

The Value of a Systematic Approach to Measurement and Analysis: An ISP Case Study

Srinivas Ramanathan, Ed Perry
Internet Systems and Applications Laboratory
HPL-98-171
October, 1998

E-mail: [srinivas,edp]@hpl.hp.com

internet performance,
service management

Internet Service Providers (ISPs) are looking to differentiate their services (e.g., Web, Email, News) by offering quality of service guarantees to their subscribers. However, the absence of effective proactive monitoring and the inability to diagnose problems in a timely manner have meant that ISPs are unable to offer guarantees for the availability and performance of their services. Instead, ISPs mainly offer network-level quality guarantees for availability and delay, that too restricted to a single administrative domain. In this paper, we describe a system for monitoring and diagnosis of ISP services that can be used by ISPs to offer service-level guarantees. Service availability and performance measures included in this system document conformance to service-level guarantees. Coupled with these high-level health reports are measures of infrastructure services such as the Domain Name Service (DNS) and Network File Service (NFS) that can be used for problem diagnosis. To further assist with diagnosis and to provide information for capacity planning, measures of server and network utilization are included in the monitoring system. Measures of the availability and usage of an ISP's Points of Presence (POPs) complete a comprehensive monitoring suite for ISPs. The integration of measurements into a service-oriented monitoring system, as described herein represents a significant advance in the state-of-the-practice for ISPs. We present experimental results obtained from a real-world ISP system that demonstrate the practical utility and effectiveness of the monitoring system in detecting and diagnosing problems in ISP environments.

Internal Accession Date Only

1 Introduction

Quality of Service (QoS), which refers predominantly to the availability and performance of Internet services, is fast emerging as a key differentiator for Internet Service Providers (ISPs) as they attempt to gain a dominant share of the Internet marketplace. In a recent PC Week study [1], 96% of respondents listed service availability as their main consideration in choosing ISPs. Furthermore, 93% of respondents perceived service performance as their second key expectation from ISPs. To meet the expectations of their subscribers and to attract new subscribers, ISPs must measure and manage the QoS of their service offerings. Doing so requires a shift in the management paradigm that ISP operations personnel are used to. From monitoring and managing network links in terms of metrics such as packet loss and delay, using tools such as *ping* and *traceroute*, ISP operations personnel must begin to look at a totally new set of metrics that represent the QoS of services from the perspective of subscribers. Examples of such metrics include accessibility of a specific Usenet newsgroup, the delay between the transmission and the reception of an Email message, the response time for retrieval of a Web page, etc.

More than ever before, in order to meet subscriber expectations, it is also becoming increasingly important for ISPs to proactively monitor their systems to detect problems in advance of subscriber complaints. Furthermore, accurate yet rapid diagnosis of the detected problems is also essential to minimize subscriber-perceived downtimes of the services. All of these requirements warrant a new generation of service quality measurement and problem diagnosis technologies.

In this paper, we describe a system that is targeted at enabling ISPs to measure and manage the quality of services offered by ISPs to residential subscribers. A heterogeneous set of measurements included in this system enable end-to-end, top-to-bottom monitoring of ISP network, server, and infrastructure components. Measurements that emulate typical subscriber accesses to Web, Email, and News services enable ISPs to assess the quality of service they offer to subscribers. Coupled to these high-level health reports are measures of infrastructure services such as the Domain Name Service (DNS) and the Network File Service (NFS) that can be used for problem diagnosis. To further assist with diagnosis and to provide information for capacity planning, measures of server and network utilization are included in the monitoring system. Measures of the availability and usage of an ISP's Points of Presence (POPs) complete a comprehensive suite for ISPs. A web-based graphical interface provides customized views of the measurement results to ISP operations personnel and capacity planners. We present experimental results from an ISP system that demonstrate the practical utility and effectiveness of the monitoring system in detecting and diagnosing problems in real-world ISP environments.

The integration of measurements into a service-oriented monitoring system, as described herein, represents a significant advance in the state-of-the-practice for ISPs. Experimental results demonstrating the effectiveness of the monitoring system in detecting and diagnosing real-world problems are testimony to the advantages of a systematic approach to measurement and analysis. The rest of this paper is organized as follows: To set the context, in Section 2, we first describe the typical architecture of an ISP offering Internet services to residential subscribers. Section 3 introduces the architecture of the monitoring system. Details of the measurements implemented in the monitoring system are described in Section 4. Section 5 presents experimental results from a real-world deployment of the monitoring system.

Section 6 contrasts this work with earlier efforts at designing monitoring systems for Internet services and infrastructure. Finally, Section 7 summarizes the lessons learnt from our work, and outlines area for further research.

2 ISP System Architecture

There are two key components of an ISP system that supports modem-based dial-in access to residential subscribers: the Points of Presence (POPs) and a server farm. These two components are elaborated next.

- **Points of Presence (POPs):** In order to allow subscribers to connect to the Internet, an ISP establishes strategically located sites that house equipment to which subscribers connect through the telephone network using dial-in modems. These locations are referred to as *Points of Presence* (POPs). The key components of the POP are (Figure 1):

- **Modem Banks:** A collection of modems is used to handle incoming subscriber calls.
- **Telco connections:** To handle incoming subscriber calls, the modems in a POP must be connected to the telephone company central office. In the older POP designs, the ISP leased individual phone lines, each of which was associated with a modem. In newer designs, ISPs lease collections of phone lines from the telco. For instance, an ISP may lease a T1 line comprising of 24 phone lines. The modems are arranged in racks, each of which is capable of demultiplexing the leased line into its constituent individual phone lines, which are then handed off to individual modems in the rack. The ISP assigns one call-in phone number for subscribers and software at the telephone company central office switch provides a “hunt group” which directs each call to the next available modem.
- **Terminal Servers:** Each group of modems in a POP is associated with a terminal server, which is a device that handles communication of packets from the POP to other locations in the ISP system.

A subscriber dialing-in to a POP is connected via a modem in the POP to a terminal server that prompts the subscriber for a login name and password. Based on this information, the terminal server contacts back-end authentication servers to perform subscriber authentication. Once authentication succeeds, the terminal server completes the establishment of a dial-in connection for the subscriber. During this process, the terminal server assigns an IP

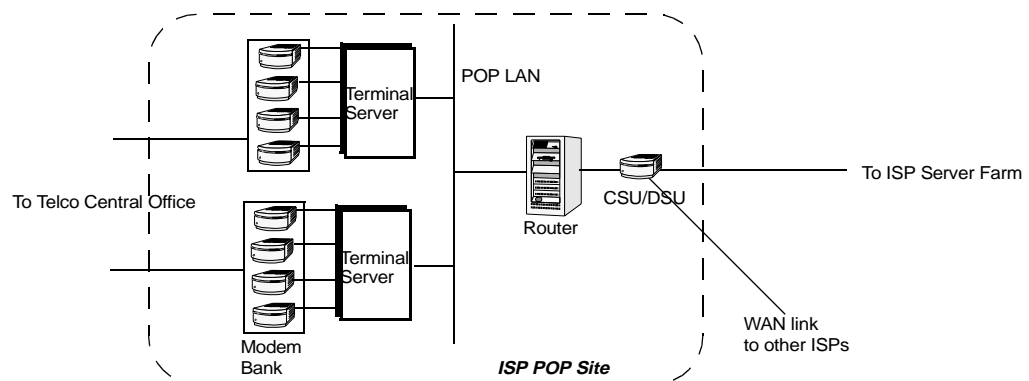


Figure 1. A typical dial-up POP architecture

address to the subscriber PC, thereby permitting communication between the subscriber PC and other servers on the Internet. Access routers and CSU/DSUs (Channel Service Unit/Data Service Unit) enable communications between the POP and other locations on the Internet over wide area network links.

- **ISP Server Farm:** A server farm is a location that houses servers that support applications such as Web, Email, and News that subscribers access. Infrastructure services such as authentication during login, domain name service (DNS), Network file service (NFS), etc., are also offered via servers located at this site. For high performance, typically, dedicated servers are used to support the different applications. Moreover, to ensure the scalability of their services, ISPs often replicate Web, Email, and News content (e.g., web pages, subscriber mailboxes, news articles) across a number of servers. A common technique for balancing load among these servers uses the round-robin scheduling capabilities of the DNS service [2]. As per this technique, replicated servers that support the same service (Web, Email, or News) are grouped together and assigned a single domain name in the DNS database. A DNS server that receives a request to translate the domain name to an IP address returns the IP address of one of the servers that is associated with the domain name. The DNS server uses a round-robin scheme to choose the IP address that it must return [2]. Since subscribers often access services using domain names rather than using IP addresses, an ISP's use of a group of replicated servers is invisible to subscribers. We use the term *front-end servers (FESs)* to refer to the different Web, Email, and New servers that directly handle subscriber requests.

While the Web, Email, and News content can be stored locally on the FESs, many ISPs prefer to store the content on distinct back-end content servers. The FESs access content from the content servers using mechanisms such as the Network File System (NFS). This architecture off-loads the task of content retrieval from the FESs. Moreover, a highly redundant architecture for the NFS content servers ensures high availability of the ISP's services.

3 Monitoring an ISP System

3.1 Motivation

There are three main motivations for monitoring an ISP system:

- *Proactive operational maintenance:* Monitoring the quality of service being offered to subscribers can enable problems to be detected and corrected in advance of subscriber complaints.
- *Customer support:* Problems detected by the monitoring system can enable customer support representatives who handle subscriber calls to filter calls that relate to known problems in the ISP system at the earliest.
- *Capacity planning:* On-going monitoring can yield baselines of performance and usage. Long-term analysis of the measurement results can indicate performance and usage trends that can form the basis for capacity planning.

The monitoring system described in this paper mainly addresses the needs of ISP operations personnel and capacity planners. While the information made available by this system can be made accessible to customer support personnel for filtering subscriber calls, an in-depth study of customer support is beyond the scope of this paper.

3.2 Measurement Technologies

Measurements are fundamental to monitoring and management. There are two predominant classes of measurements that can be employed in a monitoring system:

- *Active measurements:* Measurements in this class explicitly stimulate traffic to networks, servers, and service applications, in order to assess various metrics of interest about these elements. In order to obtain measurements that are representative of the real world, the synthetic workload generated by the active measurements must be determined based on observations of subscriber workloads obtained while the system is in operation. Furthermore, to truly reflect subscriber-perceived problem artifacts, the active measurements must use the same network path and the same set of infrastructure services (DNS, NFS) that subscribers use.

The main advantage of active measurements is that they do not rely on any monitoring capabilities built into the networks, servers, and service applications. On the other hand, since they introduce additional traffic, if they are not carefully engineered, these measurements can impact the QoS of services they are intended to measure! Moreover, since the measurements assess the ISP system and services from an external viewpoint, these measurement results do not provide sufficient information about usage of the services by subscribers and usage of resources by service applications to enable detailed problem diagnosis and capacity planning. Consequently, the utility of these measurements is limited to service quality assessment and preliminary diagnosis of problems.

- *Passive measurements:* As the name suggests, measurements in this class utilize instrumentation built into network, server, and application components. Since they do not require the explicit stimulation of traffic, these measurements have minimal impact on ISP systems. Consequently, they also scale better for larger ISP systems. Moreover, since they use built-in instrumentation that tracks real subscriber traffic, the measurements truly reflect the QoS perceived by subscribers. Besides being useful for service quality assessment, passive measurements can also provide information about resource and service usage, which is critical for detailed problem diagnosis and capacity planning. Passive measurements can be obtained from different sources:
 - *From the servers:* Web, Email, and News services are based on the reliable Transmission Control Protocol (TCP) that involves the explicit communication of acknowledgments back from subscriber client applications back to the servers. Based on acknowledgment feedback from the clients, the service applications can not only detect the start of different service transactions from clients, but also the end of these transactions. As part of their normal operation, the service applications can measure the service performance in terms of response times for different service transactions. This information can be stored in log files for further analysis by management applications. Alternatively, the information can be exposed via management information bases (MIBs). Resource and service usage information can also be obtained from the servers.
 - *From subscriber clients:* Monitoring capabilities built into subscriber client applications, or available using special-purpose instrumentation software such as *Net.Medic* [3] can provide an assessment of service quality.
 - *From special-purpose network probes:* External network probes can be used to snoop network transmissions and analyze the captured packets to yield measurements of services and networks.

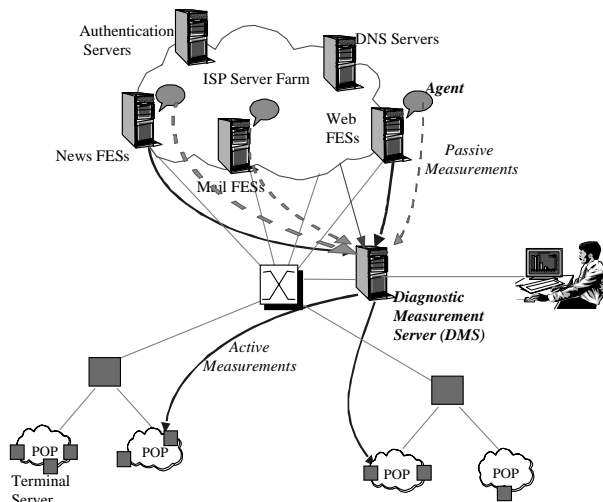


Figure 2. Components of the ISP monitoring system: deployment of a Diagnostic Measurement Server (DMS) and measurement agents in an ISP Server farm.

While the problems associated with the distribution of special-purpose software make client-side passive measurements less attractive to ISPs, the complexity and cost of implementation make the network probe-based approach less viable for service quality assessments. Since a great majority of service applications (Web, Email, and News servers) do not have sufficient instrumentation capabilities built in to enable the passive assessment of service quality, our ISP monitoring system incorporates active measurements for service quality assessment. However, the ISP monitoring system also includes passive measurements from the server, since these measurements can provide a wealth of information for detailed problem diagnosis and capacity planning.

3.3 Monitoring System Architecture

Figure 2 illustrates the components of the ISP monitoring system. In our prototype implementation, a diagnostic measurement server (DMS) serves as a host for agents that make active measurements of service quality. Software agents operating on the ISP servers track resource and service usage passively. The measurements performed by all of the agents are specified via configuration files stored locally on the hosts on which the agents execute. The measurement results are transferred from the agents to the DMS, aggregated for storage in a common data repository, and made available for analysis and display to ISP personnel for operations, subscriber support, and capacity planning. Data from the agents can either be pulled by the measurement server or pushed by agents. Mechanisms such as FTP, *rcp* (remote copy), and *scp* (secure copy) are supported.

The format of the results of each of the measurements is represented in a *data dictionary* that is stored on the DMS. Using the data dictionary's contents, a *view generator* module enables analysis and correlations of the measurement results stored in the data repository. The measurement results can be analyzed based on time of day, based on specific ISP servers, or for specific services. For display of the measurement results, the prototype includes a *view dictionary* that enables a user to configure views of interest (e.g., indicating which subset of the measurements must be displayed together on the same time-scale). From a web browser, a user of the prototype can choose one of

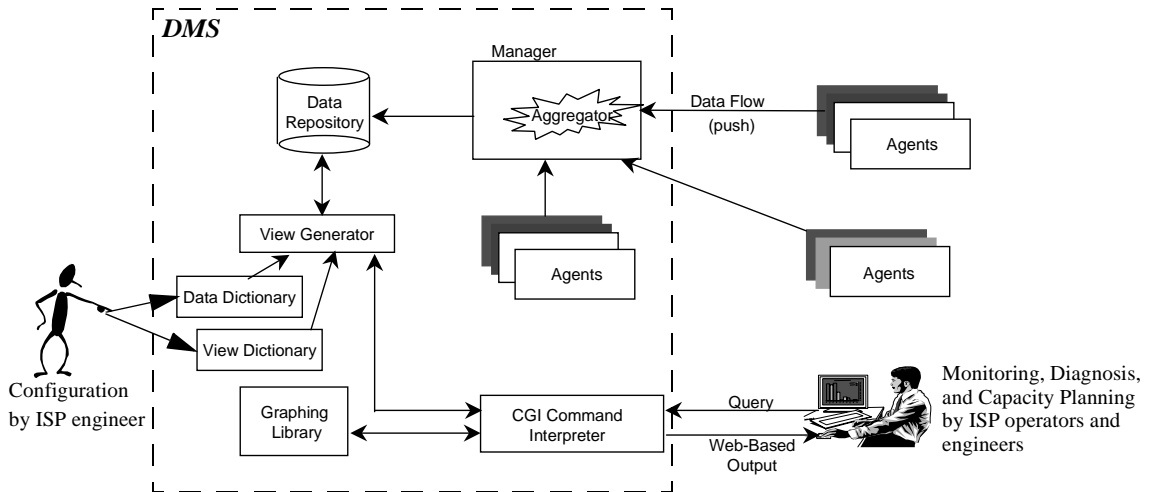


Figure 3. Architecture of the diagnostic measurement server and its interaction with measurement agents.

the views defined in the view dictionary and specify the criteria for analyzing the measurement results stored in the data repository (e.g., parameters such as the time-period of relevance and identities of the target(s) etc.). Using the common gateway interface, the selection criteria are passed to the view generator for analysis and the results displayed using a custom graphing library in graphical or tabular form via a web browser. To enable problem reports to be proactively emitted to ISP operations personnel, the monitoring system also includes threshold monitoring capabilities. SNMP alerts can be issued to existing ISP network operations consoles.

A unique feature of this system is the ability it offers for ISPs to incorporate custom measurements and views. To integrate custom measurements that an ISP may already be making (e.g., BGP routing flap measures, router utilization measures, etc.) into the monitoring system, an ISP must develop the corresponding measurement agent(s) and define the output schema in the data dictionary. The ISP can then use a high-level scripting language supported by the view dictionary to create custom views of the measurement results. The next section describes the measurements incorporated in the ISP monitoring system in more detail.

4 Details of Measurements

4.1 Dial-Up Availability Measurements

A great majority of operational problems that ISPs have faced as they attempt to scale their systems relate to their dial-up POP sites. Consequently, measurement and management of the dial-up POP sites is essential for ISPs to ensure that subscribers can connect to their systems at all times. The major questions of interest to ISPs with respect to their POP sites are:

- *POP Utilization:* Are there times when any of the POPs are operating at capacity or close to capacity? That is, how frequently are all the modem lines in a POP simultaneously in use?

- *POP Status*: Is each POP site functioning normally, i.e., are all terminal servers and modem lines working? Is each POP site connected to the ISP's internal network and to the Internet? Are the terminal servers able to authenticate new connections?

SNMP MIBs supported by the terminal servers in the POPs are the basic source of information about the POPs. From a terminal server's MIB, it is possible to determine the instantaneous state (e.g., active or idle) of each of the modem lines of the terminal server, and a cumulative number of dial-in connections to each modem line. By tracking this information periodically for all the terminal servers in all the POP sites, the ISP monitoring system determines the utilization of each of the POPs. Usage trends derived from this information serve to proactively alert the ISP of the likelihood that some of the POPs may be at capacity in the near future, and therefore require that additional modem line capacity to be provisioned. Furthermore, by comparing the rates of connections to different modem lines of a terminal server and across modem lines of different terminal servers, the monitoring system determines if all the modem lines and terminal servers are operational. The precise heuristics used to determine malfunctioning modem lines and terminal servers are described in Section 5 using illustrative examples.

4.2 Service Quality Measurements

To measure the QoS of ISP services, the ISP monitoring system includes the following measurements:

- *Web service measurements*: The monitoring agent for Web services uses active measurements to assess the availability and performance of Web services. Like a typical web client retrieving a web page, this agent first resolves the IP address of a target web server, establishes a TCP connection with the web server, and then issues a GET request to a specified static web page (see Figure 4). By interpreting the HTTP response header returned by the server, the web service agent determines the availability of the web service. In addition, the agent determines the overall response time for retrieving the web page as well as the components of the overall response time: the time for DNS resolution, TCP connection establishment time, server response time (which is the time between when the GET request is issued and the time when the HTTP response header is received), and the data transfer time (which is the time between when the HTTP response header is received and the web page retrieval is completed).

The response time components can reveal potential bottlenecks. Clearly, the characteristics of the network links (e.g.,

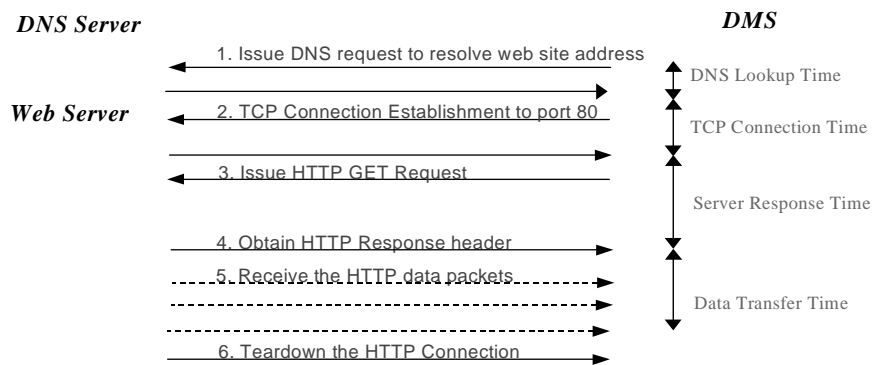


Figure 4. Operation of the web service measurement agent

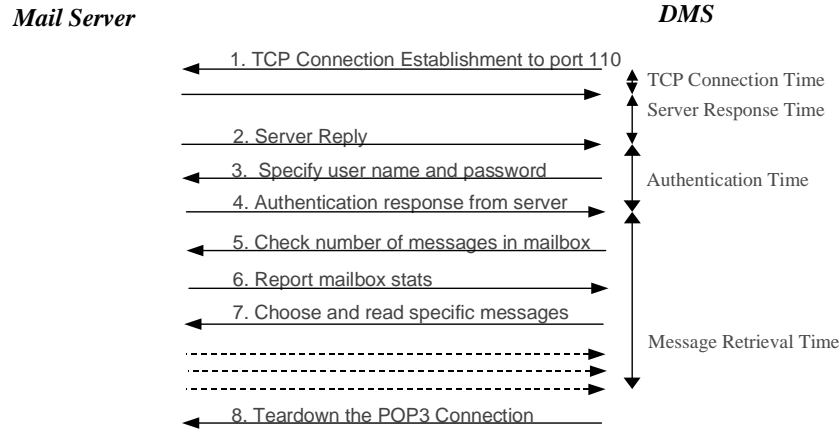


Figure 5. Operation of the mail retrieval measurement agent

delay and packet loss) interconnecting the measurement location (e.g., the DMS) to the target web server impact all of the response time components. Assuming little change in the intervening network's characteristics, an increase in the DNS response time is an indicator of DNS related problems. A large TCP connection establishment time is typically indicative of bottlenecks in the server host (note that although the TCP connection is established to the port that the web server application is listening to, TCP connection establishment is entirely handled in the kernel, rather than in the web server application). The server response time is indicative of the web server application processing delays. Included in this measure could be queuing delays at the server waiting for a web server application process to be allocated to handle the new request, delays for any DNS processing required by the server, and file system access times for accessing the requested web page.

- *Email service measurements:* The ability of a subscriber to send mail to other subscribers (of the same ISP or of a different ISP) is distinct from the ability of the subscriber to receive mail. To assess the availability and performance of these distinct Email operations, the monitoring system includes the following agents:
 - A mail delivery agent assesses the capabilities of an ISP's Email system to accept mail messages for delivery to one or more target destinations. Since the ISP system may use different mail servers to handle messages destined for local subscribers and those destined to external locations on the Internet, to verify the different mail delivery paths, different local and remote destinations can be specified as targets to the mail delivery agent. This agent uses the Simple Mail Transfer Protocol (SMTP) to communicate with one or more ISP mail servers, transmits the specified mail messages, and reports the availability and performance of the mail servers.
 - A mail retrieval agent emulates a subscriber, by accessing a test mailbox stored on the ISP's mail server. While doing so, this agent assesses the availability and performance of the mail retrieval components of the Email system independent of the mail delivery components. In addition, this agent can also measure the delay between the transmission of a mail message and its reception at the intended destination. Figure 5 depicts the operation of this agent which uses the Post Office Protocol Version 3 (POP3) for communication with an ISP's POP3-

based mail server. In its default mode of operation, this agent retrieves all the messages in a test mailbox. To obtain consistent measures over time, a constant sized mailbox must be used for retrieval. Like the web service agent, the mail retrieval agent tracks the overall response time for accessing the mailbox as well as the individual components of response time: TCP connection establishment time, time to authenticate the subscriber, and the time to retrieve the subscriber's mailbox. These response time components can be used to identify potential bottlenecks that may exist during retrieval of mail messages in the ISP system.

While retrieving messages from a mailbox, the mail retrieval agent can be instructed to selectively recognize messages with specific filter criteria (e.g., specific subject field in the mail headers) and to perform additional actions for these messages. For instance, the agent can be instructed to estimate the delay between the transmission and reception of the message by comparing the origination time of the mail message (included in the mail header) and the retrieval time of the message.

- *News service measurements:* Availability and performance of the news service is measured using an agent that operates in much the same way as the other active measurement agents described earlier (see Figure 6).

The ability of the Web, Email, and News service quality measurements to assess the subscriber-perceived QoS is dependent on the location of the respective measurement agents. In an extreme configuration, the agents can be located on a host in a Point of Presence (POP) site and can use a dial-up line to make measurements. In this case, the service quality measurements are representative of subscriber-perceived performance. However, locating host machines at the hundreds of POP sites of an ISP specifically for measurements is likely to be overly expensive. Moreover, simultaneous active measurements from a number of agents at the POP sites can also end-up congesting the servers they measure. At the other end of the spectrum of possibilities, the measurement agents can be centrally located, possibly in an ISP's server farm. A single set of measurements that assesses the service quality offered by the ISP's Web, Email, and News servers minimizes the overheads of active measurement. However, since in this case, the agents are located close to the ISP's servers, the service quality measurements they provide do not reflect the impact of the networks interconnecting the ISP's POP sites to the server farm. In fact, since there are different network paths connecting the different POP sites to the server farm, the QoS observed by subscribers is dependent

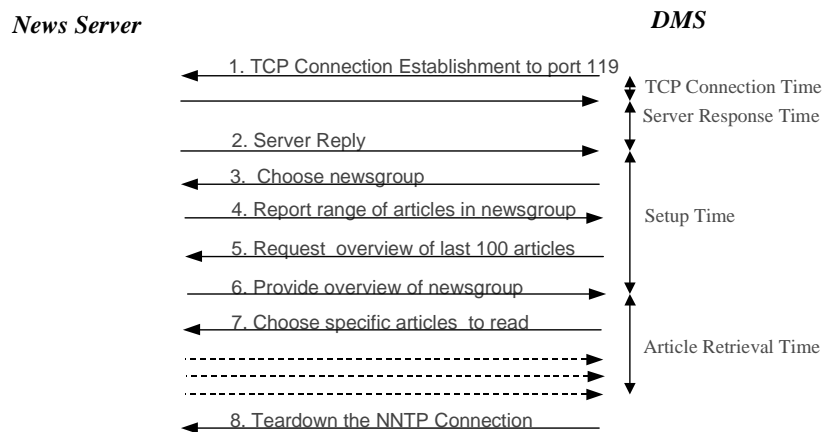


Figure 6. Operation of the news service measurement agent

on the POP sites that they use to connect to the ISP. To obtain an end-to-end perspective of service quality, the service quality measurements must be complemented with measurements of the network's effects on service quality.

4.3 Network Measurements

ISPs extensively use tools such as *ping* and *traceroute* to detect and diagnose network problems. The monitoring system incorporates ping's capabilities to assess network connectivity and provide end-to-end delay information that is useful for problem diagnosis. While ping and traceroute give an indication of changes in network characteristics, these tools are not sufficient to quantify the impact of these changes on the availability and performance subscriber-visible services. For services that use the Transmission Control Protocol (TCP) for reliable communication, a useful network performance metric is throughput defined as the rate of reliable transmission of packets between a source and a destination. The throughput achievable between any source-destination pair is a complex function of several factors such as the socket buffer size in use, the characteristics of the source and destination's TCP implementation, processing capabilities of the source and destination, the burstiness of packet losses (if any), round-trip packet transmission delays, etc. The complexity of this relationship makes it almost impossible to estimate the throughput achievable based on packet loss and delay measurements available from ping.

Many public domain tools such as *throughput TCP (tcp)* and *netperf* [4] have been widely used for measuring throughput. A key drawback of these tools is the need for custom software applications to be executed at the source and destination in order to enable the measurement. To enable throughput measurements without requiring special-purpose instrumentation at the targets, the monitoring system's throughput monitoring agent builds on the concept of the *Traceroute Reno (Treno)* tool [5]. Treno emulates a TCP-based data transfer using User Datagram Protocol (UDP) packets to transmit data to a target. The UDP packets are sized equivalent to typical TCP packets and are directed at a non-existent port on the target. Any IP addressable target responds to such packets by transmitting an ICMP error message that is almost equivalent in size to the TCP acknowledgment packets that a subscriber client machine transmits in response to a TCP data packet. By incorporating TCP's delayed acknowledgment, window opening, and retransmission schemes in the server application, Treno emulates data transfers over a TCP connection. In its basic form, Treno was intended to permit subscribers to compare the performance offered by different network providers, and for network providers to monitor the performance of their networks.

In an effort to measure the best possible throughput achievable when using TCP as the transport protocol over an IP network, Treno continues to open its transmission window, which represents the amount of data for which acknowledgments are outstanding at any time, until it encounters a packet loss. In the extreme scenario, Treno can continue to grow its transmission window to an extent that its packet transmissions flood the network for the duration of the test. Moreover, in practice, the TCP stacks on subscriber PCs are configured with a default maximum window size that restricts the amount of data transmitted simultaneously. Since Treno has no window size restriction built in and since the throughput achieved over a TCP connection is known to be proportional to the window size used, Treno can over-estimate the throughput achieved by a subscriber during a data transmission.

To overcome these limitations, a throughput monitoring agent included in our monitoring system extends the concepts of Treno by building in restrictions on the maximum window size. By doing so, the throughput monitoring

agent ensures that it does not flood the network. Furthermore, by restricting the amount of data transferred during each measurement to match typical data transfer sizes that subscribers use when retrieving Web, Email, and News content, the throughput monitoring agent ensures that its measurements reflect subscriber perceptions of network throughput. Since it does not rely on custom software in the targets for throughput measurements, the throughput monitoring agent can measure throughput from its source location to any IP-capable device, such as application servers, routers, terminal servers in the POP sites, or even subscriber PCs and terminals. An approximation of end-to-end service quality can be obtained by combining the Web, Email, and News service quality measurements with throughput measurements directed to subscriber PCs and terminals. However, such a measurement is most likely to be limited by the bandwidth bottleneck on the dial-up lines. A more useful metric for ISPs to track is the throughput achievable during transmission from the server farm to the different POP sites. The terminal servers at the POP sites (see Figure 2) can serve as targets for such a measurement. A reduction in throughput is usually a good early indicator of problems that may impact subscriber-perceived QoS. In the event that a reduction in network throughput is detected, ping and traceroute can be used to isolate the root-cause of the problem.

4.4 Infrastructure Service Measurements

Although they are not directly visible to subscribers, infrastructure services such as the Domain Name Service that handles translation of host names to IP addresses and the Network File Service that enables distributed data access, play a crucial role in the operation of Internet services. Any deterioration in availability and performance of these infrastructure services usually manifests as subscriber-visible problems with one or more of the Internet services. To monitor these infrastructure services, our monitoring system includes the following measurements:

- **DNS measurements:** Truly accurate measurements that assess DNS performance as perceived by subscribers can only be obtained using passive measurements that track address resolutions requested by subscriber applications. Such measurements can be made either at the DNS server (requiring modifications to the server application), at DNS clients (e.g., some web proxy servers log DNS response time as one of the components of the overall response time for subscriber requests), or using non-intrusive external probes. In the absence of such passive monitoring capabilities in most ISP systems, our monitoring system uses active measurements to emulate typical DNS requests and assess DNS availability and performance.

Since the ability of a DNS server to service requests from its cache is independent of the ability of the server to resolve mappings not in the cache, the DNS agent includes separate measures of DNS cache hits and cache misses for a server. By issuing a non-recursive query to a DNS server, this agent forces the server to respond based on its local state alone, and thereby emulates a DNS cache hit at the server. By requesting address resolution to a well-known Internet site and by observing the status returned by the DNS server in its response, the DNS agent measures the availability of the DNS server. To measure the performance of DNS cache misses, the DNS agent issues a request for address resolution for a randomly chosen host. The response time for cache misses are noted for those responses that are detected as not being returned from the DNS server's cache (this can be detected by checking the DNS response header). Comparison of DNS response times with the network response times provides an indication of whether the network or the DNS server is a performance bottleneck.

- **NFS measurements:** As explained earlier, many ISPs use back-end NFS servers to store Web, Email, and News content. NFS monitoring experience developed over the years in enterprise networks can be directly applied to assess the health of the NFS sub-system. There is a richer set of NFS metrics that can be collected using passive measurements made at the client end as compared to the server end. On the server end, the number of NFS calls handled by the server over time can be an indicator of potential overload conditions. On the client end, the rate of NFS calls issued by the client, the percentage of duplicate responses received from the server(s), the percentage of time-outs observed by the client, and the percentage of retransmissions can all be tracked over time to detect NFS anomalies. A rule of thumb often used is that during normal operation, the percentage of retransmissions and duplicate responses observed by the client should be much less than 5%. When this is not the case, a comparison of the percentage of retransmissions with the duplicate responses is an indicator of the cause of the problem. If very few duplicate responses have been received, the interconnecting network is potentially lossy. On the other hand, if the percentage of duplicate responses is almost equal in magnitude to the percentage of retransmissions, this implicates a slow-down at the server end as the cause of the problem (for example, because there are too few NFS server (*nfsd*) processes executing on the server). A look at the client statistics of the number of NFS requests waiting for service can indicate client-side bottlenecks. For example, if there are too few NFS client (*biod*) processes executing on the client to handle remote NFS accesses, the count of waiting requests grows.

4.5 Server Measurements

Measurements taken on ISP servers can yield a wealth of information that is critically important for problem isolation. Our monitoring system's agents executing on ISP servers passively track various statistics of interest. Among the measurements obtained using these agents are CPU utilization of the server, free memory available, virtual memory page scanning rate, and the packet transmission and reception rate through each of the interfaces of the server. For servers that support TCP-based services, the rate of connections to and from the server and the number of currently established connections on the server are useful measures of the server's workload. A breakdown of connections based on the specific ports to which they relate is also available, thereby enabling workload to be determined for each service (SMTP, POP3, HTTP, etc.), even in cases when the same physical host supports multiple services. By monitoring the different states of the TCP connections to a server, it is possible to detect abnormal conditions such as the malfunctioning of specific servers or the occurrence of SYN attacks [6].

5 Experimental Evaluation

This section presents experimental results from a deployment of the monitoring system in a real-world ISP system.

5.1 Dial-Up Availability Monitoring

The ISP system serving as a testbed for deployment of our monitoring system has over two hundred POP sites, located in different US cities. Each POP site is equipped with a number of terminal servers, each of which is associated with a modem bank of 16 modem lines. To assess the state of the POPs, SNMP MIB variables that represent the state of each of the modem lines and the number of connections serviced by each modem line were continu-

ously monitored via the terminal servers. The polling period was set to 30 mins, to enable monitoring of each of the ISP’s two hundred POPs without generating excessive load on the ISP network.

As indicated earlier, the primary goal of a dial-up ISP is to ensure that subscribers dialing in to a POP site do not experience busy signals. To achieve this goal, the ISP must ensure that the utilization of a POP site, which is represented by the percentage of modem lines at a POP site that are simultaneously active, is well below 100%. A modem line is deemed to be active during a polling period either if the modem was in use when it was last polled or if the modem had received any new connections during the last polling period. A comparison of the utilization of different POP sites reveals the POP sites that are in greatest danger of running out of capacity. Figure 7 shows a summary of the percentage of active modem lines for the ISP’s POP sites that have peak utilization above 80% during a 24 hour period¹. Figure 8, which provides details of the variation in active lines with time of day for POP_M, indicates an urgent need to increase the number of available modems in this POP site.

Even when a POP is operating below capacity, subscribers may experience busy signals when one or more of the modem lines in the POP begin to malfunction. Failures of this form are a major irritant to ISPs since there are usually tens of thousands of modem lines in an ISP system and it is not feasible to individually monitor and test each of the modem lines. Sometimes, terminal servers may fail too, making all of their associated modem lines unavailable. The key idea that we have used to discover failed modem lines and terminal servers is to exploit the knowledge of the scheduling discipline that is used to allocate incoming subscriber calls to terminal servers and modem lines. An ISP may choose to employ one of many scheduling algorithms to determine how to assign an incoming subscriber call to one of the terminal servers in a POP site. In our testbed network, the POP sites have modem banks that are external to the terminal servers. These modem banks themselves do not decide the association of modem lines to subscriber calls. Instead, the scheduling policy for subscriber calls is set by the telco’s hunt-group configuration. Observing the various POP sites in our testbed network, we noticed that there are two predominant

POP site	% Active Lines [Max]	%Active Lines [Median]	Connections/hr [Max]	Connections/hr [Median]
POP _M	100.00	98.95	367	152
POP _W	100.00	94.44	69	36
POP _P	99.48	88.36	280	191
POP _N	98.96	98.18	298	196
POP _K	96.87	58.33	228	63
POP _H	93.05	73.27	319	131
POP _G	93.05	40.10	172	32
POP _J	91.70	54.86	76	32
POP _S	87.50	47.22	56	23
POP _F	85.42	29.18	59	11
POP _B	84.37	47.22	41	22
POP _T	80.55	47.40	99	48

Figure 7. Activity Summary of an ISP’s POP sites during a 24-hour period.

1. Because of the proprietary nature of the data, the true locations of the POP sites are not divulged here. The IP addresses of the terminal servers and other ISP host machines have been re-mapped to private addresses for this purpose.

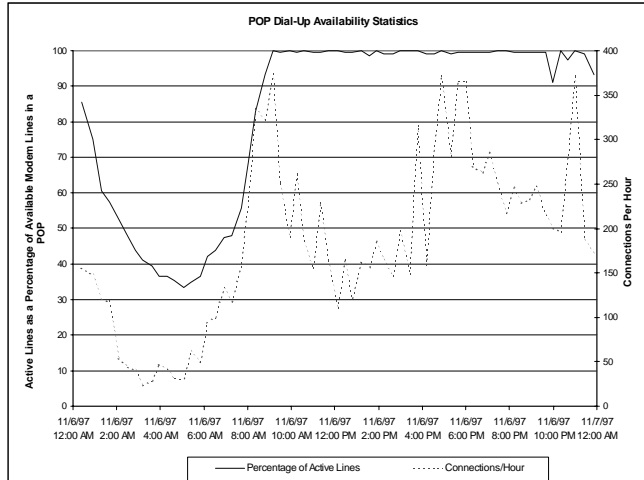


Figure 8. Activity of POP_M during a 24 hour period. Note that this POP site is operating at close to 100% capacity for most of the day.

policies in use. In a sequential scheduling policy, the terminal servers in a POP site are numbered sequentially and all incoming subscriber calls are forwarded to the first terminal server in the sequence that has an idle modem line. Alternatively, incoming calls can be allocated to terminal servers on a round-robin basis.

For POPs that use sequential scheduling, it is fairly straightforward to detect connection problems to specific terminal servers or to specific modem lines of certain terminal servers. Figure 9 shows how non-responsive modem lines of a terminal server can be detected by considering the maximum percentage of active modem lines associated with the terminal server at any instant of time. Figure 9(a) shows a scenario in which all terminal servers and their modem lines are functioning correctly. In this case, incoming subscriber calls are first assigned to the first terminal server until all its modem lines are active, and only then are any calls assigned to the second terminal server. Figure 9(b) shows the case in which the second terminal server (10.138.70.4) of POP_M has received calls although the peak modem line usage of the

POP Site	Terminal Server	% Active Lines [Max]	% Active Lines [Median]	Number of Connections
POP _D	10.133.38.3	100.00	100.00	499
POP _D	10.133.38.4	100.00	25.00	188
POP _D	10.133.38.5	37.50	0.00	24
POP _D	10.133.38.6	0.00	0.00	0

(a)

POP Site	Terminal Server	% Active Lines [Max]	% Active Lines [Median]	Number of Connections
POP _M	10.133.70.3	93.75	93.75	390
POP _M	10.133.70.4	100.00	43.75	230
POP _M	10.133.70.5	43.75	6.25	20

(b)

Figure 9. Activity status comparison of two POP sites that implement sequential call scheduling: (a) shows the ideal distribution of calls. (b) shows a case when the second terminal server receives calls before the first terminal server is 100% active. This indicates the existence of a potentially faulty modem line in the first terminal server.

Modem Line	% Activity	Number of Connections
1	95.38	31
2	89.23	38
3	100.00	25
4	81.54	55
5	100.00	0
6	76.92	21
7	98.46	24
8	0.00	0
9	100.00	0
10	72.31	46
11	76.92	42
12	100.00	0
13	96.92	37
14	69.23	20
15	67.69	18
16	100.00	33

Faulty Modem Line →

Figure 10. Activity comparison of the individual modem lines of terminal server 10.138.70.3. Modem line 8 which has seen no connections is potentially faulty.

first terminal server (10.138.70.3) is not yet 100%. This potentially indicates that the first terminal server (10.138.70.3) has one or more modem lines that is not answering subscriber dialins. Since the monitoring system tracks the activity of each modem line and the number of connections serviced by each line, it can compare the activity and connection rates for all the modem lines of each terminal server. Such a comparison for terminal server 10.138.70.3 (see Figure 10) reveals that for a 24 hour period, while all the other lines were either active or were serving one or more calls, line 8 alone was idle throughout this period and did not serve any calls. Hence, line 8 potentially represents a broken modem. Modem lines 5 and 9 in Figure 10 may have been active throughout the day without servicing any new connections either because they were serving long-lived connections that were initiated during a prior period, or because these lines may have been operationally “busied out” because they had malfunctioned in the past.

Through the SNMP MIB variables that it monitors, the monitoring system also tracks changes in the connections serviced by each of the modem lines in a POP site over time, and uses this information to determine the rate of connections

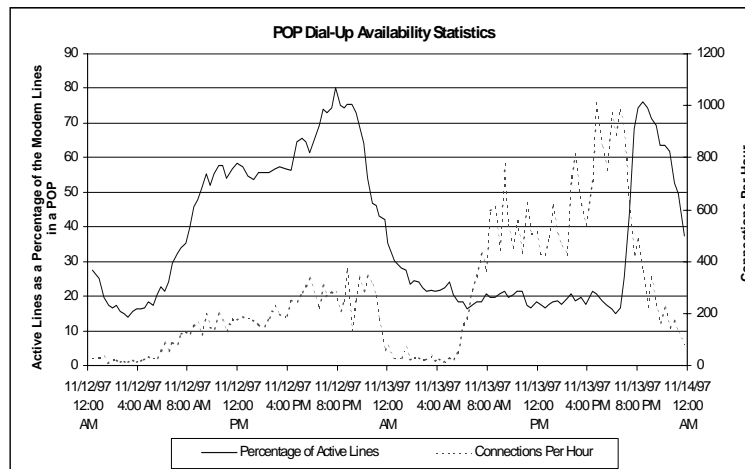


Figure 11. Dial-in statistics for a POP site for a 48 hr period. The sudden increase in connections to the POP site on the 13th was attributable to authentication problems experienced at the POP site.

being serviced by each POP site. In our testbed network, a clear time-of-day and day-of-week correlation in the rate of connections to a POP site was observed. Hence, a significant drop or a significant increase in the rate of connections as compared to an expected nominal rate of connections to a POP site is a good indication that subscribers are likely to be experiencing problems accessing the ISP's services via the POP site. Such changes in connection rate to a POP may be attributable to a variety of problems (e.g., modem line or terminal server faults, authentication problems, network inter-connection problems to the server farm, etc.). Figure 11 shows a case when the connection rate to a POP site increased significantly on 11/13 (as compared to the normal, indicated by the data for 11/12), when the ISP was experiencing problems authenticating subscribers from the POP sites. In their frustration at being unable to connect to the ISP's services, subscribers dialed-in more often, thereby causing over a four-fold increase in the connection rate.

5.2 News Problem Diagnosis

Next, we consider the diagnosis of a problem with the News service being offered over the testbed network. Typical response times for retrieving news articles amounting to 10KB of data from the FESs in the testbed network were in the range of 1-6 seconds. In the scenario under consideration, the response times for news access were observed to have increased dramatically to over 100 seconds at certain times of the day. Further analysis of the measurement results (see Figure 12(a)) revealed that one of the three FESs in use was performing significantly worse than the others. Restrictions on our access to this FES prevented further diagnosis of the problem.

The poor performance continued for several days, as the ISP's operations personnel failed to even detect that the problem existed, let alone fix the problem. Interestingly, a week later - on 8/22, the performance measured to all three FESs had degraded significantly (see Figure 12(b)). Analysis of the breakdown of the response time for one of the FESs (see Figure 13) indicated that on 8/22, while the set-up time dominated during the period from 7am to 10am, the server response time was the dominant component during the period from 10am onwards. The changes in the dominant response time components suggest that there were two different problems (with potentially different sources) that were

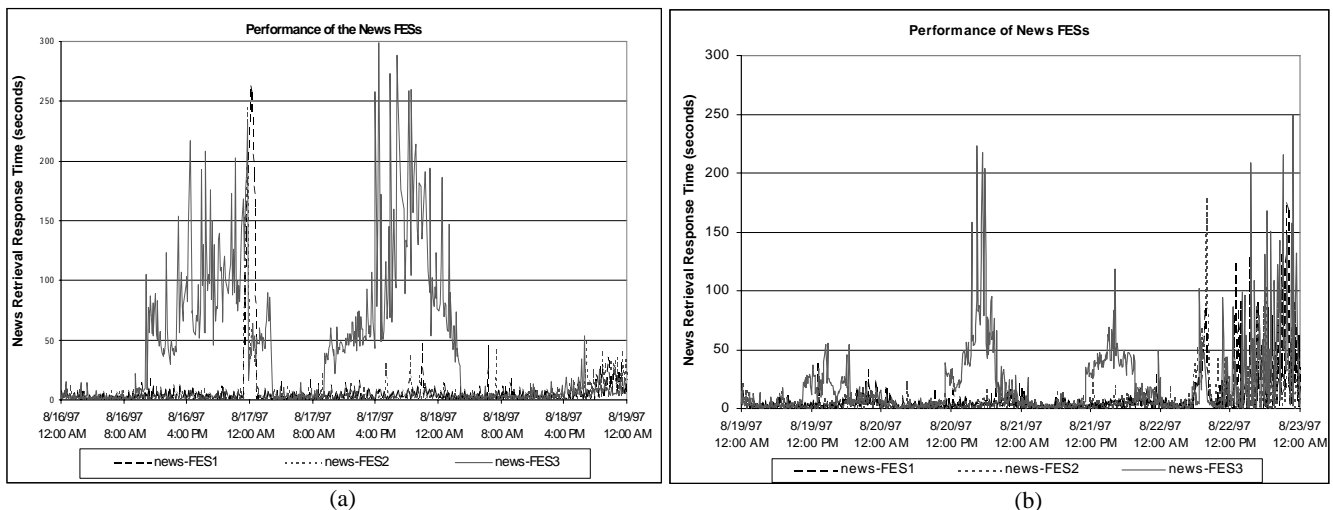


Figure 12. (a) Performance of the news service offered over the testbed network. Three front-end servers (FESs) are used for this service. (b) Performance of the three news FESs during a subsequent period. While news-FES3 performed poorly in the period 8/19-8/21, all three FESs experienced problems on 8/22.

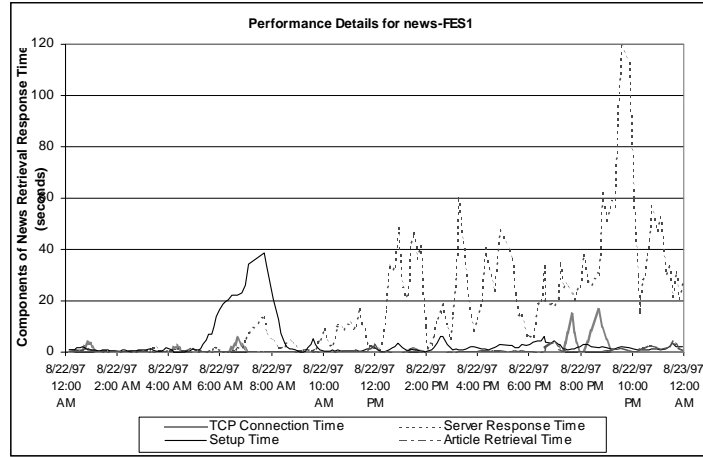


Figure 13. Detailed breakdown of the performance of news-FES1 on 8/22.

affecting news performance on this day. A similar behavior was observed for the other FESs as well, indicating that the same problems were affecting all the FESs.

Passive measurements available from agents operating on news-FES1 provided additional information essential for diagnosis. Comparison of the rates of TCP connections handled by news-FES1 on this day with baselines computed from historical data indicated that there had not been any increase in the number of connections to or from news-FES1. Analysis of the rates of packets to and from news-FES1's external network interface too did not indicate any change in subscriber-generated news workload. However, analysis of the rates of packets to and from news-FES1's internal interface revealed the source of the problem. Figure 14(a) depicts the packet traffic handled by news-FES1's internal interface over a week. There were two traffic peaks - one on 8/17 and the other on 8/22 - corresponding to an incoming traffic rate of over 45000 pkts/min, both of which corresponded to times when news perfor-

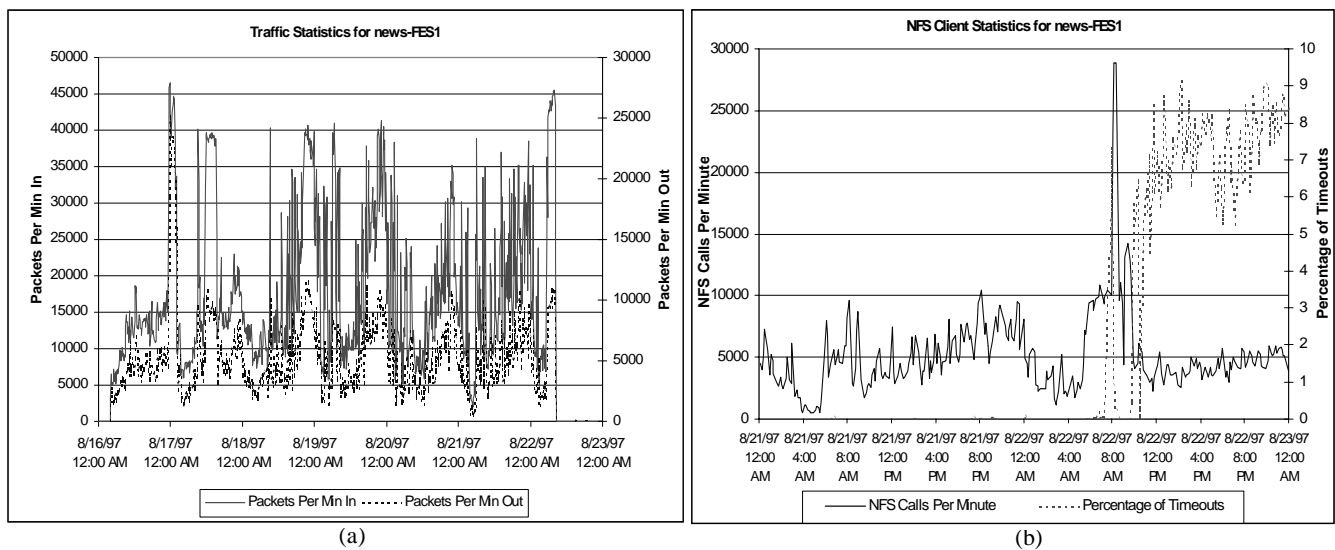


Figure 14. (a) Packet traffic handled by the internal interface of news-FES1. (b) Client measurements of NFS obtained from news-FES1.

mance observed using the news service agent had degraded. The second peak coincides with the 7am-10am time period on 8/22. The above analysis leads to the conclusion that the source of the first problem is congestion on news-FES1's internal network interface. Rerouting of traffic from a 10Mbps interface to a 100Mbps interface around this time removed the performance bottleneck.

The second problem observed on 8/22 from 10am onwards was attributable to NFS problems. Figure 14(b) compares the NFS statistics gathered from news-FES1 for 8/21 and 8/22. While there had not been a detectable increase in NFS calls per minute issued by news-FES1 to the NFS server, the percentage of NFS time-outs had increased dramatically from close to 0 to almost 10%. Comparison of network delays between news-FES1 and the NFS server did not indicate any change from normal. This leads to the conclusion that a slowdown of the NFS server was the cause of the problem. Since the same NFS server is used by all the news FESs, performance problems were observed when accessing each of the news FESs.

5.3 Email Performance Tuning

Figure 15 depicts the performance of mail delivery in the testbed network. In this case, the monitoring system was used to track the delay between the transmission of a message and its receipt at a local subscriber mail account. As is evident from the figure, starting with a mail delivery delay of about 10 mins on 8/27 when the measurement was deployed at this site, mail delivery performance of the ISP's mail system continued to degrade. By 9/11, mail delivery delay had reached close to 35 mins at peak times. This rapid increase in mail delivery delays was attributable to a failure of the mail queue processing algorithms implemented by the mail server application. As is noticeable from the figure, tuning of the mail server application, together with the introduction of new mail handling policies by the ISP (e.g., to queue mail messages for no more than 3 days), resulted in a significant drop in the mail delivery time noticeable to subscribers. By 9/18, mail delivery delays were about 5 mins. This scenario is representative of the utility of the ISP monitoring system. Besides providing measurements that can enable ISPs to assess the quality of services they offer, the monitoring system provides a means by which ISP operation staff can directly observe the impact of performance tuning and system reconfigurations that they perform.

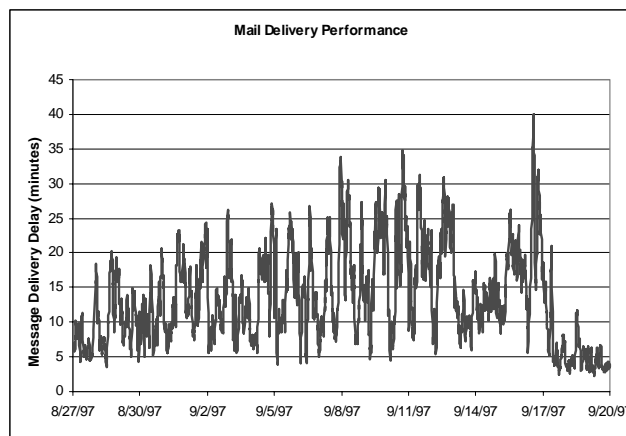


Figure 15. Performance of the testbed network's mail delivery system.

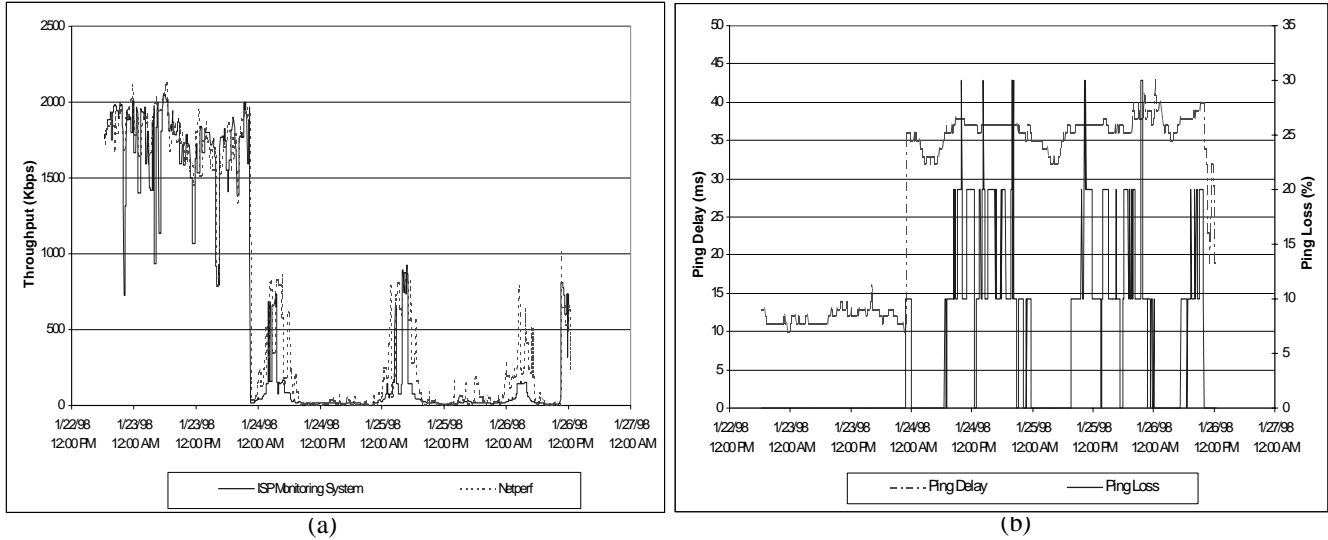


Figure 16. (a) Comparison of throughput measurements made by the ISP monitoring system with those obtained using *netperf*. (b) Packet loss and delay measured during the same period.

5.4 Network Monitoring

To evaluate the integrity of the network throughput measurement that the ISP monitoring system incorporates, we compared the throughput measured by the throughput monitor with that achieved over a TCP connection using the *netperf* public domain tool, for the same source-destination pair. The data transfer size was set to 40Kbytes, and the TCP socket buffer size used was 8Kbytes. Packet loss and round-trip delay measurements were made at the same time for correlation with the throughput measurements. Figure 16 compares the results of throughput measurements obtained using *netperf* and using the ISP monitoring system. In order to evaluate the suitability of the ISP monitoring system's throughput measurement as a service-level metric that can be used to offer a reasonable minimum guarantee of performance, the 90 percentile values that represent the throughput that is achieved at least 90% of the time were chosen for comparison. The percentile values were computed over a 30 min time window.

Figure 16(a) illustrates a remarkable degree of correlation between the throughput measurements observed using *netperf* and that observed using the ISP monitoring system. While the instantaneous measurements (not shown in Figure 16) can vary because of the dynamic nature of network phenomena, the 90 percentile values of throughput show a 97% correlation. Figure 16(b), which shows the variation in network round-trip delays and packet loss values during the same time period, serves to explain the reasons for the different drops in throughput. The drastic drop in throughput around 1/24 happens around the same time that network delay increases by about three-fold, because of a routing change in the network infrastructure. Subsequent to this period, there are three valleys in the throughput graphs that map directly to periods of high packet loss. The three peaks around 800 Kbps correspond to periods of high network delays with no packet loss.

6 Comparison to Related Work

Conventional network management systems (HP OpenView, IBM Tivoli, CA Unicenter, etc.) provide very limited capabilities to assist an ISP in detection and diagnosis of problems relating to Internet services. Over the years, several measurement tools that have been developed and made available in the public domain for monitoring IP networks and servers. Tools such as *ping* and *traceroute* have been used for network monitoring and diagnosis, and those such as *top*, *vmstat*, and *iostat* have been used for monitoring servers. The lack of an integrated, easy-to-use monitoring system has meant that ISP operations personnel spend a significant amount of time manually troubleshooting problems.

The Virtual Adrian toolkit from Sun Microsystems [7] aims to simplify the task of problem diagnosis for ISPs by providing a variety of statistics relating to the performance of their servers. The focus on management of services, rather than just management of servers, is one of the distinguishing characteristics of our monitoring system.

At the other extreme of the spectrum of products is the *WhatsUp* monitoring solution [8]. Using only active measurements, WhatsUp tracks the availability and performance of various services. However, its measurements are limited in functionality. Each measurement attempts to establish connections to the respective TCP ports and issue a single service-specific command. This may not represent typical subscriber accesses to the ISP services. Consequently, WhatsUp's measurements are not sufficient to assess the service quality of ISP services. Moreover, since it predominantly uses active measurements, WhatsUp does not provide sufficient information for problem diagnosis.

One of the early systems to provide an integrated solution for ISPs is the network operations console, NOCOL [9]. Using a distributed architecture that includes agents for different types of measurements, NOCOL provides information essential for problem detection and diagnosis. The lack of an easy-to-use user interface, and the absence of sufficient service-specific measurements (NOCOL tests service availability by merely connecting to the service port; Additional diagnostic information as well as service performance measurements are lacking) has been a major deterrent in its widespread adoption. The ISP monitoring system described in this paper uses a similar manager-agent architecture as NOCOL. The main novelty of this system is the support it offers for problem identification and diagnosis. By incorporating a comprehensive suite of measurements that track the availability, performance, and usage of subscriber-visible services (Web, Email, and News), as well as infrastructure services (such as DNS, NFS, etc.), the monitoring system provides the information essential for problem diagnosis. An easily customizable Web-based user interface enables ISP personnel easily identify the root-cause of problems.

Recently, Seshan et al. have developed a toolkit to allow clients to passively measure and report performance information back to a performance server [10]. There are two key challenges in applying this approach in the ISP environment. First, to enable client-side monitoring, ISPs must support special-purpose client software for all the different platforms (Windows, Mac, Unix, etc.) that subscribers use. Second, the performance reported by clients in an ISP environment is determined by the speed of the dial-up modem link that subscribers use. In order to avoid being biased by client-side bottlenecks, the performance data must be analyzed by taking into account the client modem's capabilities. Lack of information about the modem capabilities of each client and the need to aggregate data from thousands of clients are some of the challenges in the practical applicability of this approach.

7 Conclusions

In this paper, we have presented the design and implementation of a monitoring system that is customized to meet the needs of Internet Service Providers. The practical utility of the monitoring system, validated by experimental evaluation in a real-world ISP network, demonstrates the value of a systematic instrumentation and analysis approach for service management. Our experience demonstrates the significant benefits that can be obtained from making even relatively simplistic measurements on an on-going basis for monitoring ISP networks. Historical QoS and usage baselines derived from on-going monitoring is invaluable in the proactive detection of problems. Furthermore, there is significant value in correlating the results of a set of measurements that cover all of the server, network, and service components involved in providing a service. Such correlation can enable speedy diagnosis of problems, thereby enabling ISPs to offer better QoS to their subscribers.

Our experience as typified by the example problem diagnosis scenario presented in Section 5 also highlights the difficulty in correlating and interpreting the results of a number of heterogeneous measurements manually. To further enhance the usability of our ISP monitoring system, we are beginning to explore correlation and analysis techniques for automating the interpretation of measurement results to identify the root-causes of problems.

References

- [1] R. Wetzel, Customers rate ISP services, *PC Week*, November 1997.
- [2] P. Albitz and C. Liu, *DNS and BIND*, O'Reilly Publications, 1992.
- [3] R. Metcalfe, Net.Medic lets you know when to pull the plug on your Internet connection, *InfoWorld*, April 14, 1997.
- [4] R. Jones, Netperf: A network performance benchmark, Hewlett-Packard Company, February 1996.
- [5] M. Mathis and J. Mahdavi, Diagnosing Internet congestion with a transport layer performance tool, *In proceedings of INET'96*, June 1996, <http://www.psc.edu/~mathis/htmlpapers/inet96.treno.html>
- [6] K. Richardson, Detecting and Understanding Performance Failures with Course-grain Monitoring, *In Proceedings of the Second Grace Hopper Celebration of Women in Computing Conference*, Sept. 1997.
- [7] A. Cockcroft, Advanced monitoring and tuning, *Sun World On-Line*, October 1995, <http://www.sun.com/951001/columns/adrian/column2.html>
- [8] M. Avery, WhatsUp Gold enhances network supervision tasks, *InfoWorld*, Vol. 20, No. 4, Jan 26, 1998.
- [9] R.M. Enger and J.K. Reynolds, FYI on a network management tool catalog: Tools for monitoring and debugging TCP/IP Internets and interconnected devices, *Request for Comments - RFC 1470*, June 1993.
- [10] S. Seshan, M. Stemm, and R.H. Katz, SPAND: Shared Passive Network Performance Discovery, *In proceedings of the first Usenix Symposium on Internet Technologies and Systems*, Monterey, CA, December 1997.