

Reliability of Internet Hosts : A Case Study from the End User's Perspective

M. Kalyanakrishnan R. K. Iyer J. U. Patel

Center for Reliable and High-Performance Computing
Coordinated Science Laboratory
University of Illinois
1308 W.Main St, Urbana, IL 61801

Abstract

This paper presents the results of a 40-day reliability study on a set of 97 popular Web sites done from an end user's perspective. Data for the study was acquired by periodically attempting to fetch an HTML file from each Web site and recording the outcome of such attempts.

Analysis of the acquired data revealed: (i) 94% of the HTML file fetch requests succeed on the average. (ii) Most failures last less than 15 minutes. (iii) The underlying network plays a dominant role in determining host accessibility: (a) Network related-outages account for a major part of the failures. (b) Some network-related outages rendered more than 70% of the hosts inaccessible. (c) Host-related failures tend to be shorter than failures that might involve the network. (v) The network connectivity is high on the average with 93% of the sites being accessible at any given time. (vi) Mean Availability of the hosts is high (0.993).

1. Introduction

With the ever-increasing use and rapid expansion of the Internet, connectivity of a host from the user site, and the reliability of the connection are becoming increasingly important. This paper focuses on the *average* Internet user. (i.e. the majority of users that access various Web sites to obtain information from them.) Attempt is made to address issues such as: (a) What is the (stationary) probability that a user request to access an Internet host succeeds? (b) On an average, what percentage of hosts might remain accessible to the user at a given moment? (c) What are the major causes of access failures as seen by the user? (d) Typically how long could a host be unavailable to the user?

The term failure, in this context, refers to the inability to contact a host and fetch an HTML file from it. The failure might be due to problems with the host (e.g. the host being too busy handling other requests) or problems with the underlying network (e.g.

non-existence of a proper route to the host). Thus the failure is an integration of the behavior of the remote host, the intermediate routers, and the local network at the user site.

2. Related Work

The issue of Internet host behavior has been the focus of active research. Long et.al [1] evaluated Mean Time To Failure (MTTF), Mean Time To Repair (MTTR), Availability and Reliability for a sample of hosts on the Internet by repeatedly polling the hosts from different sites at exponentially distributed time intervals.

Paxson [2] [7] described in great detail the vagaries of end-to-end Internet routing. The *traceroute* [13] [14] program was extensively utilized to study, in detail, routing behavior in the Internet. Major routing pathologies and properties were identified.

Chimoy examined NSFNET backbone routing in [4]. The study was based on the trace information collected over a 12-hour period on all the NSFNET nodes on the T1 backbone. It was shown that only a small number of networks have a highly volatile connectivity to the backbone; the majority of the networks have a fairly stable connectivity to the NSFNET backbone. Paxson [5] obtained metrics that could characterize Internet performance. Arlitt et.al [3] have studied Web server workloads and obtained invariants that characterize it. Thakur et.al [8] presented a methodology to study failures in a LAN of Unix workstations. This methodology provides valuable insights on how to characterize, and measure Internet host reliability.

3. Data Collection and Analysis

We selected 97 among the top 100 Web sites, as rated by the PC Magazine [9], for the study. The data collected during our study comprises of success/failure logs corresponding to attempts to connect to, and fetch an HTML file from a Web site. On a failure, the

cause of the failure, if recognizable, was also recorded. A sample of the logs is shown below:

```
Server www.cnn.com time: 5-3-97 : 7:48:23;
Downloaded 22009 bytes successfully
```

```
Server www.columbiahouse.com time : 5-3-97 : 7:48:24;
server returned error 302, Trying with new URL:
http://www.columbiahouse.com/?86264561812817419329132317
Downloaded 11260 bytes successfully
```

```
Server www.americangreetings.com time : 5-3-97 : 7:48:35;
Connection failed with error code Connection timed out
```

```
Server www.egghead.com time : 5-6-97 : 6:17:25;
server returned error 500
```

The logs reveal the following information: (1) server (Web site) name (2) date and time of the attempt (3) result of the attempt. The sample above depicts the outcome of four separate access attempts, each to a different Web site, as explained below:

The first attempt, (*www.cnn.com*) was a success. The second attempt (*www.columbiahouse.com*) though successful, required a re-attempt as the HTTP server, first, returned an error response. The third attempt (*www.americangreetings.com*) was a failure due to a problem with connection establishment. The last entry in the sample depicts another failure (*www.egghead.com*), in this case, due to an error response from the HTTP server at the Web site.

The data collection was performed by a tool (a Perl script) that ran continuously at the test site (CRHC, University Of Illinois at Urbana-Champaign). The operation of the tool is described below.

The tool maintains a list of the Web sites under study. It periodically iterates through the list. During an iteration, the tool attempts to access each one of the Web sites, one after the other. Web sites that are found to be inaccessible during a regular iteration are handled by the tool as follows: Initially, the tool attempts to access each of these sites once every five minutes. If such an attempt succeeds, the frequency of access attempts to that Web site is reduced to the default rate (once every two hours). However, if failures continue, the frequency of attempts to access that particular Web site is gradually reduced until it reaches the original rate of once every two hours and then the whole process repeats. In addition, the first few re-attempts to access a Web site (i.e. after a failure was encountered) are supplemented by simultaneous runs of the traceroute program on the same site. This aids in detecting routing problems that might be the cause of failures.

The motivation behind this variation of rate of access attempts is to get a tight upper bound for the

duration of the failure and progressively "back off" if the failure seems to be of a long duration.

At the test site, there were 5 instances of the tool running simultaneously, each recording success/failure logs for a set of about 20 Web sites each. This split-up helped reduce the iteration time of the tool so that re-attempts to Web sites that failed could be made as soon as possible (within 5 minutes of the failure).

4. Overview of the Results

We commence the discussion with a brief examination of the parameters that, we believe, reflect the accessibility and reliability of Internet hosts as seen from the user's perspective. Table 1. lists these parameters and their corresponding measured (mean) values (The parameters were evaluated using the SAS [12] tools). In the remainder of this section, each parameter listed in Table 1. is briefly discussed with respect to what it represents, the implications of its observed values, and how useful it might be in reflecting the state of the Internet as a whole.

Table 1: Parameters to Describe Accessibility/Reliability of Internet Hosts

Parameter	Value
Average Success Frequency	94.4
Failure duration	43.022 min (mean)
Inter-failure time	4.016 days (mean)
Hourly Failure Rate	2.16 (mean)
Modes of failure	Connection Timeouts (42.8%) Connection Refusals (27.0%)
Mean Availability	0.993

The *Average Success Frequency* (ASF) for a Web site is defined as the percentage of HTML file-fetch attempts to that Web site that are successful. It functions as an indicator of host accessibility since the greater the ASF, the more accessible a host is. We obtained a high mean ASF indicating that, on an average, host accessibility is high.

The *Failure Duration* is the amount of time for which a particular remote host is inaccessible to the user. It represents the collective failure of the host and the network and is a measure of the ability of the host/network to recover. Though we observed the mean failure duration to be around 43 minutes, most (70%) of the failures lasted less than 15 minutes.

The *Inter-failure time* is the time period between two consecutive instances of failure of the same remote host. It provides an estimate of the degree of nonstop operation of the host. Though the mean value was around 4 days, individual sites differed greatly in their mean inter-failure time.

The *Hourly Failure Rate* is the percentage of Web sites that fail at least once during any given one hour

interval. It is a measure of the fraction of the network unavailable to the user at any time (the granularity being an hour). We observed that, on an average, only 2% of the hosts are unavailable at any time.

The *modes of failure* provide insights into the actual causes of the failures. Our study found that failures due to Connection timeouts and Connection refusals are the two most frequently observed failure modes. Also, majority of the connection timeout failures appear to be network-related.

The *Mean Availability* is defined as the ratio of inter-failure time (mean) to the sum of the failure duration (mean) and the inter-failure time (mean). It denotes, on an average, what fraction of the time a remote host is available. The data showed that the Mean Availability was very high with a small variation.

5. Average Success Frequency (ASF)

The ASF is a measure of the stationary probability of a given file-fetch request to a host succeeding. We initially obtained a mean ASF of 85.4. However, it was observed that 11 Web sites almost always responded with error codes *302/404* (which are HTTP-specific errors) or with a *Connection Refused* error message and hence they were excluded from further analysis. The Average Success Frequency for the remaining 86 servers is tabulated below:

Table 2: Measured ASF

Parameter	Value	Confidence Interval	Confidence
mean	94.4	± 3.4	50%
		± 4.3	75%
		± 6.3	90%
median	95.8		
standard deviation	5.6		

Thus the stationary probability of an HTML file-fetch attempt being successful is 0.944 (94.4/100). Figure 1. shows the distribution of ASF for this reduced set. The distribution appears highly skewed with most of the Web sites (95.3%) showing an ASF of 80 or more. A high percentage (83.7%) of the sites showed an ASF of 90 or more. A small fraction (about 4.7%) exhibited a lower ASF (less than 80) and were consequently less accessible than the rest. We also examined the distribution of the Average Failure Frequency, defined as:

$$\text{Average Failure Frequency} = (1 - \text{ASF}) / 100.$$

The distribution fitted an exponential distribution at a significance level of 0.05.

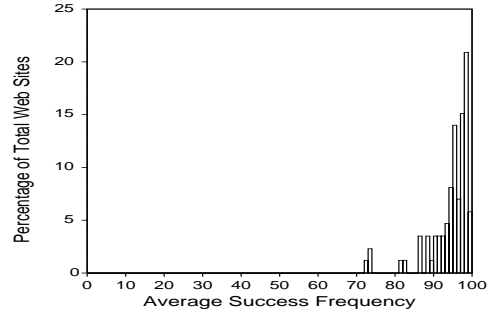


Figure 1: Distribution of ASF

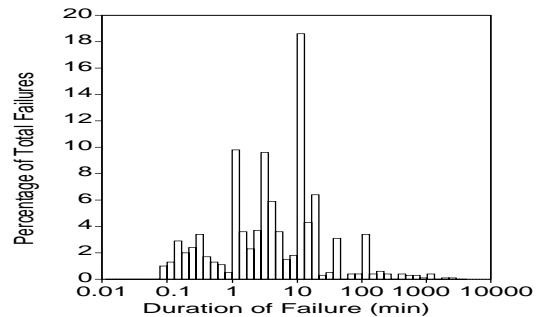


Figure 2: Distribution of Failure Duration

6. Failure Duration

The failure-duration parameter is an upper bound on the duration of an instance of failure. As discussed earlier, this parameter indicates the ability of the network/remote host to recover from a failure. In this context, each instance of failure represents the string of consecutive unsuccessful attempts to access the same Web site. (i.e. successive failed attempts were coalesced into a single instance of failure.) Thus the failure duration truly represents the time period for which the host is inaccessible. The table below (Table 3.) summarizes the results for the measured failure duration. The observed Mean Failure Duration (MFD) is distorted somewhat by a small fraction of failures that lasted much longer (a day or more) than the rest. The median however, is only slightly greater than 5 minutes. Figure 2. shows the distribution of the failure duration. More than a third (37.2%) of the failures lasted less than 5 minutes. Nearly a third (33.3%) of the failures lasted between 5 and 15 minutes. Approximately 18% of the failures lasted longer than 30 minutes. There were only 3 instances of failure that lasted longer than a day. The longest failure recorded lasted 3.375 days.

The peak at the 5-6 minute range was predominantly due to two kinds of failures : Failures due to

a connection refusal by the server, and failures due to a timeout¹ while reading the HTTP header that precedes the HTML file.

Table 3: Measured Failure Duration

<i>Failure Instances</i>	<i>Mean (min) / Confidence</i>	<i>Standard Deviation</i>	<i>Median (min)</i>
783	43.022 ±37.573 (50%) ±41.990 (75%)	233.504 min	6.667

In order to understand the failures better, we further categorized them based on the number of unsuccessful access attempts made during the failure period. Specifically: A *One-time failure* is defined as an instance of failure that involves only one unsuccessful access attempt. On the other hand, a Multiple failure is defined as an instance of failure that involves two or more unsuccessful access attempts.

We observed that 44.8% of the failures were Multiple failures. This would suggest that, once a failure is encountered, the probability of an immediate re-attempt being a success is not high. We examined the results of all re-attempts (on encountering a failure) made within 5 minutes of encountering a failure. Out of the 1268 such re-attempts, only 38.2% succeeded, indicating that the probability that an immediate re-attempt succeeds is very low. The value for MFD obtained in our study is significantly different from the values for the Mean Time to Repair (MTTR) obtained in [1]. The discrepancy, we feel, is due to the difference in the types of hosts evaluated.

7. Classification of the Causes of failure

Table 4. summarizes the distribution of failures based on the error message associated with the failure. In this case, as opposed to Section 6, each failed access attempt was considered as a unique instance of failure. (i.e. no coalescing of failures was performed.)

Table 4: Cause-wise Distribution of Failures

<i>Cause of Failure</i>	<i>Number of Failures</i>	<i>Percentage</i>
Connection timed out	1073	42.8
Connection refused	677	27.0
Could not read HTTP header within the timeout interval	531	21.2
No route to host	92	3.7
Server returned error 500	76	3.0
Server returned error 401	51	2.0
Network is unreachable	1	0.0

¹Once the tool successfully establishes connection with the HTTP server process on the remote host, it waits for a pre-defined timeout interval to receive the HTTP header from the server.

The following error categories are clearly related to the host:

1. Connection refused (The network was able to reach the host but the host did not respond favorably to the request.)

2. Server returned error 500/401. (Errors 500 (internal error) and 401 (unauthorized request) indicate a problem with the HTTP server on the host.)

The following error categories are clearly related to the network:

1. No route to host
2. Network is unreachable

Connection Timeout failures and failures due to timeout while reading the header could be due to either the host or the network. They are examined in detail below.

7.1 Connection Timeouts and Timeouts while Reading Header

Connection timeouts could arise due to (among other factors): a route establishment problem, the server having too many pending connect requests at that time, or congestion along the route to the server. Timeouts while reading header may occur due to sudden failure of the host/network or (more plausibly) temporary congestion along the route. To categorize these failures as a network-related problem or a host-related problem, we utilized the results of the traceroute runs. More specifically, we examined all traceroute runs to each Web site that were executed within 5 minutes before/after facing a connection timeout or a timeout while reading header. Table 5. shows the results of such traceroutes.

Table 5: Summary of Traceroute Runs Made while Experiencing Timeouts

<i>Parameter</i>	<i>Connection Timeout</i>	<i>Header-Read Timeout</i>
Total instances of failure	1073	531
Total traceroutes executed within 5 minutes of failure	659	436
Number of such traceroutes that failed to reach the host	434 / (66%)	100 / (23%)

Evidently, majority of the traceroute runs executed on facing Connection timeouts failed. Moreover, most of such failures occurred at one of the intermediate routers, (i.e. not near the host) clearly indicating a network problem. Thus majority of these traceroutes failed due to network-related problems. Since each such traceroute was executed within minutes of a Connection timeout failure, we concluded that the corresponding Connection timeout also failed due to a network problem. We extrapolate this result to cover

all the Connection timeouts. Thus about 66% of the Connection timeouts were network-related.

Considering the timeouts while reading the header, however, most traceroute runs in this case reached the host. Since a successful traceroute run implies proper functioning of both the host and the network, the actual cause of failure is not very evident. However, since the host did respond favorably to a connection request (The tool waits for the header only if it is able to successfully connect to the HTTP server.) a few seconds before the failure, and since such hosts (dedicated to maintaining the Web site) typically have some built-in fault-tolerance, a sudden host failure seems less probable. A more plausible explanation is a temporary network problem such as congestion along the route to the host that might have delayed the reception of the header. Thus with a small margin of error, we concluded that almost all of these failures were due to network problems (though the problem may not be as severe as in the case of Connection timeouts, which lasted much longer). Thus approximately 53% ($42.8 * 0.66 + 21.2 + 3.7$) of the failures were identified as network-related.

7.2 Comparison of the Major Classes of Failure

A comparative study of the three major classes of failure (i.e. Connection Timeouts, Connection Refusals and Timeouts while reading header) yields useful insights. We focus on the failure duration and the inter-failure time for each class of failure. (For this study, all consecutive access attempts that failed with the same error message were coalesced into a single instance of failure.) Tables 6 and 7 summarize the failure duration and the inter-failure time respectively.

Table 6: Comparison of Failure Duration for Different Classes of Failures

<i>Cause of Failure</i>	<i>Mean (min)</i>	<i>Median (min)</i>	<i>Standard Deviation</i>
Connection timeout	46.558	12.117	134.106 min
Connection refused	30.259	6.742	191.815 min
Timeout while reading header	17.486	2.725	76.714 min

Table 7: Comparison of Inter-Failure Time for Different Classes of Failures

<i>Cause of failure</i>	<i>Mean (days)</i>	<i>Median (days)</i>	<i>Standard Deviation</i>
Connection timeout	5.652	4.002	6.077 days
Connection refused	7.274	4.076	7.980 days
Timeout while reading header	2.695	1.080	3.882 days

The tables reveal the following:

1. Connection timeouts last longer than most other failures. Since majority of the connection timeouts are network-related, it appears that the network is relatively slow in recovering from failures as compared to the remote host.
2. Timeouts while reading the header lasted for a much shorter period than the other failures. (In fact nearly 37% of these failures lasted less than a minute.)
3. Connection Refusals occur less frequently when compared to the other two classes.

7.3 Summary of the Discussion on Failure Classes

To summarize, we observed the following:

1. Connection Timeouts (majority of which are network related) and Connection Refusals (host related) account for nearly 70% of the failures.
2. Network-related failures marginally outnumber host-related failures.
3. Network related failures last longer than the host related ones.

8. Mean Availability

To quantitatively describe the behavior of the hosts from the user's perspective, we evaluated the Mean Availability of the hosts. The Mean Availability is defined as the ratio :

(Mean Inter-Failure Time / (Mean Failure Duration + Mean Inter-Failure Time))

The Mean Availability evaluates to 0.993. To obtain a range for this parameter, we calculated the Availability for each Web site using the corresponding mean values of failure duration and inter-failure time for that web site. The results of this computation are shown in Table 8².

Table 8: Measured Availability

<i>Parameter</i>	<i>Value</i>	<i>Confidence Interval</i>	<i>Confidence</i>
Mean	0.993	± 0.005	50%
		± 0.006	75%
		± 0.067	99%
Standard Deviation	0.013		
Median	0.997		

Figure 3. shows a histogram of the availability distribution over all the 86 participant sites. The distribution is skewed with all the sites having an availability of at least 0.92. A significant percentage (40%) of the sites exhibited an availability of 0.99 or more. However, there was no site that had an availability of 1.00.

²It is seen that the mean of the availabilities of individual sites has the same value as the overall mean availability (computed using MFD and the mean inter-failure time)

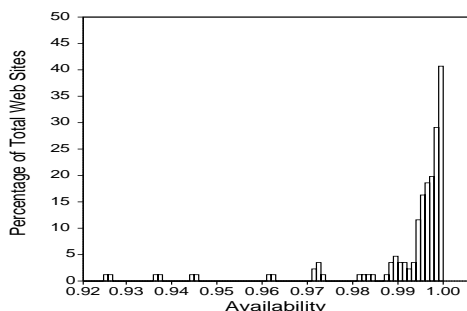


Figure 3: Availability Distribution

Comparing the results obtained here with the results obtained in [1], the hosts in our study seem to exhibit a higher availability than those in [1] (where the mean availability was 0.926).

9. Conclusion

We have performed a reliability study on a set of popular Web sites. The measurements were based on the results from actual data transfer requests, which are a more accurate reflection of host availability than an examination of the host being up or down. Our conclusions are summarized below:

1. Most of the user requests to the Web servers are successful, hence there is a high stationary probability (0.94) that a given request goes through successfully.
2. On an average, only about 2% of the servers fail within a given one hour interval.
3. Connectivity problems seem to play a major role in determining the accessibility of the hosts. Network-related failures tend to outnumber host-related failures. Also, the network appears to be slower in recovering from failures as compared to the hosts.
4. Majority (70.5%) of the failures are short (less than 15 minutes). However, a few failures spanned over a couple of days.
5. On an average, Web sites stay up for over 4 days without any failures though a good fraction of them fail at least once within a day.
6. The connectivity to the hosts seems to be good on the average. However, we did observe a few major network-related failures that rendered nearly 70% of the hosts inaccessible for a significant period of time.
7. The Mean Availability of the hosts is very high (0.993).

Acknowledgments

This work was supported by National Aeronautics and Space Administration under grant NAG-1-613, in cooperation with the Illinois Computer Laboratory for

Aerospace Systems and Software (ICLASS).

References

- [1] D.Long, A.Muir and R.Golding, "A longitudinal survey of Internet host reliability," *Proc. Symposium on Reliable Distributed Systems.*, IEEE Computer Society Press, pp. 2-9, September 1995.
- [2] V.Paxson, "An Analysis of End-to-End Internet Dynamics - A Partial Draft of a Dissertation to be submitted to U.C.Berkeley," University of California, Berkeley 1996.
- [3] M.F.Arlitt and C.L.Williamson, "Web Server Workload Characterization: The Search for Invariants," *SIGMETRICS '96*, pp. 126-137.
- [4] B.Chimoy, "Dynamics of Internet Routing Information," *Proc. SIGCOMM '93*, pp. 45-52, September 1993.
- [5] V.Paxson, "Towards a Framework for Defining Internet Performance Metrics," *Proc. INET '96*.
- [6] HTTP: A protocol for networked information, URL:<http://www.w3.org/pub/WWW/Protocols/HTTP/HTTP2.html>
- [7] V.Paxson, "End-to-End Routing Behavior in the Internet," *Proc. SIGCOMM '96*, August 1996.
- [8] A.Thakur and R.K.Iyer, "Analyze-NOW-An environment for Collection and Analysis of Failures in a Network of Workstations," *Proc. Seventh International Symposium on Software Reliability Engineering*, October 1996.
- [9] *PC Magazine*, Vol. 16, No. 4, pp. 101-124, February 18, 1997.
- [10] J-C.Bolot, "End-to-End Packet Delay and Loss Behavior in the Internet," *Proc. SIGCOMM '93*, pp. 289-298, September 1993.
- [11] K.Claffy, H-W.Braun and G.Polyzos, "A Parameterizable Methodology for Internet Traffic Flow Profiling", *IEEE JSAC*, 13(8) pp. 1481-1494, October 1995.
- [12] SAS Institute Inc., *SAS INTRODUCTORY GUIDE*, Third Edition, Box 8000, Cary, NC 27511-8000.
- [13] W.R.Stevens, *TCP/IP Illustrated, Volume 1: The Protocols*, Addison-Wesley, 1994.

[14] V.Jacobson, traceroute, <ftp://ftp.ee.lbl.gov/traceroute.tar.Z>, 1989.