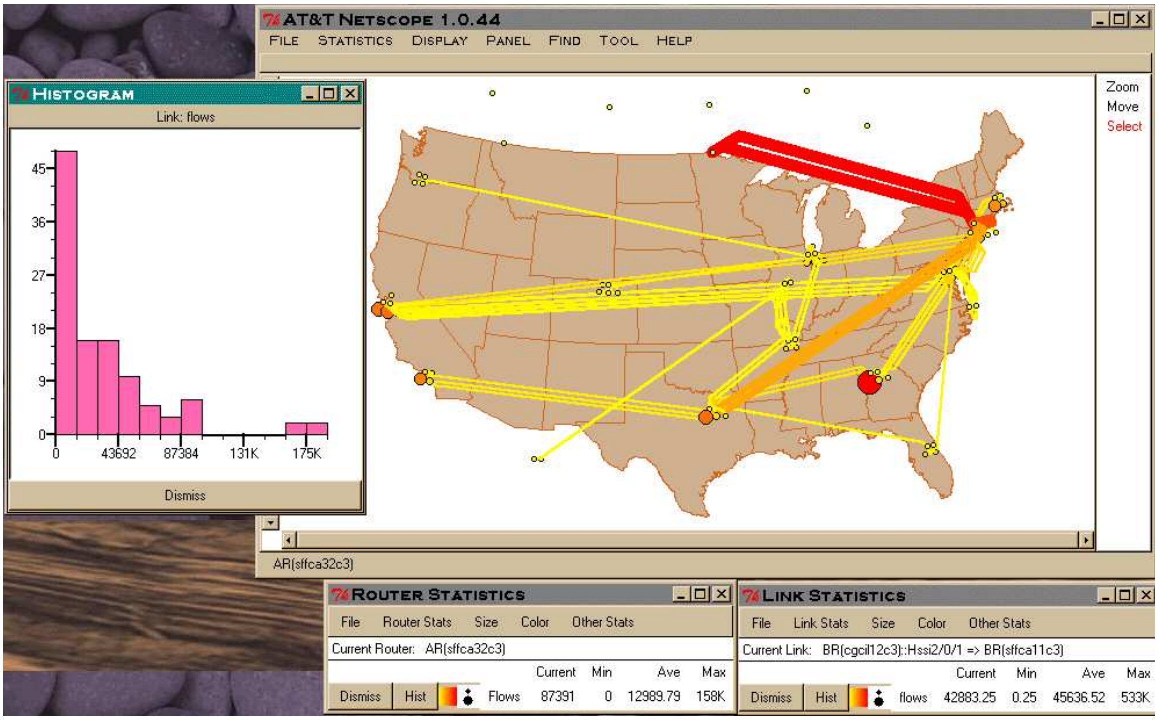

Figure 3: A NetScope view in a traffic engineering application.

4.3 Web Proxy Caching

We used our PacketScope traces to study the performance of Web proxy caching. Much work in this area has focused on high-level metrics such as hit rates, but has ignored low-level details such as "cookies," aborted connections, and persistent connections between clients and proxies as well as between proxies and servers. These details have a strong impact on performance, particularly in heterogeneous bandwidth environments where network speeds between clients and proxies are significantly different than speeds between proxies and servers.

We evaluated through detailed simulations the latency and bandwidth effects of Web proxy caching in such environments. We drove our simulations with packet traces from the PacketScope in Bridgeton, MO, as well as from inside our research laboratory. There were three main results. First, caching persistent connections at the proxy can improve latency much more than simply caching Web data. Second, aborted connections can waste more bandwidth than that saved by caching data. Third, cookies can dramatically reduce hit rates by making many documents effectively uncachable, since cookies accompany personalized documents. Additional detail is available from our recent paper on this topic [7].

4.4 Peering

Source and destination IP addresses in a given Internet packet often belong to different Autonomous Systems (ASs), administered by different Internet Service Providers (ISPs). The path such a packet takes through ASs depends on the physical connectivity between ISPs and their customers, together with policies for interdomain routing applied locally within each AS. Interdomain connectivity and routing policy depend, in turn, on business agreements between ISPs, commonly termed *peering* agreements. Roughly, if flows between two ISPs A and B are sufficiently large and balanced (e.g., the ratio of the traffic that A sends B to the traffic that B sends A is near 1) then the two ISPs evenly share the costs of interconnecting. If the flows are not balanced (e.g., if the traffic from A to B is significantly smaller than in the other direction), then A appears a customer of B, and pays B accordingly. (Details of peering relationships are considered proprietary.) In negotiating a private peering agreement between two ISPs, each would typically want to collect its own, independent measurements.

Consider two ISPs, A and B, initially having no direct connections, but entering into negotiations to establish direct connections. A must estimate traffic volumes to B and from B. Let us consider the latter, thornier problem. A would like to estimate the traffic volumes for BGP routes that B uses to get traffic to A. In general, A only sees BGP routes in the other direction (from A to B). However, using data such as BGP tables and flow statistics collected at border routers, rough but workable bounds can be found. To obtain a lower bound, A can account for all flows that originate in the ASs belonging to B. An approximate upper bound, can be obtained under the assumption of symmetric BGP routes: extract from the BGP tables from A all prefixes that contain B in their BGP list, and sum all flow volumes whose source prefix is in this set of prefixes. Though BGP routes are commonly asymmetric, estimates we have obtained by this technique have nevertheless proved quite useful. In practice, using this technique in the role of A estimating inflows from B, we have obtained estimates close to those obtained by B through direct measurements of outflows to A.

Table 1 shows data derived from one such investigation. Flow measurements were taken at ISP A's border routers, with ISP A representing AT&T WorldNet, and ISP B another ISP that at the time was reachable only via the four NAPs listed: AADS, MaeEast, PacBell, and Sprint. These measurements, taken over the course of a week, showed flows between the two ISPs to be sufficiently balanced that a private peering agreement was struck, with costs split evenly.

NAP	$A \to B$		$B \to A$		Ratio
	Bytes	Mbps	Bytes	Mbps	
AADS	18G	9.64	25G	13.4	1.39
MaeEast	30G	15.8	17G	9.06	0.57
PacBell	24G	11.9	65G	32.5	2.73
Sprint	22G	11.5	16G	8.53	0.74

Table 1: Traffic volumes between two ISPs, A and B, considering a peering arrangement.

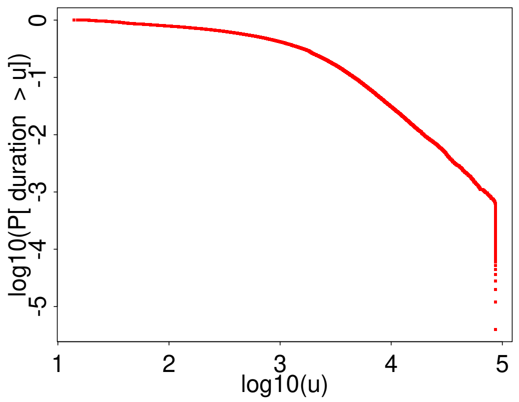

Figure 4: Tail distribution of dial session duration.

4.5 Dial Service Pricing

How should an ISP market dial access, to gain and hold market share, while making an adequate return on investment? To make their service attractive, ISPs must place dial POPs (Points of Presence) at locations that potential customers can reach via an unmetered local call, i.e., with price independent of duration. If the ISP charges a flat (usage insensitive) fee for access, customers may have little incentive to log off during prolonged idle periods.

Figure 4 depicts the measured distribution of call holding times of AT&T WorldNet dial sessions (obtained from the Call Record Journal) measured when the service had a flat fee. The linear slope on a log-log scale is characteristic of a *heavy-tailed* distribution. (The sharp cut off at the upper end is due to a disconnect policy for sessions lasting 24 hours, in force at the time.) Closer analysis joining customer registration data with the session usage data, carried out in early 1998, revealed that 3% of the users (those with monthly usage exceeding 150 hours) accounted for about 30% of the usage minutes.

As a result of analyses of this type, AT&T WorldNet offered a new pricing plan, which became effective in June of 1998, providing up to 150 hours of service for a fixed monthly fee, and a usage sensi-

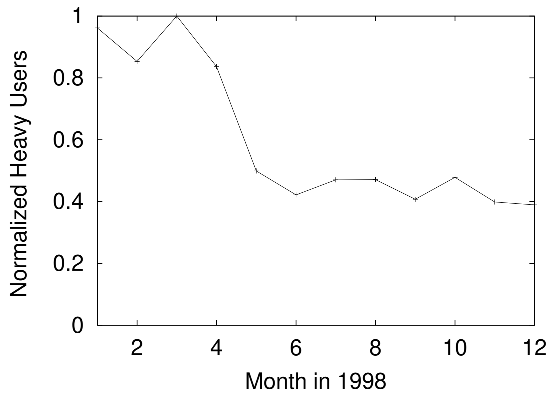

Figure 5: (Normalized) number of WorldNet heavy users, i.e., those using service more than 150 hours per month. Usage sensitive pricing beyond 150 hours was announced 4/98, effective 6/98.

tive charge for service beyond 150 hours. The result was dramatic. Starting from the time the service was announced (and before it became effective), usage patterns began to change; in particular, heavy users began logging off, effectively cutting the heavy tail. This is illustrated in Figure 5. Internal and independent measurements indicated significant improvements in access availability; the chance of connecting to the service on the first call attempt increased significantly.

Of course, pricing plans come and go in response to market forces. However, this study helps to understand the qualitative effects of the plans. Unlimited usage billing leads to the growth of a small community of heavy users consuming a disproportionate share of resources. Limited usage billing can effectively reduce the impact of such heavy users.

4.6 Other Studies

In this section we have focused on some applications of our measurement data to network design, engineering, operations, and marketing. Many other studies have been conducted using these and other data sets that we have not yet described. We mention some of these studies here.

Placing Points of Presence (POPs): Consumers are expected to find the dial service more attractive if they can make an unmetered telephone call to a POP. This study uses information about local calling areas to optimally place new POPs in order to maximize the number of telephone lines that can make local calls to them.

Churn: New dial customers represent an investment for any ISP. Understanding the reasons subscribers churn (i.e., leave the service) provides information about how to improve service and can help target marketing. Predicting churn rates is part of forecasting expenses and revenues. Based on the the Call Record Journal and the Registration Records, this study suggests that churning and non-churning customers show quite different usage profiles early on in their subscription history.

Scaling Analysis: Analysis of the time series of packet arrivals has revealed new scaling structure in Internet traffic: self-similarity at long time scales; a change in scaling behavior at intermediate timescales; and multifractal scaling at short timescales. These phenomena can be related to the presence of closed-loop controls, such as in TCP, or open-loop controls, such as in UDP applications.

5 Conclusions

We have presented a measurement and analysis infrastructure for characterizing the usage and behavior of IP networks. The infrastructure comprises both passive and active measurement elements deployed in AT&T's commercial IP network. These measurement elements continuously feed data to a large, high-performance data repository at AT&T Labs. We have also described a number of studies carried out using the different forms of data in the repository. The variety of measurement data at our disposal, together with the ability to correlate multiple sets of data, enable new and important characterizations of many aspects of the Internet.

References

[1] J. Apisdorf, K. Claffy, K. Thompson, and R. Wilder, "OC3MON: Flexible, Afford-

able, High Performance Statistics Collection," `http://www.nlanr.net/NA/Oc3mon`

[2] J. E. van der Merwe, R. Cáceres, Y. Chu, and C. J. Sreenan, "mmdump - A Tool for Monitoring Internet Multimedia Traffic," *AT&T Labs–Research TR 00.2.1*, February 2000.

[3] Cisco NetFlow; for further information see `http://www.cisco.com/warp/public/732/netflow/index.html`

[4] Cooperative Association for Internet Data Analysis, "Internet Measurement Efforts," `http://www.caida.org/Tools/taxonomy.html#InternetMeasurement`

[5] A. Feldmann, "Continuous online extraction of HTTP traces from packet traces", *Proc. W3C Web Characterization Group Workshop*, November 1998.

[6] A. Feldmann, J. Rexford, and R. Cáceres, "Efficient Policies for Carrying Web Traffic over Flow-Switched Networks," *IEEE/ACM Transactions on Networking*, Vol. 6, No. 6, December 1998.

[7] A. Feldmann, R. Cáceres, F. Douglis, G. Glass, and M. Rabinovich, "Performance of Web Proxy Caching in Heterogeneous Bandwidth Environments," *Proc. of IEEE Infocom '99*, March 1999.

[8] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, "NetScope: Traffic engineering for IP Networks", IEEE Network Magazine, March 2000.

[9] R. Greer, "Daytona And The Fourth Generation Language Cymbal", *Proc. ACM SIGMOD '99*, June 1999. For further information see `http://www.research.att.com/projects/daytona`

[10] V. Jacobson, C. Leres, and S. McCanne, `tcpdump`, available at `ftp://ftp.ee.lbl.gov`, June 1989