

# Original Investigations

# JAMIA

Research Paper ■

## Evaluating Internet End-to-end Performance: Overview of Test Methodology and Results

---

FRED B. WOOD, DBA, MBA, VICTOR H. CID, MS, ELLIOT R. SIEGEL, PHD

**Abstract** **Objective:** An evaluation of Internet end-to-end performance was conducted for the purpose of better understanding the overall performance of Internet pathways typical of those used to access information in National Library of Medicine (NLM) databases and, by extension, other Internet-based biomedical information resources.

**Design:** The evaluation used a three-level test strategy: 1) user testing to collect empirical data on Internet performance as perceived by users when accessing NLM Web-based databases, 2) technical testing to analyze the Internet paths between the NLM and the user's desktop computer terminal, and 3) technical testing between the NLM and the World Wide Web ("Web") server computer at the user's institution to help characterize the relative performance of Internet pathways.

**Measurements:** Time to download the front pages of NLM Web sites and conduct standardized searches of NLM databases, data transmission capacity between NLM and remote locations (known as the bulk transfer capacity [BTC]), "ping" round-trip time as an indication of the latency of the network pathways, and the network routing of the data transmissions (number and sequencing of hops).

**Results:** Based on 347 user tests spread over 16 locations, the median time per location to download the main NLM home page ranged from 2 to 59 seconds, and 1 to 24 seconds for the other NLM Web sites tested. The median time to conduct standardized searches and get search results ranged from 2 to 14 seconds for PubMed and 4 to 18 seconds for Internet Grateful Med. The overall problem rate was about 1 percent; that is, on the average, users experienced a problem once every 100 test measurements. The user terminal tests at five locations and Web host tests at 13 locations provided profiles of BTC, RTT, and network routing for both dial-up and fixed Internet connections.

**Conclusion:** The evaluation framework provided a profile of typical Internet performance and insights into network performance and time-of-day/day-of-week variability. This profile should serve as a frame of reference to help identify and diagnose connectivity problems and should contribute to the evolving concept of Internet quality of service.

■ *J Am Med Inform Assoc.* 1998;5:528–545.

Over the last two years, use of the Internet by the biomedical community has increased dramatically. The National Library of Medicine (NLM), like many others,<sup>1,2</sup> is making the transition to use of the Internet and the World Wide Web ("Web") for the delivery of services. In June 1997, the NLM announced that access to MEDLINE\* would be free using the World Wide Web. Since that time, the total number of Web "hits" or connections to the MEDLINE search sites has more than doubled and is currently increasing at 10 to 15 percent per month.<sup>3</sup> Usage appears to be split relatively evenly between biomedical researchers, health care professionals, and health care consumers. The NLM is making most of its other databases similarly accessible via Web sites. Indeed, the NLM's intent is that the Internet and the Web be the primary means for information access and delivery, in the hopes of facilitating wider, easier, and less expensive database access and use.

Surveys conducted by the NLM have indicated that the vast majority of traditional NLM customers now have access to the Internet.<sup>4</sup> But just having Internet access does not guarantee satisfactory Internet performance. As Internet use increases domestically and globally, complaints about erratic performance have multiplied, especially during peak business hours. The size, heterogeneity, and complex dynamics of the Internet infrastructure present challenges that are still not well understood even by the Internet technical community.

In 1997, the NLM formally established an Internet connectivity evaluation project (InterCEPt). At the most basic level, InterCEPt is intended to improve our understanding of the role that end-to-end Internet performance plays in facilitating (or hindering) access to the NLM and other biomedical databases. However, Internet performance evaluation is a nascent field, as yet with very limited development in terms of established methodologies or metrics. So an important aspect of InterCEPt is to develop and refine methods and metrics for end-to-end Internet testing and evaluation. Our goal in doing so is to help end users better understand the performance aspects of their current and planned Internet services. For this purpose, we have proposed an Internet quality-of-service concept, which encompasses our efforts to define, describe, and evaluate the quality of Internet services in a way that is meaningful and practical to end users. We hope that our work on Internet quality of service

can help close the gap between the performance levels that Internet service providers and the technical Internet community can deliver and the educated expectations of the Internet user community.

Research on quality-of-service aspects of Internet performance has been designated a priority research area by the Presidential Advisory Committee on High Performance Computing, Communications, Information Technology, and the Next Generation Internet. Also, the National Science Foundation is requiring its grantees and contractors to address Internet performance issues. In a related area, the G7 Global Healthcare Applications Project has approved a subproject activity on international Internet connectivity evaluation, led by the NLM for the United States.

This paper presents and discusses the results of the first phase of the NLM's InterCEPt research. In this research, the methodologic findings are as important as the test results, since the methods and metrics are still evolving. As such, the research presented must still be considered exploratory, even though we believe it constitutes the most complete profile of Internet performance to date for a cross-section of biomedical users and institutions.

## Methods

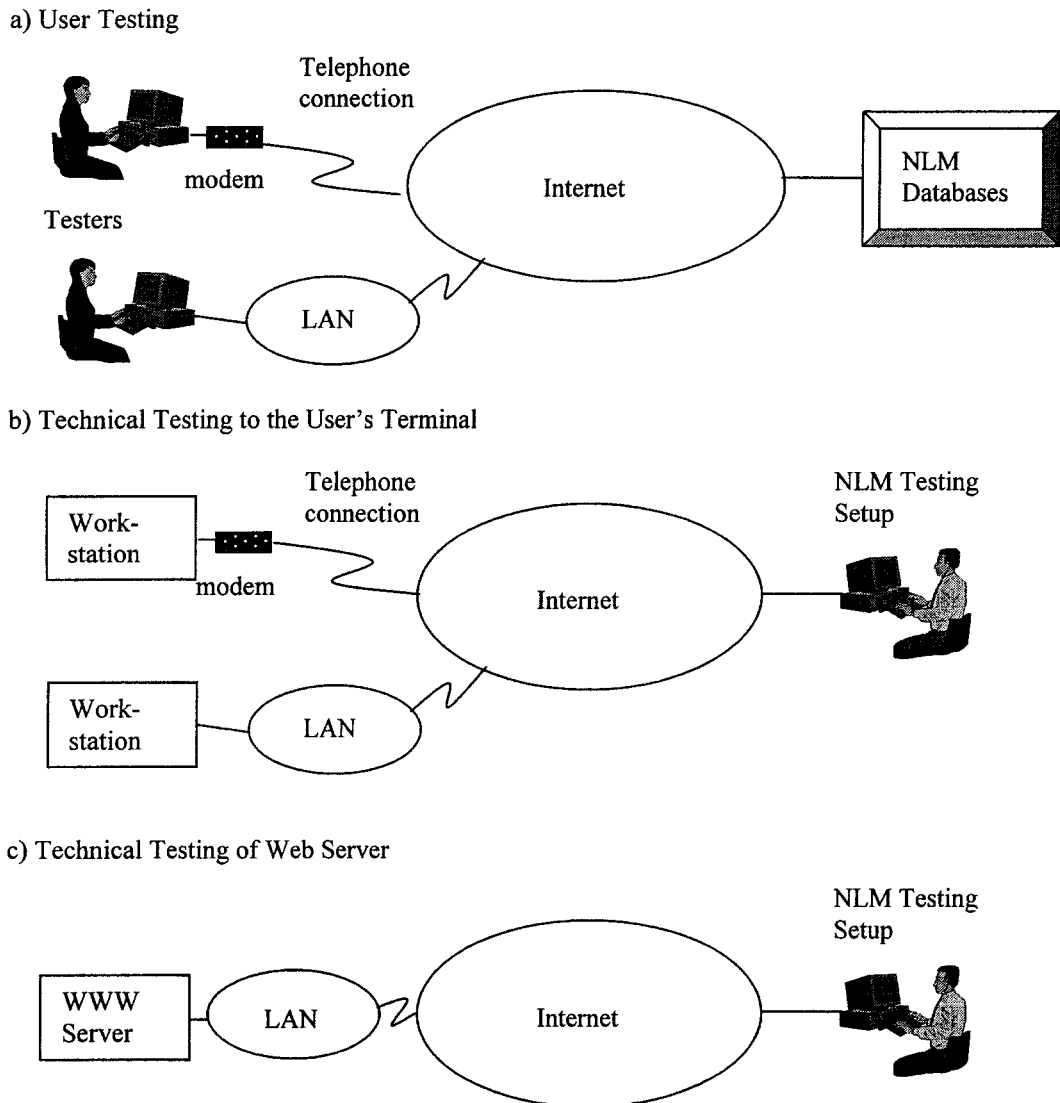
At the outset of the project, the NLM reviewed the field of Internet performance evaluation, including related government, university, and commercial sector activities.<sup>5-13</sup> We identified numerous specific testing tools and methods in use but no consensus yet on a suite of tools and set of metrics that would provide the most meaningful and useful results to our user community. After considerable experimentation and pilot testing, we developed a three-level methodology for Internet connectivity evaluation (Figure 1).

## User Testing

The first level is what we call "user testing." By this we mean tests conducted directly by users, or by their surrogates, with the computers and Internet connections typically used to conduct searches of biomedical databases. For purposes of this project, we focused on Internet-based searching of NLM databases and access to other NLM information, all available from NLM Web servers. At the time of this research, the NLM was connected to the Internet through a T-3 link. The metric used was response times to standardized queries and searches. Response time was defined as

---

\*MEDLINE is a registered trademark and "Internet Grateful Med" and "PubMed" are service marks of the National Library of Medicine.



**Figure 1** End-to-end Internet performance testing strategies.

the time in seconds it takes to completely download the information requested from the appropriate NLM Web site. Each test included measurement of the response times to download the front pages of the Web sites for the NLM, National Center for Biotechnology Information, PubMed (a Web-based interface to searchable databases), and Internet Grateful Med (another Web-based database search interface). The tests also included the measurement of response times for conducting standardized PubMed and Internet Grateful Med searches (the search terms were "discoid lateral meniscus"). The users measured response times with a standard watch having a stopwatch function or with a stopwatch. Test measurements to the nearest second were of sufficient accuracy for purposes of this research.

Testing was conducted at 16 different user locations in 12 states (Table 1). All but three of the test locations were at medical or health sciences, hospital, or health care provider libraries. One was at the office of an information researcher who does significant biomedical information research, and two were at the homes of medical librarians. A total of 347 tests were conducted. Each tester conducted one to several days of testing, usually on an hourly basis, for a half or full day at a time. The total number of individual test measurements was 3,733. Testers were provided a detailed test protocol. For each testing session, the testers recorded specific information on the computer and Internet connection used and filled out a data entry worksheet to record the date and time of day along with the measured response times for each test con-

ducted. Also, testers recorded the details of any problems or error messages that they encountered (e.g., request timed out, no response, server down). Testers were requested to clean their "Web cache" every time before their measurements, to ensure that the Web pages were completely downloaded every time.

### Technical Testing to User Terminal

The second level we call "technical testing to the user terminal." Here we refer to technical tests of the Internet path between the NLM Internet test center computer and the user's desktop computer. This is computer-to-computer testing with no direct involvement of the user, except in setup. The test metrics used were bulk transfer capacity (BTC), to measure the data transmission capacity of the Internet pathway between two locations; "ping" round-trip time as an indicator of the latency or propagation delay in the net-

work for data traveling from one end of the network to the other and back; and the number and sequencing of links or hops from origin to destination (network routing). We also measured packet loss, defined as the percentage of data packets for which the testing software did not receive an acknowledgment of successful transmission. Packet loss is especially useful in describing the performance of congested networks, and we have included data on this metric in the analysis of some specific performance problems. The test tools used were TReno, Ping, and Traceroute, all of which have been successfully tested in the network engineering community.<sup>14</sup>

This level of testing requires user cooperation in determining the Internet protocol address of the specific computer terminal being used in the test and in ensuring that the Internet connection is open for the duration of the technical testing. Also, the user needs to

Table 1 ■

User Testing, Site Details

Location	Institution	1998 Testing Dates	Number of Tests	Number of Test Measurements*
Los Angeles, Calif.	Louise M. Darling Library, UCLA	Apr 3,6	16	176
Oakland, Calif.	Kaiser Permanente Medical Center	Mar 30, Apr 4	18	197
Middletown, Conn.	Tremeine Medical Library, Middlesex Hospital	Apr 2	9	99
Peoria, Ill.	Health Sciences Library, University of Illinois Medical College	Mar 30, 31; Apr 1	27	292
Baton Rouge, La.	Information Research	Apr 8, 9	15	163
New Orleans, La.	Tulane University Medical Center Library	Apr 6, 7	18	198
Salem, Mass.	Health Science Library, North Shore Medical Center	Apr 2, 3	17	186
Omaha, Neb.	McGoogan Library, University of Nebraska Medical Ctr.	Apr 3, 6	16	174
Buffalo, N.Y.	Health Science Library, Millard Fillmore Hospital	Apr 22-24	14	154
New York, N.Y.	Samuel L. Wood Library, Cornell University Medical College	Apr 1-3, 6	17	184
New York, N.Y.	Medical librarian's home computer	Apr 1-3, 6	25	267
New York, N.Y.	Medical librarian's home computer, text-only†	Apr 1-4	25	247
Oklahoma City, Okla.	Integrus Baptist Medical Center Library	Apr 6	9	99
Portland, Ore.	Western States Chiropractic Hospital	Apr 3, 6-8	9	98
Gladwyne, Pa. (Philadelphia suburb)	Medical librarian's home computer	Mar 5-15	68	748
Greenwood, S.C.	Medical Library, Self Memorial Hospital	Apr 1, 2	18	190
Greenwood, S.C.	Community Health Information Center, Self Memorial Hospital	Apr 2	9	96
Seattle, Wash.	Health Sciences Library, University of Washington	Mar 31, Apr 1, 2	<u>17</u>	<u>165</u>
Totals			347	3,733

\*Every test included 11 different measurements. This column lists the total number of successful measurements over the complete testing period.

†The user downloaded only the text content of Web pages. Normally, the entire contents of the Web pages (text and graphics) were downloaded in the experiments.

Table 2 ■

## Technical Testing, Site Details

Location	Institution	Hostname/Address of Computer Tested	1998 Test Date(s)
User terminal testing			
Baton Rouge, La.	Information Research	202-150-246.ipt.aol.com*	Apr 7
New Orleans, La.	Tulane University Medical Center Library	rsl.tcs.tulane.edu	Apr 7
New York, N.Y.	Samuel L. Wood Library, Cornell University Medical College	mac103789.med.cornell.edu	Apr 2
Greenwood, S.C.	Self Memorial Hospital	User28.ais-gwd.com*	Apr 2
Gladwyne, Pa.	Medical librarian's home computer	209.107.11.162*	Mar 24
Web-site testing: Domestic			
Los Angeles, Calif.	Louise M. Darling Library UCLA	www.library.ucla.edu	Mar 27–Apr 10†
Denver, Colo.	St. Joseph's Hospital	saintjosephdenver.org	Mar 27–Apr 10
Middletown, Conn.	Middlesex Hospital	www.midhosp.org	Mar 27–Apr 10†
Macon, Ga.	Mercer University	gain.mercer.edu	Mar 27–Apr 10†
Chicago, Ill.	Alzheimer's Association	www.alz.org	Mar 27–Apr 10
Peoria, Ill.	University of Illinois Medical College	www.uicomp.uic.edu	Mar 27–Apr 10
New Orleans, La.	Tulane University Medical College	www.mcl.tulane.edu	Mar 27–Apr 10
Omaha, Neb.	University of Nebraska Medical College	www.unmc.edu	Mar 27–Apr 10
New York, N.Y.	Cornell University Medical College	www.med.cornell.edu	Mar 27–Apr 10
Buffalo, N.Y.	Millard Fillmore Health System	www.mfhs.edu	Mar 27–Apr 10
Oklahoma City, Okla.	Integrus Baptist Hospital	www.integrus-health.com	Mar 27–Apr 10
Burien, Wash.	Highline Community Hospital	www.halcyon.com	Mar 27–Apr 10
Seattle, Wash.	Health Sciences Library, University of Washington	www.hslib.washington.edu	Mar 27–Apr 10
Web-site testing: Domestic versus international comparisons			
Ottawa, Canada	Canadian Institute for Scientific and Technical Information (CISTI)	www.cisti.nrc.ca	Feb 6–24
Paris, France	Institut National de la Sante et de la Recherche Medicale (INSERM)	www.inserm.fr	Feb 6–24
Köln, Germany	Deutsches Institut für Medizinische Dokumentation und Information (DIMDI)	www.dimdi.de	Feb 6–24
Rome, Italy	Istituto Superiore di Sanita (ISS)	www.iss.it	Feb 6–24
Tokyo, Japan	Japan Science and Technology Corporation (JST)	www.jst.go.jp	Feb 6–24
West Yorkshire, UK	The British Library (BL)	portico.bl.uk	Feb 6–24
Farmington, Conn., USA	University of Connecticut Health Center	www3.uchc.edu	Feb 6–24
Seattle, Wash., USA	University of Washington Health Sciences Library	www.hslib.washington.edu	Feb 6–24

\*Computer hostname/address was dynamically assigned on a per-connection basis.

†Intermittent testing.

ensure that the terminal is not used for other purposes while the testing is underway. However, once the test is properly configured, users need not be actively involved. The actual testing was done remotely from the NLM. For this phase of the research, the NLM was able to successfully test the pathways to user terminals at five locations—three with dial-up Internet connections and two with fixed Internet connections (Table 2). Each of these sites was tested typically at 10-minute intervals for one to several hours on the test day. Technical testing was not possible at several locations because of restrictions imposed by local “firewalls” and other security measures. A firewall is legitimately intended to keep out hackers and other undesired or malicious users but also can limit attempts to test end-to-end Internet performance. Un-

fortunately, the same tools used to evaluate Internet performance can also be used to compromise the security of remote systems.

### Technical Testing to Web Sites

The third level tests the technical performance of Internet paths between the NLM and computers hosting Web services at user locations. Web hosts typically use the same Internet link and internal network infrastructure as the end users at that institution. So long as the Web host is in the same local area network as the user and both share the same Internet link, the Web hosts can serve as a reasonable surrogate in lieu of testing direct to the user's terminal. Also, as with the user terminal testing, firewalls and other security

measures can and did prevent or impair testing in some situations.

The metrics and tools used to test the performance of Internet communications to Web hosts were the same as those used to test user terminals. The metrics were BTC, round-trip time, network route distance/stability, and packet loss. The tools were TReno, Ping, and Traceroute. Eight Web hosts were fully tested, including five of the host locations where end-user testing was conducted and two of the host locations where user terminal testing was carried out (Table 2). Another five Web hosts (for a total of 13) were only partially tested, because of problems with firewalls and other security measures as noted; end-user testing was conducted at all five of these hosts. Web host testing was conducted typically for two to three weeks, with data collected automatically at about ten-minute intervals.

Eight other Web hosts (plus one of the previous 13) were tested for a longer period of time (two to three months) in connection with the G7 Global Healthcare Applications subproject mentioned earlier. These hosts were selected from the NLM's international affiliates (known as International MEDLARS [Medical Literature and Retrieval System] Centers) in the G7 countries and from the NLM's regional medical libraries in the United States.

## Results

### User Testing Results

The user testing results (Tables 3 and 4) showed wide variability in measured response times. Test results are presented in terms of median values, which we found best represent typical response times. No responses (meaning that the Internet connection could not be made or the search request received no response or timed out) were not included in these tables, but were treated separately. The median time per test site to download the NLM front page ranged from a low of 2 to a high of 59 sec (Table 3). The response times and their variability for the front pages of Internet Grateful Med, National Center for Biotechnology Information, and PubMed were much faster, ranging from 1 to 14 sec, 1.5 to 24 sec, and 1 to 12 sec, respectively (Table 3). The special testing from a New York City dial-up site with and without graphics—that is, downloading text and graphics versus downloading only the text of the Web page—showed a substantial reduction in NLM front page response times with text-only downloading (4.6 sec with text only, compared with 23.4 sec with graphics).

The response times for conducting standardized searches ranged from 2 to 14 sec for PubMed, and from 4 to 18 sec for Internet Grateful Med to get the search results (Table 4). The IGM search results took somewhat to significantly longer to obtain, with a response time typically 10 to 30 percent longer. Since the Internet connections, computer platforms, and search terms were held constant and search times were similar (during the business day), the differences appear to be attributable to the search engines and Web site architectures of NLM's Web servers. Also, during the test period, Internet Grateful Med was being upgraded to address known capacity problems. The differences in response times for downloading the first abstract are less significant, with PubMed and Internet Grateful Med responding in a similar fashion (Table 4).

As a group, those sites with high-bandwidth Internet connections (10+ Mbps up to T-3, 45 Mbps), and fast computers (Pentium 200 Mhz) had the shortest median response times—2 to 4 sec to download front pages and 2 to 6 sec to obtain search results. The sites with dial-up Internet connections (14.4 or 28.8 Kbps) generally had the longest response times, regardless of the type of computer platform. Sites with T-1 (1.5 Mbps) bandwidth had mixed results. Some T-1 sites with faster computers had response times almost as quick as the T-3 sites tested, whereas other T-1 sites with slower computers had much longer response times. All the links, except the dial-up links tested, were subject to the load of several users using Internet services concurrently over the same link.

The user tests also provided data on the number and type of problem connections. The overall problem rate was about 1 percent; that is, on the average users experienced a problem incident once every 100 test measurements. The types of problems fell into two broad groupings. Problems that seem related to Internet Grateful Med or to its Web site included the message "error caused by break in connection to MEDLARS system; please retry or edit your search," which accounted for about one fourth of total incidents, and other error messages (e.g., "internal server error," "IGM system has timed out; you must restart"). Problems that could not be clearly attributed to NLM's back-end software and that could be related to Web server software or the general Internet included very slow response time (defined as 60+ seconds), which accounted for about one fourth of all problem incidents (spread across the NLM, National Center for Biotechnology Information, PubMed, and Internet Grateful Med Web sites); no response (the user restarted or quit before receiving any response), which

Table 3 ■

## User Testing Results: Front-page Download Response Times\*

Location/Institution	Internet Link†	Computer Platform‡	Median Response Time (sec)§			
			NLM	NCBI	PubMed	IGM
Buffalo, N.Y. Millard Fillmore Hospital	T-3	P5/233/64	2	1.5	1	1
Los Angeles, Calif. UCLA	T-3	P5/200/64	2	4	1	2
New York, N.Y. Cornell University	T-3	PM/200/64	2	1.5	1.5	1.2
Seattle, Wash. University of Washington	10 Mbps	P5/200/32	2	2	2	1
New Orleans, La. Tulane University	10 Mbps	P5/200/32	2	1.4	1.2	0.9
Salem, Mass. North Shore Medical Center	T-1	P5/133/32	6	6	3	3
Oakland, Calif. Kaiser Permanente	T-1	P5/133/32	8	15	11	7.5
Middletown, Conn. Middlesex Hospital	T-1	PM/80/26	21	8	7	7
Oklahoma City, Okla. Integrus Baptist Hospital	T-1	I4/66/32	12	24	12	14
Peoria, Ill. University of Illinois	T-1	I4/66/16	18	7	7	7
Omaha, Neb. University of Nebraska	4 Mbps	I4/66/16	8	7.5	5.5	4
Portland, Ore. Western States Chiropractic Hospital	56 K FR	P5/200/32	2	2	2	1
New York, N.Y. Medical librarian's home computer	28.8 Kbps dial-up	Mc/G3/96	23.4	7	6.5	5.1
New York, N.Y. Medical librarian's home computer, text only	28.8 Kbps dial-up	Mc/G3/96	4.6	2.6	2.5	2.6
Greenwood, S.C. Self Memorial Hospital, Medical Library	28.8 Kbps dial-up	P5/133/72	27.5	9	8	6
Gladwyne, Pa. Medical librarian's home computer	28.8 Kbps dial-up	P5/90/24	26	8	9	5
Baton Rouge, La. Information Research	28.8 Kbps dial-up	I4/33/20	12	11	8	7
Greenwood, S.C. Self Memorial Hospital, Community Health Information Center	14.4 Kbps dial-up	P5/200/64	59	12	11	9

\*Response times are given for all sites, which are grouped by type of Internet link and computer platform.

†Advertised Internet link type or link speed is expressed in kilobits per second (Kbps) or megabits per second (Mbps), as indicated. Actual connection speed may vary. FR indicates frame relay link.

‡Platform description consists of processor type/processor speed (in megahertz) or model/memory size (in megabytes). P5 indicates Pentium; PM, Power Macintosh; I4, Intel 486; Mc, other Macintosh.

§Response times indicate how long it took to download the front page of the specified site. NLM indicates the National Library of Medicine; NCBI, National Center for Biotechnology Information; IGM, Internet Grateful Med.

||User downloaded only the text content of Web pages (i.e., automatic downloading of graphics in Web browser was disabled).

Table 4 ■

## User Testing Results: Standardized Search Response Times\*

Location/Institution	Internet Link†	Computer Platform‡	Median Response Time (sec)			
			PubMed		IGM	
			Full Results	1st Abstract§	Full Results	1st Abstract§
Buffalo, N.Y. Millard Fillmore Hospital	T-3	P5/233/64	3	1	4.5	3
Los Angeles, Calif. UCLA	T-3	P5/200/64	2	2	4	3
New York, N.Y. Cornell University	T-3	PM/200/64	2.9	1.4	4.7	2.5
Seattle, Wash. University of Washington	10 Mbps	P5/200/32	3	2	5.5	4
New Orleans, La. Tulane University	10 Mbps	P5/200/32	2.4	1.7	4.8	2.5
Salem, Mass. North Shore Medical Center	T-1	P5/133/32	4	4	7	4
Oakland, Calif. Kaiser Permanente	T-1	P5/133/32	8	7	16.5	10.5
Middletown, Conn. Middlesex Hospital	T-1	PM/80/26	11	6	15	7
Oklahoma City, Okla. Integrus Baptist Hospital	T-1	I4/66/32	6	11	18	13
Peoria, Ill. University of Illinois	T-1	I4/66/16	8	3	10	4
Omaha, Neb. University of Nebraska	4 Mbps	I4/66/16	5	5.5	5.5	4
Portland, Ore. Western States Chiropractic Hospital	56 K FR	P5/200/32	4	3	8	4
New York, N.Y. Medical librarian's home computer	28.8 Kbps dial-up	Mc/G3/96	6.3	3.2	12.1	4.3
New York, N.Y. Medical librarian's home computer, text only	28.8 Kbps dial-up	Mc/G3/96	5.3	2.4	6.5	3.6
Greenwood, S.C. Self Memorial Hospital, Medical Library	28.8 Kbps dial-up	P5/133/72	7	4	10	5
Gladwyne, Pa. Medical librarian's home computer	28.8 Kbps dial-up	P5/90/24	6	7.5	10.5	4
Baton Rouge, La. Information Research	28.8 Kbps dial-up	I4/33/20	13	7	17	10
Greenwood, S.C. Self Memorial Hospital, Community Health Information Center	14.4 Kbps dial-up	P5/200/64	14	10	15	7.5

\*The test search was for "discoid lateral meniscus" in PubMed and Internet Grateful Med (IGM). Response times are given for all sites, which are grouped by type of Internet link and computer platform.

†Advertised Internet link type or link speed is expressed in kilobits per second (Kbps) or megabits per second (Mbps), as indicated. Actual connection speeds may vary. FR indicates frame relay link.

‡Platform description consists of processor type/processor speed (in megahertz) or model/memory size (in megabytes). P5 indicates Pentium; PM, Power Macintosh; I4, Intel 486; Mc, other Macintosh.

§Time indicates how long it took to display the first abstract in the full result set.

||User downloaded only the text content of Web pages (i.e., automatic downloading of graphics in Web browser was disabled).



Table 5 ■

## Technical Testing Results

Location/Institution	Internet Link	Median BTC*	Mean Ping RTT (msec)†	Number of Hops‡	Stability of Route§
User terminal test results:					
Baton Rouge, La. Information Research	28.8 Kbps dial-up	40 Kbps	424	9	74%
New Orleans, La. Tulane University Medical Center	10 Mbps	321 Kbps	76	14	100%
New York, N.Y. Cornell University Medical College	T-3	1.99 Mbps	18	11	73%
Greenwood, S.C. Self Memorial Hospital, Medical Library	28.8 Kbps dial-up	126 Kbps	353	12	100%
Gladwyne, Pa. Medical librarian's home computer	28.8 Kbps dial-up	60 Kbps	160	13	NA
Domestic Web-site test results:					
Los Angeles, Calif. Darling Library, UCLA	T-3	2.67 Mbps	70	13	80%
Denver, Colo. St. Joseph's Hospital	T-1	954 Kbps	69	13	100%
Middletown, Conn. Middlesex Hospital	T-1	1.04 Mbps	52	12	100%
Macon, Ga. Mercer University	10 Mbps	NA	96	13	100%
Chicago, Ill. Alzheimer's Association	T-1	1.1 Mbps	33	13	69%
Peoria, Ill. University of Illinois Medical College	T-1	161 Kbps	65	13	76%
New Orleans, La. Tulane University Medical College	10 Mbps	1.66 Mbps	76	14	100%
Omaha, Neb. University of Nebraska Medical College	4 Mbps	NA	NA	12+	~100%
New York, N.Y. Cornell University Medical College	T-3	4.91 Mbps	18	16	25%
Buffalo, N.Y. Millard Fillmore Health System	T-1	1.15 Mbps	55	12+	~100%
Oklahoma City, Okla. Integris Baptist Hospital	T-1	1.05 Mbps	83	12	100%
Burien, Wash. Highline Community Hospital	NA	2.84 Mbps	126	12	55%
Seattle, Wash. University of Washington Health Sciences Library	10 Mbps	3.74 Mbps	83	12	56%

\*Bulk transfer capacity (BTC) was measured in kilobits per second (Kbps) or megabits per second (Mbps), as indicated.

†Round-trip time (RTT), as reported by Ping.

‡Number of hops, or steps, in the most frequently used Internet pathway (route) between the NLM and the remote location, as reported by Traceroute. A "plus" sign (+) indicates that the actual value is greater than or equal to the value shown.

§Percentage of times an observed route (as returned by Traceroute) from the NLM to the indicated location was the same as the immediately preceding observed route to the same location. An "approximately" sign (~) indicates that the value shown is estimated.

||NA indicates data that were not available or could not be measured (e.g., because of a firewall).

accounted for another one fourth of the incidents (mostly involving Internet Grateful Med search steps); and various error messages (e.g., "network error occurred; reset by peer," "unexpected error has caused break in connection," "the URL is not reachable; the server could be down").

### User Terminal Testing Results

While the user tests measure the response times actually experienced by the users, the user terminal test results (Table 5) provide insight into some of the technical reasons for the performance of Internet connections. The TReno test results showed the large differences in BTC between dial-up and fixed Internet connections. Median dial-up capacity was in the range of 40 to 126 Kbps, as would be expected for 28.8-Kbps modem connections with data compression. This does highlight that modem connections can considerably exceed the rated modem speed through compression of the data being transferred. Notice that BTC values are presented using median rather than mean values in order to more accurately characterize the typical values and minimize the distortion caused by outliers. Median BTC for the T-1 connection was about 320 Kbps and for the T-3 connection about 2 Mbps. These values are considerably below the theoretic bandwidth of the Internet link.

The Ping test results showed the corresponding differences in the latency of the Internet pathways tested (Table 5). The average round-trip times for the dial-up connections were in the range of 160 to 400 msec, compared with about 18 msec for the T-3 connection successfully tested.

The network routing test results illustrate the relatively large number of hops, or steps, in typical Internet paths (9 to 14 for this sample) and the high stability of some Internet routes. The dial-up connection to Greenwood, South Carolina, and the T-1 connection to New Orleans, Louisiana, both had about the same number of hops (12 and 14, respectively) and 100 percent route stability. This means that for all the measurements taken, the Internet route between the NLM and the test sites was exactly the same. For the other two sites tested—to New York City and to Baton Rouge, Louisiana—the route stability was about 73 percent, i.e., changes in the Internet route were observed in about one quarter of the measurements.

For the connection between the NLM and the user terminal at the Cornell University Medical College in New York, the packets of data made the round trip to the New York region (seventh hop) in less than 5 msec (Figure 2). But the median round trip to the user ter-

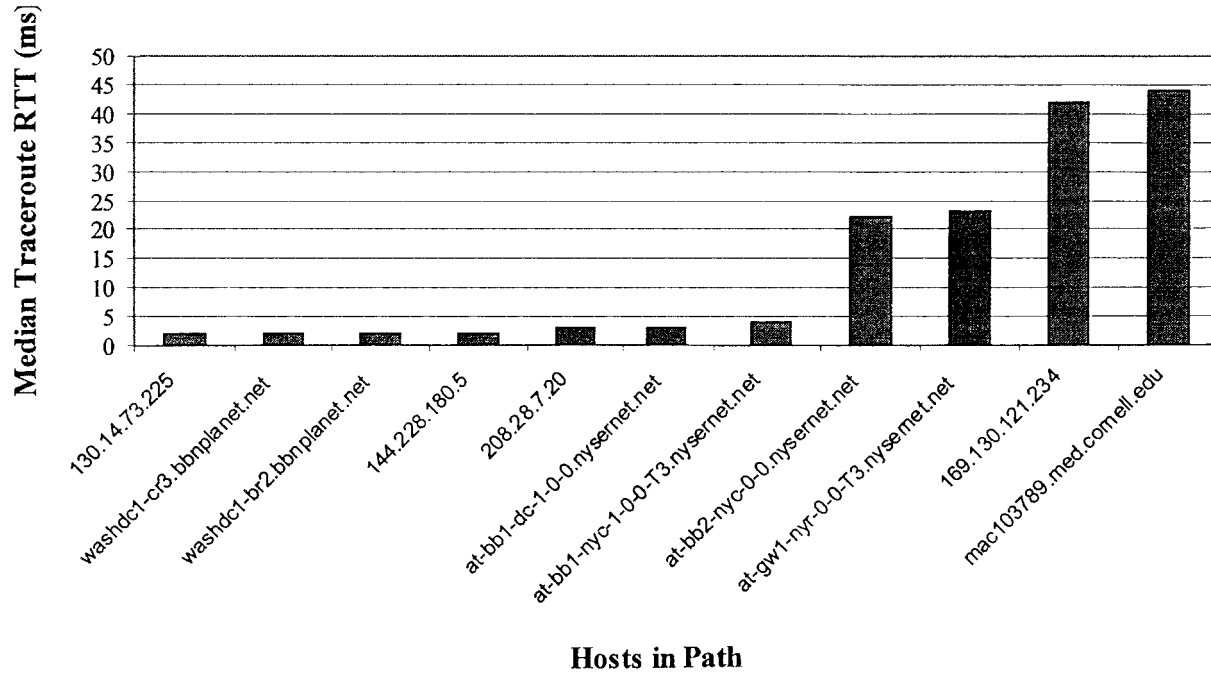
terminal was about 44 msec, indicating that about 90 percent of the total time occurred during the last few hops within the local Internet provider and campus networks and the end-user terminal. Likewise, for the pathway from NLM to the computer terminal of a medical librarian in Gladwyne, Pennsylvania (a suburb west of Philadelphia), the median round-trip time was 8 msec to the Philadelphia area (ninth hop) and 20 msec to the local Internet provider (12th hop), but 160 msec to the end-user terminal (Figure 3). Thus, in this example the so-called last mile over local phone lines accounted for about 88 percent of the total round-trip time.

### Web Site Testing Results: Domestic

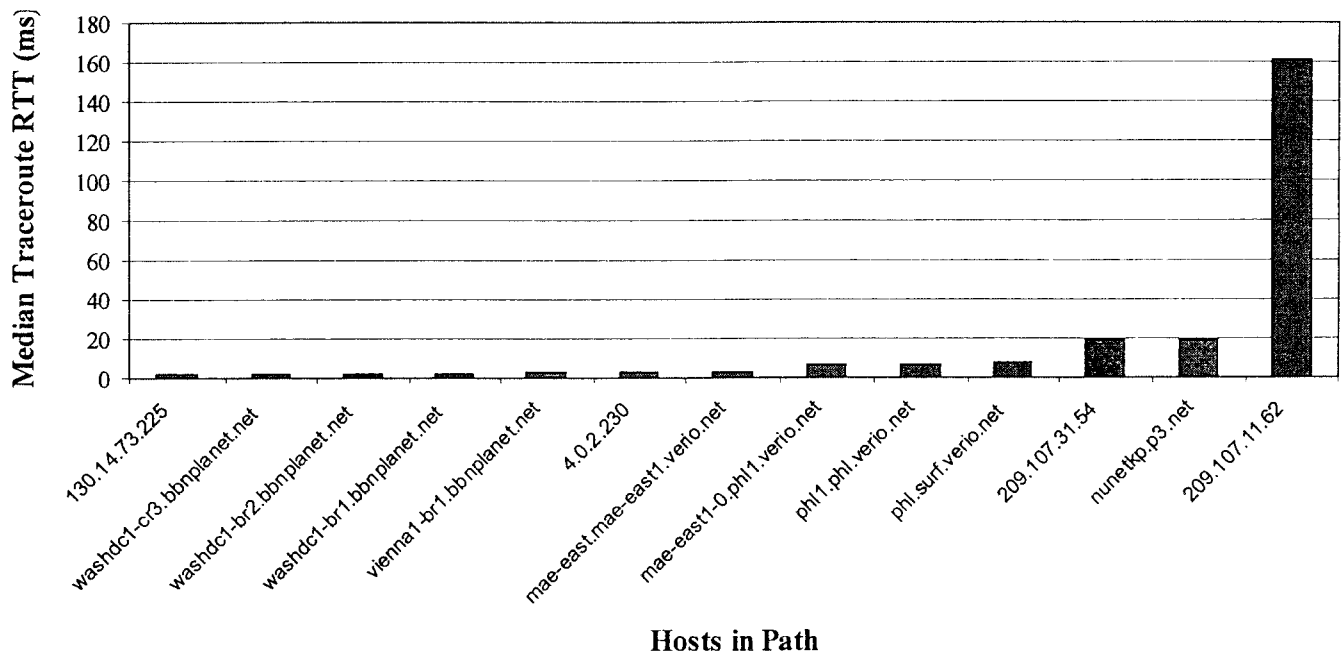
Testing of Internet pathways to Web hosts is technically straightforward and less labor intensive than either user testing or user terminal testing. Web host testing is entirely automated, with minimal human intervention necessary. The eight sites fully tested and five sites partially tested helped further complete the profile of Internet connectivity (Table 5). Regardless of the rated Internet connection or the geographic location, the number of hops in the pathway or route ranged from 12 to 16. The stability of the pathway varied: Several were 100 percent stable for the period of testing, whereas others ranged from 25 to 80 percent stable. There is no obvious correlation between the number and stability of the hops, the speed of transmission, and the round-trip latency of the pathway. For example, the pathway between the NLM and the Cornell Medical College Web site in New York City has the most hops (16) and lowest stability (25 percent) but the fastest average round-trip time (20 msec). Comparing sites with T-1 connections, the number of hops ranged from 12 to 13, route stability from 69 to 100 percent, and round-trip times from 33 to 83 msec.

The pathway to the University of Illinois College of Medicine at Peoria had a considerably smaller median BTC than other sites also using T-1 links—about 160 Kbps compared with a range of 900 Kbps to 1.1 Mbps for the other T-1 sites. Packet loss for the NLM–Peoria pathway was particularly high during peak load periods, exceeding 30 percent, compared with between 0 and 5 percent in off-peak periods, which could be an indication of a highly congested pathway and could explain, at least in part, the lower overall performance.


The test results clearly show the wide variability in BTC as a function of time of day and day of week. This is dramatically illustrated by the BTC measurements for the pathway to the Integris Hospital Web




**Figure 2** Median Traceroute round-trip times (RTT, in milliseconds;  $n = 120$ ) from NLM to hosts in path to mac10389.med.cornell.edu (Cornell University Medical College, New York City; T-3 link). Tests were performed on Apr 2, 1998, from 10:36 A.M. to 4:50 P.M..




**Figure 3** Median Traceroute round-trip times (RTT, in milliseconds;  $n = 63$ ) from NLM to hosts in path to 209.107.11.62 (Gladwyne, Pennsylvania; dial-up link). Tests were performed on Mar 24, 1998, from 2:10 P.M. to 5:30 P.M..



**Figure 4** Bulk transfer capacity (BTC, in kilobits per second) versus time from NLM to www.integris-health.com (Integris Baptist Medical Center, Oklahoma City, Oklahoma; T-1 link). Tests were performed from midnight Mar 27 to 1:00 P.M. Apr 7, 1998. Packet size was 8,166 bytes.



**Figure 5** Bulk transfer capacity (BTC, in kilobits per second) versus time for data transfer from NLM to www.med.cornell.edu (Cornell University Medical College, New York City; T-3 link). Tests were performed from midnight Mar 27 to 1:00 P.M. Apr 7, 1998. Packet size was 8,166 bytes.



**Figure 6** Bulk transfer capacity (BTC, in kilobits per second) of selected Web servers, by country. Tests were performed Feb 1998. Values shown, from top to bottom for each location, are maximum, 75th percentile, median, 25th percentile, and minimum. Packet size was 8,166 bytes. NOTE: UWASH indicates University of Washington, Seattle, Washington; UCHC, University of Connecticut Health Center, Farmington, Connecticut; NRC, National Research Council, Ottawa, Canada; INSERM, Institut National de la Sante et de la Recherche Medicale, Paris, France; DIMDI, Deutsches Institut für Medizinische Dokumentation und Information, Köln, Germany; ISS Istituto Superiore di Sanita, Rome, Italy; JST, Japan Science and Technology Corporation, Tokyo, Japan; BL, The British Library, West Yorkshire, United Kingdom.

host in Oklahoma City (Figure 4). Here, during the 12-day test period, the off-peak nighttime BTC was about 1.35 Mbps, close to the maximum capacity for a T-1 link (1.544 Mbps). During weekdays, the BTC shrank to about 250 Kbps and sometimes to less than 100 Kbps. On weekends, the BTC was usually higher than on weekdays. Packet loss varied from 0 percent off-peak to 18 percent during peak load periods, with one instance of 100 percent packet loss recorded during the testing period. The pathway to the Cornell University Medical College Web host shows a similar BTC pattern (Figure 5). The off-peak nighttime BTC was about 5.4 Mbps, often dropping to less than 3 Mbps during weekdays, compared with BTC values that are seldom less than 4 Mbps on weekend days. Packet loss varied between 0 percent during nighttime hours up to around 3 percent during weekday business hours, and on occasion it exceeded 5 percent during a period of higher congestion between 74 and 96 hours into the test (Figure 5).

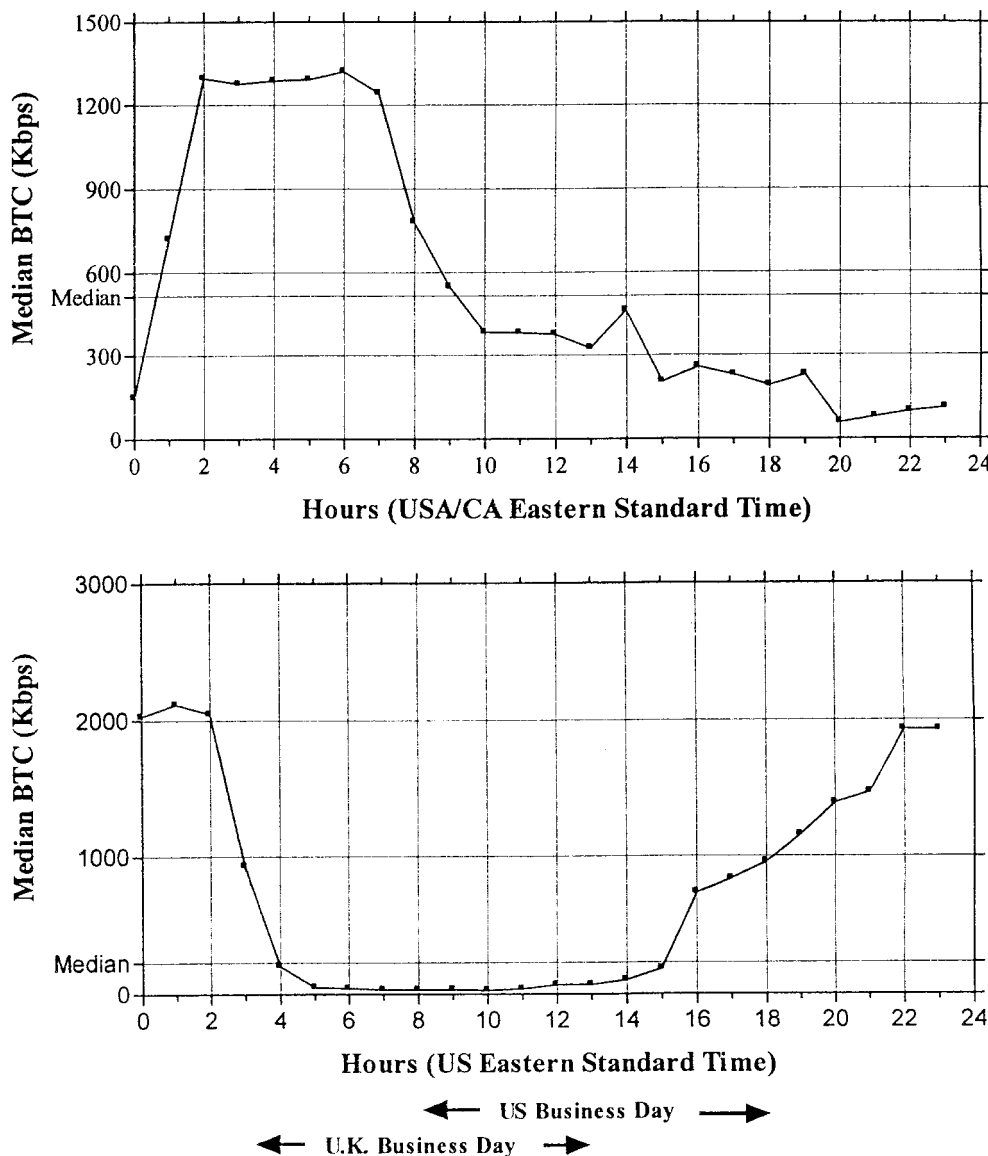
#### Web Site Testing Results: Domestic versus International

For comparison purposes, the NLM tested pathways to selected international Web hosts in the G7 countries

(Table 2). The test results showed a similar range of values. For median BTC (Figure 6), the two U.S. pathways tested at 3.4 and 1.3 Mbps, followed by the United Kingdom and Canadian sites at 640 Kbps and 537 Kbps, respectively. The German, Italian, French, and Japanese hosts tested at about 126, 44, 37, and 38 Kbps, respectively. However, all hosts tested exhibited at least one order of magnitude variation between maximum and minimum BTC measured during the test period. One U.S. host (University of Washington Health Sciences Library) and the Canadian host (National Research Council) evidenced two and one-half orders of magnitude variation, and the British host (British Library) three orders of magnitude variation.

The Internet pathway between the NLM and the British Library (London metropolitan area) dramatizes the variability. The off-peak nighttime BTC was typically around 2.1 Mbps, although it was sometimes as large as 3.4 Mbps. In comparison, during peak daytime weekday hours, the BTC dropped to less than 100 Kbps. The median weekday BTC for each of the eight sites was averaged by hour over the test period to highlight the time-of-day variations. The results shown for the Canadian and British sites were typical (Figure 7a). For the Internet path to the Canadian Na-


**Figure 7** Bulk transfer capacity (BTC, in kilobits per second) for paths from the NLM to selected Web servers in, *top*, Canada (www.cisti.nrc.ca at the National Research Council, Ottawa) and, *bottom*, the United Kingdom (portico.bl.uk at The British Library, West Yorkshire), based on hourly median BTC values over an 18-day test period in Feb 1998.



tional Research Council, the median BTC averaged 1.3 Mbps between 2:00 and 7:00 A.M. EST and then dropped precipitously to about 0.75 Mbps by 8:00 A.M. The BTC shrank further to about 0.38 Mbps by 10:00 A.M., 0.2 Mbps by 3:00 P.M., and 0.06 Mbps (60 Kbps) at 8:00 P.M. For the Internet path to the British Library, the median BTC averaged about 2 Mbps from 4:00 to 8:00 A.M. GST (London local time) and then dropped to about 0.9 Mbps by 9:00 A.M., 0.2 Mbps by 10:00 A.M., and bottomed out at about 0.05 Mbps (50 Kbps) by 11:00 A.M. until 5:00 P.M., when the median BTC starts increasing again. By overlaying the U.S. and U.K. business days (Figure 7b), it is apparent that almost all the measured BTC reduction is due to peak hour Internet congestion in the United Kingdom (probably

in the London metropolitan area). The BTC for the NLM-British Library path reached a minimum three hours before the U.S. business day started.

In order to explore whether this pattern is more universal, the weekday median BTC measurements were averaged by hour for all eight test locations. For each Web site tested, the hours (local time) during which the median BTC hourly average fell above or below the overall host median were determined (Figure 8). The results show clearly that all hosts evidenced a similar pattern, with a reduction in median BTC occurring sometime between 8:00 A.M. and 10:00 A.M. local time and not recovering until sometime between 4:00 P.M. and midnight local time.



**Figure 8** Profile of Internet path loading for selected Web servers, by country, as a function of host local time, based on daily median bulk transfer capacity (BTC) values over an 18-day test period in Feb 1998. Gray areas indicate BTC above median; black areas indicate BTC below median. See Figure 6 legend for an explanation of abbreviations.

## Discussion

### Preliminary Methodology Framework

The first contribution of this research is the development of a preliminary framework for testing end-to-end Internet performance. This research has demonstrated that a multilevel approach, using multiple tools, methods, and metrics, is probably necessary. Also, the set of tools and methodologies needs to be perfected to allow better understanding of all aspects of Internet performance that affect end-user service and to make possible improved problem diagnosis. The tools tested have some limitations. For example, TReno did not record valid BTC values when the packet loss rate was high, and Traceroute could not estimate a network path very precisely when routes changed often. However, the overall results demonstrated the feasibility of using the available tools to evaluate the performance of end-to-end Internet communications, including pathways to workstations with dial-up as well as fixed Internet connections.

We also are testing and developing methods for emulating end-user testing. Some private companies have emulated parts of the user testing, for example by automated testing of selected Web sites and Internet providers via dial-up (e.g., Inverse Network Technology, Inc.) or fixed (e.g., Keynote Systems, Inc.) Internet connections.

Symmetric testing is an important goal for future methodologic development. For the NLM's information services (and those of other biomedical libraries and information providers), technical testing from NLM to the user is probably most important, since the bulk of the information transfer flows from NLM to the user. Users typically use only a few words or phrases to make a query or initiate a database search,

but receive in response a much larger volume of information. Ideally, however, technical tests should be two-way, in order to explore asymmetries in the Internet pathways and their relative performance between any two points. The NLM is currently collaborating with domestic and foreign partners to carry out various forms of symmetric testing and to develop tools for this purpose.

### Exploratory Internet Performance Profile

The second contribution is an exploratory profile of typical Internet performance, based on the results of the testing reported here (see Tables 3, 4, and 5). In January 1998 the total number of hosts in the Internet was estimated to be more than 29 million and growing quickly, and the network technology and traffic patterns change very dynamically.<sup>15,16</sup> It is very difficult to test a truly representative sample of Internet hosts and obtain results that remain meaningful over time. While the profile obtained in this initial study needs to be validated by additional research, it provides a benchmark against which individual users can compare their own test results. The performance profile can be used in work with system developers and network providers (including Internet service providers) to identify, diagnose, and solve Internet performance problems.

This initial profile shows, for example, that often performance of Internet communications through a T-1 link can be more susceptible to performance degradation than even much slower dial-up connections, probably because of the sharing of the T-1 link among many concurrent users (whereas dial-up links typically have a single user). The profile also indicates that the available throughput capacity via a network path can be consistently much smaller than the maximum theoretic size of the Internet links connecting the users to the Internet. For example, the University of Illinois

Medical College at Peoria has a T-1 link but an effective median BTC of less than 11 percent the size of this link (Table 5).

To compare the actual versus the predicted test results, we used the following simple formula to estimate a theoretic download time of the Web pages:

Download time in seconds =

$$\frac{\text{Size of Web page in bits}}{\text{Bandwidth in bits per second}}$$

This formula does not consider factors such as protocol overhead and data compression, among others. However, the formula provides a useful approximate reference value. For example, the NLM's 70 KB (roughly 700 Kb) home page was typically downloaded by dial-up users using 28.8-Kbps modems in about 25 sec, which is consistent with the estimated value of  $700 \text{ Kb}/28.8 \text{ Kbps} = 24.3 \text{ sec}$ . Some testers were able to download the page a little bit faster, probably because of data compression by the modem, as illustrated by the BTC measurements shown in Table 5. The download time for two of the dial-up users (Information Research and Self Memorial) was usually longer than the expected performance based on BTC. We believe that this may be related to the high latencies of these dial-up connections.

The actual performance of some higher-bandwidth connections also was far below the predicted performance. For example, the downloading of the NLM home page from Tulane Medical Center should, under ideal conditions, take about  $700 \text{ Kb}/10 \text{ Mbps} = 0.07 \text{ sec}$ , but it took 2 sec. However, this download time is consistent with the BTC measurement of 321 Kbps for this site:  $700 \text{ Kb}/321 \text{ Kbps} = 2.2 \text{ sec}$ . In general, all non-dial-up links underperformed. We believe that this slower performance was related to the shared nature of these links and their high congestion during business hours, as suggested by our reduced BTC and increased packet loss measurements.

The test data showed the advantages of reducing the size of the Web pages on Web servers and the load imposed by graphic components on the downloading of Web content over slow links. For example, for dial-up users the 70 KB NLM home page consistently downloaded much more slowly than any of the other NLM Web pages tested (which are all smaller than the NLM's home page), and the downloading of text-only information reduced the download time by up to 80 percent (Table 3). The data also show that despite the availability of very-high-speed Internet links at the two ends of an Internet link (e.g., the path between

the NLM and Cornell University Medical College), the effective network speed between these two endpoints can be considerably slower than the speed of these end links, and the communications can still be heavily affected by time-of-day performance variations. We believe that this behavior probably reflects the congestion in the Internet infrastructure between the two locations.

The data indicated that some Internet routes are often less stable than expected. For example, the route from the NLM to Cornell Medical College was stable for only 25 percent of the test observations (Table 5). Only about half the routes to the sites shown in Table 5 remained stable during all the observations. These excessive route changes are generally attributed to "route flaps" caused by suboptimal routing protocol behavior, network infrastructure failures, or "load balancing" strategies used to improve network performance. Route flaps present a significant problem in today's Internet and are the focus of current research.<sup>17</sup>

### Time-of-day and Day-of-week Variability

The third contribution is further confirmation of preliminary findings that Internet paths consistently show large time-of-day and day-of-week variability. The effective transmission capacities of network pathways are generally reduced during the business day. This appears to be directly related to the much higher level of use during the day, as is reflected in traffic patterns for use of the NLM's databases. The transmission capacities are the greatest during the night or early morning hours. Typically, for Web sites tested in the United States and other G7 countries, the effective transmission capacities are reduced on the order of 40 to 95 percent during local business hours compared with weekend and other off-peak hours. The data suggest that U.S. business hours do not drive the performance degradation of international Internet communications for users in the United States as much as the business hours in the international locations tested.

The implication is that, within certain limits, designing and building a significant margin of excess transmission capacity may help minimize peak-hour delays. The data suggest that even high-bandwidth (e.g., 10 to 45 Mbps) Internet pathways may suffer from significant capacity reductions because of traffic congestion during peak hours, probably reflecting a loaded Internet infrastructure in some geographic regions (see Figure 5, for example). The performance of Web servers and other information services and the bandwidth of Internet links are important but partial



factors in the end-to-end performance of Internet communications.

## Conclusions

This research on evaluating Internet connectivity has laid a groundwork for exploring what is required to ensure that biomedical information systems, which are increasingly dependent on the Internet, perform adequately in meeting user needs. This challenge will only get more difficult as the volume of users and applications further increases and as the applications become more information intensive.

The initial Internet testing methodology and performance profile presented here should contribute to the concept of Internet quality of service. It will be important for the biomedical community to be an active participant in the process of defining and setting standards for the quality of Internet services, so that Next Generation Internet and its variations will be responsive to the needs of biomedical and health researchers, practitioners, and consumers. The research reported here focused on applications involving retrieval of biomedical information from Web-based computerized databases and files, such as those maintained by the NLM. This type of application constitutes the vast majority of Web use at the NLM and, we believe, most other biomedical information providers. However, the research needs to be extended to more information- and computation-intensive applications such as the transfer of large-scale human genome databases, virtual laboratories (where remote users electronically access and operate high-end biomedical laboratory equipment), and telemedicine (such as real-time tele-radiology or teledermatology consultation and diagnosis at a distance).

A fully developed testing methodology to diagnose and evaluate Internet performance requires further research. However, simple testing procedures, such as those used in this exploratory research, can be of immediate help to users. We would suggest that users begin with a series of download tests against the NLM Web sites and databases. The test results, when compared with the data in Tables 3 and 4 as a frame of reference, should provide a basic indication of relative Internet performance.

If performance looks sluggish, the user could then ask their system developers to conduct detailed technical tests. Traceroute testing should show the approximate path followed by the communications through different links and network devices and the stability of the path over time. Bulk transfer capacity testing should be able to measure the available bandwidth (over

time) through the complete path and to different points along the path. Ping testing should be able to determine the end-to-end connectivity of the Internet pathways and their approximate packet loss rates. These test results could be compared with the Table 5 data to identify possible performance problems. This in turn may provide a basis for consultation with Internet service providers, network specialists, and others who can assist in problem solving. Meanwhile, the NLM is researching further improvements in the metrics and methods for evaluating Internet end-to-end performance and providing meaningful reference information to its users and the broader Internet community.

Finally, this research has documented the extremely limited performance of some Internet links during peak business hours (sometimes 80 percent or more below the off-peak performance) and the importance of designing Internet access facilities and services with appropriate excess capacity. As the demand for Internet-based information resources increases, biomedical users can rightfully expect the Internet infrastructure to be adaptable and to meet new performance requirements. To help ensure that this happens, the evaluation of end-to-end Internet performance should be an ongoing and permanent activity within organizations making substantial use of and offering Internet services.

The authors gratefully acknowledge the support and encouragement of Donald A. B. Lindberg, MD, Director, NLM, during the course of this research. They have also benefited from the assistance of Karen Wallingford, MLS, and the other medical librarians who participated in end-user testing and user terminal testing. They also acknowledge the cooperation of other NLM colleagues: Dennis Benson, Larry Kingsland, Kathi Cansese, and Ed Sequeira. Finally, they appreciate the cooperation of the various domestic and international organizations whose Web sites were included in the testing program.

## References ■

1. Cimino, JJ. Beyond the superhighway: exploiting the Internet with medical informatics. *J Am Med Inform Assoc.* 1997; 4:279-84.
2. Lowe HJ, Lomax EC, Polonkey SE. The World Wide Web: a review of an emerging Internet-based technology for the distribution of biomedical literature. *J Am Med Inform Assoc.* 1996;3:1-14.
3. From 7 million to 70 million . . . what a difference a year makes. *Gratefully Yours.* Bethesda, Md.: National Library of Medicine, Mar-Apr 1998:4-5.
4. Wood FB, Wallingford KT, Siegel ER. Transitioning to the Internet: results of a National Library of Medicine user survey. *Bull Med Lib Assoc.* 1997;85:331-40.
5. Brakmo LS, Peterson LL. Tcp vegas: end-to-end congestion

- avoidance in a global Internet. *IEEE J Selected Areas in Communication*. 1995;13:1465–80.
6. Carter RL, Crovella ME. Measuring Bottleneck Link Speed in Packet-switched Networks. Boston, Mass.: Boston University, 1996. Technical Report BU-CS-96-006.
  7. Keshav S. An engineering approach to computer networking: ATM networks, the Internet, and telephone networks. Reading, Mass.: Addison-Wesley, 1997.
  8. Mahdavi J, Paxson V. Connectivity. Internet Engineering Task Force Web Site. Nov 1997. Available at: <http://www.ietf.org/internet-drafts/draft-ietf-ippm-framework-03.txt>.
  9. Paxson V. Toward a Framework for Defining Internet Performance Metrics. Berkeley, Calif.: Lawrence Berkeley National Laboratory, 1996. Technical Report LBNL-38952.
  10. Paxson V. End-to-end routing behavior in the Internet. *IEEE/ACM Trans Networking*. 1997;5:601–15.
  11. Paxson V. Measurements and Analysis of End-to-end Internet Dynamics [PhD thesis]. Berkeley, Calif.: University of California, 1997.
  12. Paxson V, Almes G, Mahdavi J, Mathis M for the Internet Engineering Task Force of the Internet Society. Framework for IP Performance Metrics. May 1998. Available from the USC Information Sciences Institute at: <ftp://ftp.isi.edu/in-notes/rfc2330.txt>.
  13. Welch V, Catlett C. Internet Performance Study. Champaign-Urbana, Ill.: National Center for Supercomputing Applications, University of Illinois, 1996. Available at: <http://notme.ncsa.uiuc.edu/People/vwelch/projects/inetperf/#details>.
  14. Mathis M, Mahdavi J. Diagnosing Internet congestion with a transport layer performance tool. *Proc INET'96*; Jun 15–28, 1996; Montreal, Canada. Reston, Va.: Internet Society, 1996. Available at: <http://www.isoc.org/conferences/inet96/proceedings/d3/d3-2.htm>.
  15. Lottor M. Network Wizards Internet Domain Survey. Menlo Park, Calif.: Network Wizards, 1998. Available at: <http://www.nw.com/zone/WWW/top.html>.
  16. Paxson V, Floyd S. Why we don't know how to simulate the Internet. *Proceedings of the 1997 Winter Simulation Conference*; Dec 1997; Atlanta, Georgia. Available at: <ftp://ftp.ee.lbl.gov/papers/wsc97.ps>.
  17. Labovitz C, Malan G, Jahanian F. Internet routing instability. *Proc SIGCOMM'97*; Sep 1997; Cannes, France. Available at: <http://www.acm.org/sigcomm/sigcomm97/papers/p109.html>.