Diversity in DNS Performance Measure

Deepak Kumar Shrestha *

Abstract

DNS is a critical component of an Internet infrastructure. Its correct functioning is paramount for various Internet applications. However, very few studies have been conducted on its performance and more so from the user's viewpoint. Only on 2002, an extensive study on DNS performance from a client's perspective was carried out at 75 different locations around the globe mostly from Americas and Europe. The study shed light on some of the underlying factors relating to the latency experienced by the user and some of the ways to improve the DNS performance. For independent validation of the previous results, new data was collected from KMITNB and a new investigation was carried out on a number of metrics in order to see if the new results support the key findings of the previous outcomes.

Keyword : DNS, performance

1. Introduction

The Domain Name System (DNS) is a distributed, dynamic, hierarchical database primarily used to resolve the human-readable domain names of remote machines to IP addresses.

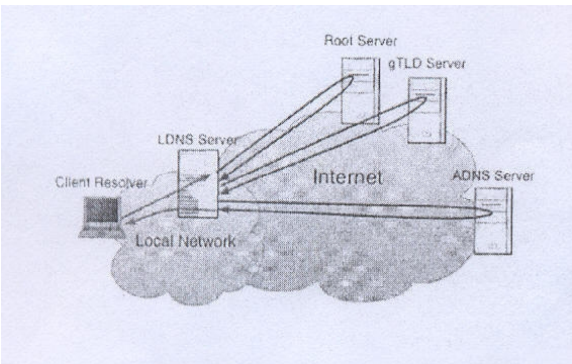

Fig. 1. Basic DNS operation

This resolution, called a domain name lookup, is typically the initial step in communication between two IP endpoints when the remote IP address is not known. Thus DNS is a critical component of the operation of many Internet applications [1] such as World Wide Web, Email etc. The domain name resolution process is shown in Figure 1.

The root and gTLD servers are crucial parts of the Internet Infrastructure. Until mid-2000 the 13 root servers also served as the gTLDs. As the growth of the Internet increased the workload on the roots, the .com, .net, .org domains were moved off the root system onto a separate layer of "global Top Level Domain" (gTLD) servers. Initially, there were 11 gTLD servers. Again, around mid-2002 an additional two gTLD servers were added in Atlanta and Seattle, US bringing the total number of those servers to 13. In Figure 2, each city that hosts root and gTLD server has a one-letter name in the bracket in the form (X:Y) where X indicates root servers and Y – gTLD servers.

Fig. 2. Placement of root and gTLD servers

One of the underlying design principles of DNS is its hierarchical database. The DNS system is structured to simplify administration of different portions of the namespace, or domains. In this way, IP addresses can be easily mapped and re-mapped by local administrators without having to notify a central authority. The organization of the namespace for domain names has a hierarchical structure, where each node in the tree has a label. Domain names are put together by walking down the namespace tree and at each node in the tree, prepending the label contained in the node with each step. A dot is used to separate each label. As an example, a small portion of the namespace tree is depicted in Figure 3.

This paper is organized as follows. In section 2, the background and importance of the problem is given. In section 3, the methodology of collecting a new data set is explained. In section 4, we discuss the results after comparing newly obtained results with the previous ones and conclusions and future work will be stated.

*Department of Information Technology, Faculty of Information Technology, KMITNB

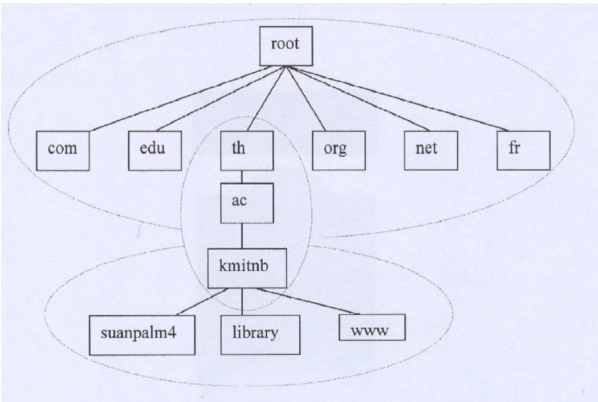

Fig. 3. Structure of the DNS namespace

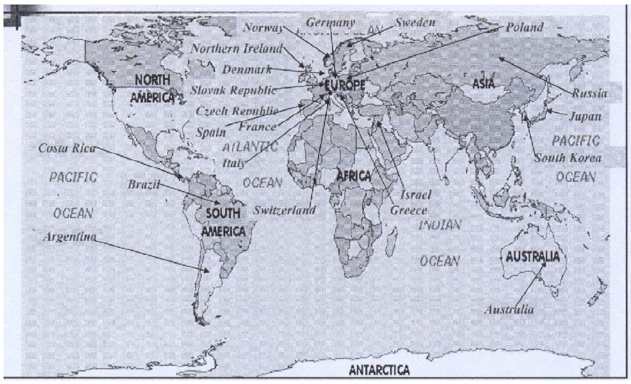

Fig. 5. Measurement locations (sites) around the globe

2. Background and Importance of the Problem

With the introduction of World Wide Web (WWW) in the early 90's, the Internet hosts have significantly increased over the last decade as shown in Figure 4 [6]. This incredible growth of the Internet can be attributed largely to two technologies one of which is DNS and the other is WWW.

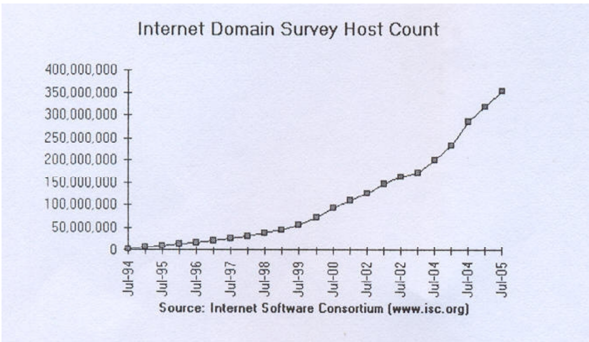

Fig. 4. Internet host count

As more and more people are using the Internet applications especially the www, users experience delays while browsing. This user perceived latency exhibits paramount concern in today's Internet driven economy with novel services ranging from real-time stock quotes to online merchandising. The success of these services is highly dependent on sufficient user-perceived network performance. Some studies were conducted solely on DNS performance [3], [7], [8]. These studies were based on measurements taken from one or two locations in the Internet Topology and did not focus on the performance from the client's viewpoint. Some other studies focused on evaluating web performance caused by the latency on the part of DNS. The most notable largest study on wide-area DNS performance from client's viewpoint was carried out in 2002 by Liston et al. [1]. Their study focused on the DNS performance at different client's locations while looking up the same non-cached domain names. They collected measurements from 75 locations around the globe as shown in Figure 5.

The data collection involved looking up around 15K second level domain names from each location. Their primary investigation was to find out the degree to which metrics for wide-area DNS performance such as mean response time, number of servers contacted and root & gTLD server performance, differ across locations in the Internet. Their results showed the followings:

- Response time from gTLD severs has much impact on the mean time to resolve non-cached domain names.
- Response time from root servers has trivial impact on the mean time to resolve non-cached domain names.
- The set of root and gTLD servers that provide the best service changes from site to site. There is more variation among gTLD servers that provide the best service than among root servers.
- There is a wide range in DNS performance for resolving non-cached names.
- The proportion of names that are aliases varies little across sites.
- The distribution of TTLs of completed lookups is not sensitive to location.

The detailed study on wide-area DNS performance study showed some important and interesting results. Without independent validation of the those results, the outcomes from the previous study could not be applied in the future reengineering of DNS as well as other global distributed systems. Therefore, verification of previous outcomes with a new one is very essential.

It is highly desirable that data be collected from as many different locations around the world as possible in order to truly reflect the performance of Internet protocols and services but this task poses serious challenges due to various difficulties associated with data collection procedures. These difficulties have prompted a new way of collecting data in which researchers develop an application that is of interest to general users and make it available for free download. The beauty behind this method of data collection is that while anonymous user runs the application for his own purpose, it also sends valuable network performance data in the background to the far away central server. Thus this new method of collecting data is very promising however the success of such methods are yet to be seen [5].

3. Methodology

As the primary task of this project was to see whether the newly obtained results from a new site support the key findings of the previous study, the tools needed to collect the new data namely a modified name server, a script file and a list of 15K domain names from the previous study were also used in this project as well. The detailed steps of collecting the new data set and analysis of the log file for investigating various metrics are explained below

3.1 Collecting a new data sets

The new data set was collected using a modified ***named*** name server [2]. This is the main tool that logs each event in finer details than the normal ***named*** name server does which is essential for subsequent analysis of the log file. Each name lookup logs the following information in a log file:

- Receipt of a request from client to resolve a name
- Sending of a request to remote servers
- Receipt of responses from remote servers
- Answer sent to the querying client
- Removal of queries from an internal queue of pending queries
- Identification of an entry in the local cache
- Identification of the type of response by the server

In order to lookup domain names, a script file is used which drives the data collection process. The script file[2] acts as a client resolver that sends name lookup request to the modified name server which then queries different name servers to resolve the name. If the resolver does not receive the response after 5 seconds, it times out and retries the request for lookup again.

The names to be looked up are the 15k domain names that were used in the previous study. The names are unique second level domain names that were filtered from about 100,000 domain names obtained from crawling the web.

The steps to collect new data set involved the following steps:

i) Installing modified DNS software - The modified DNS that was used to collect data is a BIND software. The software was installed on one of the machines in a KMITNB lab.

ii) Running the name server - The name server was run continuously for about 4-6 hrs without any interruption. During the operation, about 15K predefined web sites were looked up and the route to different servers during lookup of each site was logged in a log file.

3.2 Analysis of the log file

The log file was analyzed to investigate the following metrics.

- Completion and success rate – This is a first and important metric to be investigated which gives the rate of success of the looked up domain names. Investigation of other metrics are based on the completed domain names i.e. the combination of successful and unsuccessful domain names.
- Mean response time for completed lookup (MRTc) – This metric finds out the mean value of the response time i.e. total time taken to resolve a domain name, of all the completed looked up names. The MRTc is affected by series of other metrics which are investigated as well.
- Connectivity – This metric finds out the mean value of the Minimum Response Time for Completed lookups (MINc). It is considered that minimum round trip time from the closest name server from among the name servers contacted during name lookup process gives the status of connectivity of a site. The lower the MINc, the better the connectivity is.
- Loss rate – The loss rate is caused by a remote server's failure to respond to the query within a specified time period in which case the lookup process is repeated. Such loss of time during name resolution process might affect the MRTc. This metric finds out the number of retries that have occurred in the critical path.

Critical path is the shortest path among the paths leading to the answer which may be the case when LDNS receives multiple answers to a query to resolve the domain name. Such a scenario is shown in Figure 6.

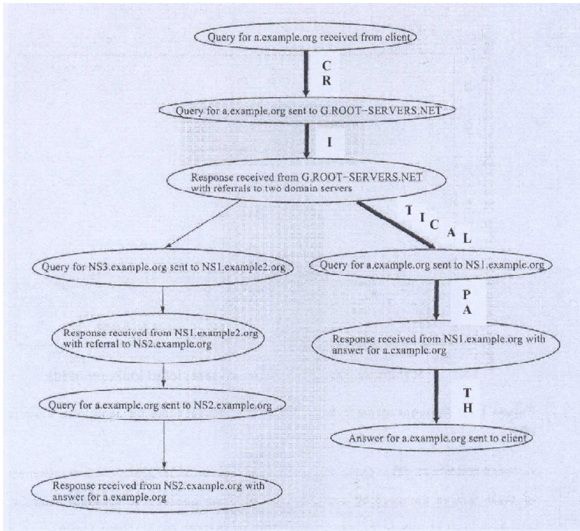

Fig. 6. Resolution tree and critical path

- Root/gTLD server performance – The performance of root/gTLD servers may have an effect on MRTc as these servers are central to resolving the domain names. This metric finds out the Mean Response Time for all queries sent to root servers (MRTr) and Mean Response Time for all queries sent to gTLD servers (MRTg).
- Root/gTLD server interactions – This metric finds out the root and gTLD servers favored by our site by investigating the distribution of queries from a site to the root /gTLD servers.
- Aliases and CNAMEs – Domain name that has an alias is indicated by CNAME (Canonical NAME) in a server's response message which is a static value set by domain administrator. This metric finds out the domain names that have an alias and whether a given CNAME acts as an alias for other CNAMEs and if an alias is assigned to different CNAMEs.
- TTLs of completed queries – Time To Live (TTL) value is a static value and is set by a domain administrator. This metric investigates the distribution of TTLs of completed queries on our site.

4. Discussion of results

This section reports results obtained from analyzing the data collected from the new KMITNB site after investigating several metrics such as completion and success rate of resolving names, mean response time for completed lookup, root and gTLD server's performance, root and gTLD server interactions as well as CNAMEs given for aliases and distribution of TTLs are presented. The newly obtained results are compared with the previous outcomes [1] and conclusions are made.

4.1 Completion and success rate

The previous study was conducted in 75 sites between January and March of 2002.

Table 1 Complete and success rate (old)

Total names probed	14.983
Successful lookups	13,900 – 14,200 (92.7% - 94.7%)
Completed lookups	14,500 – 14,700 (96.4% - 98.1%)
Failed lookups	450 (3%)

The new study was conducted in KMITNB on August 2005.

Table 2 Complete and success rate (new)

Total names probed	14.983
Successful lookups	13,152 (87.78%)
Completed lookups	14,646 (97.75%)
Failed lookups	337 (2.2%)

The previous study showed 0.6% of domain names went invalid over the period of two months and concluded that completed and successful lookups are time as well as location sensitive.

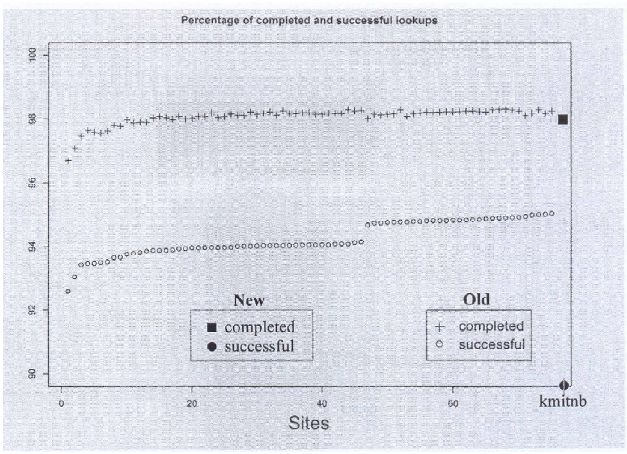

Fig. 7. Number of completed and successful lookups

The outcomes of the previous study [1] showed that the number of unsuccessful lookups increased from 90 (0.6%) to 1494 (9.97%) over the period of 40 months. However, those unsuccessful names were combined with successful lookups for the investigation of other metrics. This inclusion of unsuccessful names in the investigation of other metrics was based on further examination of the log file which revealed that almost all of the unsuccessful lookups went through the same lookup process for name resolution as for the successful lookups. This may have been due to the change only in the host portion of the domain names in the authoritative zone of the respective domain keeping intact the domain names in the root servers' database. Our study supported an assertion from the previous study that number of successful lookups for a static set of domain names is time-sensitive.

4.2 Mean response time of completed lookup (MRTc)

The previously conducted study with 75 sites showed MRTc ranging from 0.95 to 2.31 sec with most of the site around 1.5 sec which was a noticeable delay.

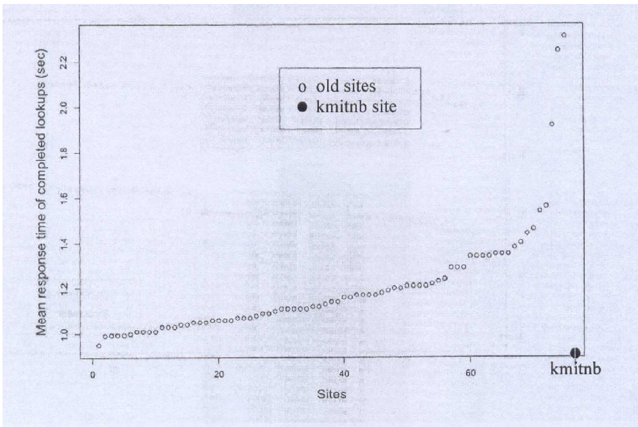

Fig. 8. Mean response time of completed lookup (MRTc)

Our site showed 0.71 sec MRTc which is much lower than the MRTc found in the previously conducted study at 75 locations. This significantly lower MRTc suggests that since 2002, the global Internet infrastructure has undergone significant improvement both in connection technology and server design. However MRTc alone can not be taken as a determining factor for a site's performance as it may have been affected by several other metrics which were further investigated.

4.2.1 Connectivity

The previous study showed that most of the 75 sites had a good connectivity however the weak correlation found between MRTc and MINc suggested that connectivity did not account for higher MRTc.

Our new site showed much higher MINc than the previous study. The higher MINc should not have affected the MRTc as the previous study conducted at 75 locations showed weak correlation between MRTc and MINc. Our conclusion is that the poor connectivity at our site may have attributed to higher MINc as there are large numbers of students using Internet through limited bandwidth.

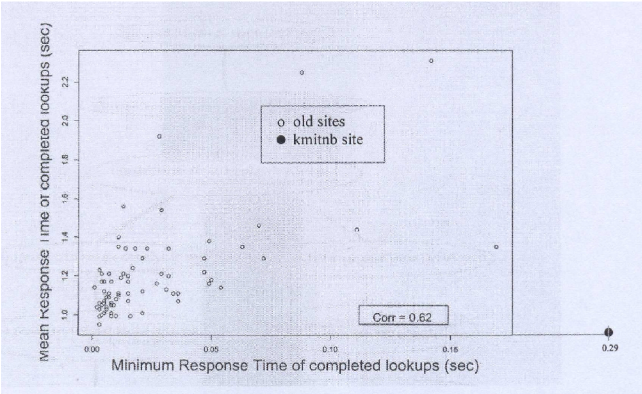

Fig. 9. Minimum response time vs. Mean response times of completed lookup (MRTc) at each site

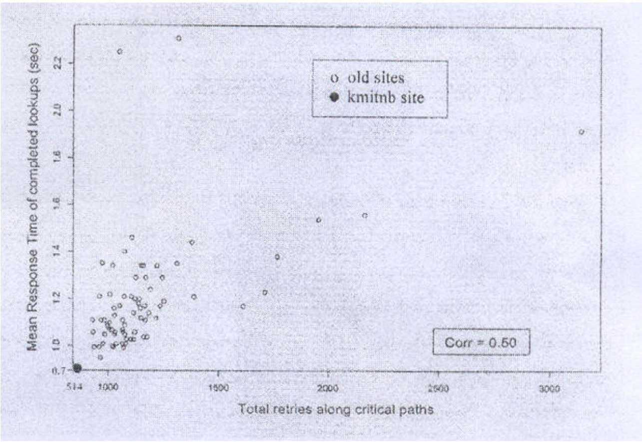

Fig. 10. Total retries along critical path vs. MRTc at each site

4.2.2 Loss rate

The loss rate is measured by the number of retries along the critical path. The previous study showed varying degree of loss rate and a weak correlation between MRTc and number of retries which suggests that loss rate should not affect the MRTc of a site.

In the new study, the loss rate dropped significantly to only 512 which is almost half of the lowest number of retries found in the previous study. However, this significantly lower loss rate should not have affected the MRTc of our site as the previous study showed weak correlation between MRTc and number of retries.

4.2.3 Root/gTLD server performance

The previous study at 75 sites showed strong correlation between mean response time of all root server queries (MRTr) & MRTc and mean response time of all gTLD server queries (MRTg) & MRTc. This prompted further investigation of root/gTLD server performance which showed that only small percentage of queries (7%) goes to root server, 60% to gTLD and 98.8% to other name servers. These findings asserted that performance of other name servers and gTLD servers have considerable effect on the client's perceived performance whereas root servers' performance have trivial effect.

The new study conducted at our site showed similar root and gTLD server's performance as in the previous study. This outcome prompted us to further investigate the performance of root/gTLD servers which showed that 95% queries went to other name servers, 54% - to gTLD servers and 40% - to root servers along the critical path. Similarly, 96% of time was spent on other servers, 47% - on gTLD servers and 24% - on root servers along the critical path. The seemingly high percentage of queries and time spent on root servers were because of the very small number of queries to root (only 20) found in the critical path as opposed to 1,397 queries to gTLD and 13,376 to other servers. As regards to gTLD and other servers, the percentage of queries and time spent on them along the critical path came out similar to the outcomes of the previous study.

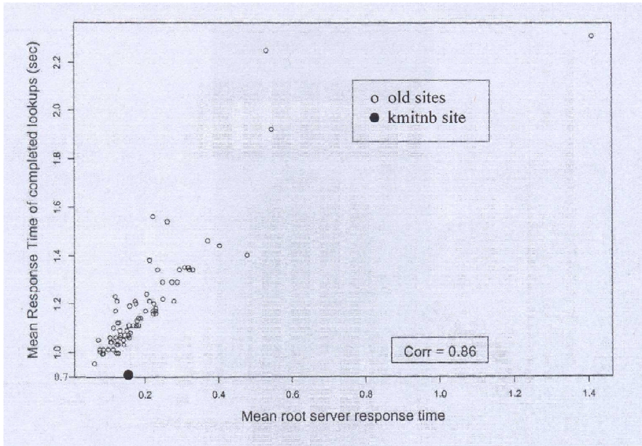

Fig. 11. Mean root server response time vs. MRTc at each site

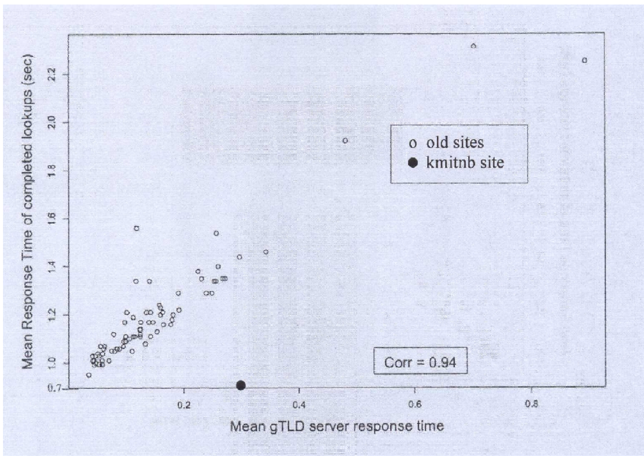

Fig. 12. Mean gTLD server response time vs. MRTc at each site

These findings evidently suggest that performance of gTLD and other name servers are very crucial and can reduce the user perceived latency considerably whereas performance of root servers are inconsequential.

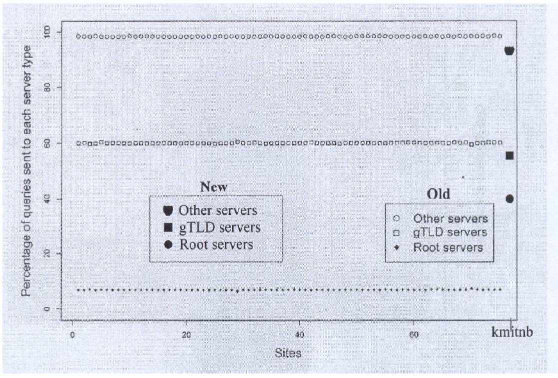

Fig. 13. Percentage of queries to root, gTLD and other servers along critical path

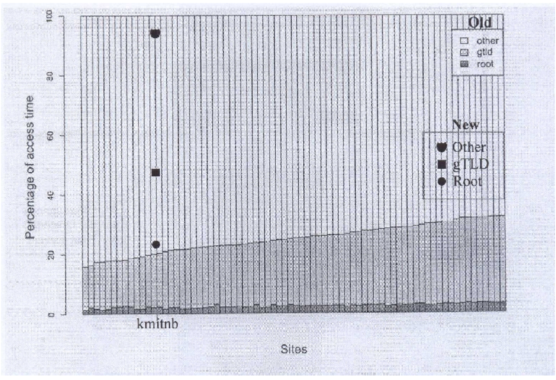

Fig. 14. Percentage of time querying root, gTLD and other servers along critical path

4.3 Root/ gTLD server interactions

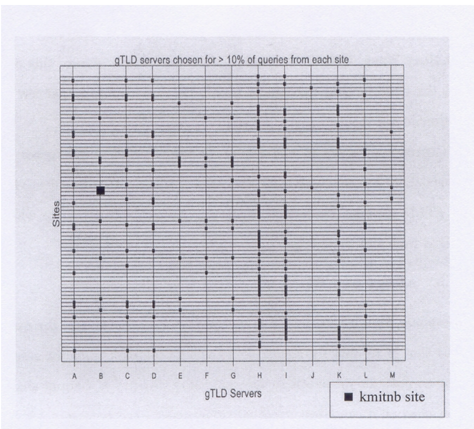

Fig. 15. Root servers favored by each site

The previous study at 75 sites showed four root servers A, D, H located in Washington DC area and I located in Stockholm favored by most of the sites. Likewise, the gTLD servers H and I located in Amsterdam and Stockholm favored by most of the sites. Other remaining root/gTLD servers were favored less. It was also found that there was more variation in which gTLD servers are favored from site to site than root servers.

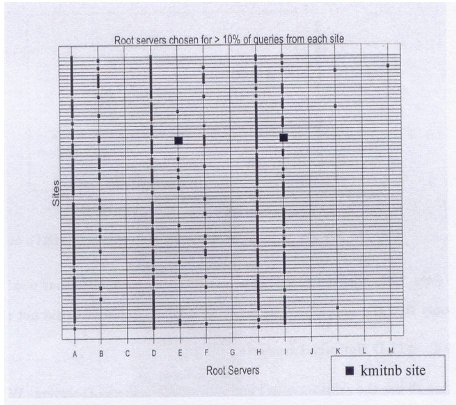

Fig. 16. gTLD servers favored by each site

The new study showed that our site favored root servers I as in the previous study and E which was not favored previously. As regards to gTLD servers, our site favored previously un-favored B gTLD server bypassing nearest J and M gTLD servers (see figure 2). Based on a single study from this part of the world, a general conclusion can not be made regarding the nature of server selection by a site.

4.4 CNAME

The previous study showed 26% of the domain names had alias and number of domain names with aliases varied only slightly across sites. The study concluded that names that are an alias is not a function of location.

Table 3 CNAME Redirections (old)

Number of redirections, X	Mean number (%) of CNAMEs with X redirections (old)	Mean number (%) of CNAMEs with X redirections (new)
1	3810 (96.3%)	2969 (99.76%)
2	138 (3.5%)	6 (0.20%)
3	8.77 (0.2%)	1 (0.03%)
4	1 (0.03%)	

Table 4 Number of different CNAMEs per alias (old)

Number of different CNAME mappings, X	Number of aliases with X different mappings
1	4230 (93.6%)
2	269 (5.9%)
3	13 (0.2%)
10	1
11	1
15	1
19	1

The new study at KMITNB site showed decrease in aliased domain names from 26% to 20% over the longer period of time. There were an insignificant number of names with 2 and 3 redirections as opposed to the previous outcomes and all aliases resolved to the same canonical name. Our findings suggests that CNAMEs being a static parameter which is set by site administrator tend to change over a longer period of time and the change is dictated by the management level of the organization.

4.5 TTL

The previous study at 75 sites showed that TTL values which are static parameters and set by a domain administrator is not a function of the location. The results from our new study initially showed that the number of domain names with TTL values which falls in most of the bins seemed to change over the longer period of time as shown in the figure 17 as opposed to the previous assertion that TTL values are not time and location sensitive.

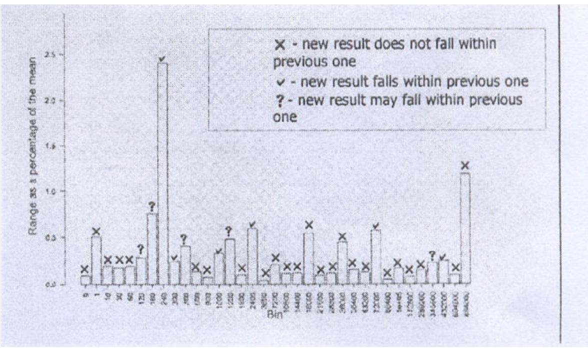

Fig. 17. Range of number of TTLs in each bin across sites, as a % of the number of TTLs in the bin

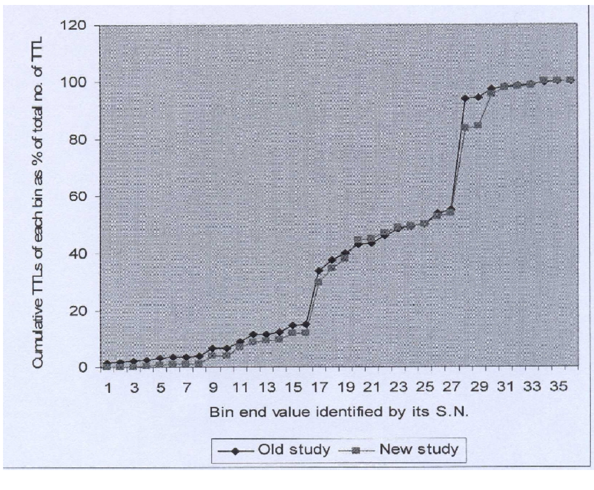

Fig. 18. Cumulative TTL values of each bin as a percentage of total no. of TTLs vs. bin end value

However, while looking at an overall trend of both new and previous study on TTL values, there has been only slight change seen as shown in the figure 18. This firmly suggests that distribution of TTLs seen at a site is constant as asserted by the previous study.

It can also be said that administrator tends to change the TTL value over longer period of time and is dictated by the management level of the organization.

Conclusions and future works

Our new study conducted at previously unrepresented part of SE Asia namely Thailand showed that key findings of the previous outcomes are supported by the new results. The outcomes of the studies both previous and new can be taken as a basis for the future improvement of the Internet infrastructure. However as data from 76 locations can not be considered representative for characterizing or verifying the performance of Internet protocols and services, it is desirable to collect and evaluate data from many more locations. Future work may involve gathering data not only from well-connected locations but also from a large and diverse set of actual locations throughout the Internet to actually uncover the additional factors that may have negatively affected on the user's perceived latency. Because of the difficulties associated with data collection procedures, the mode of collecting data has changed in recent years however the success of such methods have yet to be seen.

Reference

[1] Liston J., Sridhar S., Ellen Z., "Diversity in DNS performance measures" Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement, 2002.

[2] http://www.cc.gatech.edu/~liston/dnsperf.html, 10 August 2005.

[3] Danzig P., Obraczka K., Kumar A., "An analysis of wide-area name server traffic: a study of the Internet Domain Name System ", in Proceeding s of ACM SIGCOMM, January 1992.

[4] Brownlee N.. Nemeth E., "DNS root/gTLD performance measurement", Proceedings of the LISA 2001 15th Systems Administration Conference, San Diego, CA, USA, Dec. 2001.

[5] Liston J., "Measuring user-perceived Internet performance in multiple locations", A thesis presented to college of computing, Georgia Institute of Technology, July 2004.

[6] http://www.isc.org, 15 September, 2005.

[7] Brownlee N., claffy K., and Nemeth E., " DNS measurements at a root server:, in Blobal Internet 2001, November 2001.

[8] Jung J., Sit E., Balakrishnan H., and Morris R., " DNS performance and the effectiveness of caching", in Proceedings of the ACM SIGCOMM Internet Measurement Workshop, 2001.

[9] Barford P. and Crovells M., "Critical path analysis of TCP transactions", in Proceedings of ACM SIGCOMM, 2000, pp. 127-138.

[10] Brandhorst C., Pras A., " DNS: a statistical analysis of name server traffic at local network-to-Internet connections", Proc. 11th Open European Summer School and IFIP WG 6.4/6.6/6.9 Workshop CEds Kloos E.D. et.al, Spain, July 2005, pp.72-78.

[11] Mockapetris P., "RFC 1035: Domain names concepts and facilities," Nov. 1987.

[12] Mockapetris P., "RFC 1035: Domain names implementation and specification," Nov. 1987.