DNS Performance and the Effectiveness of Caching

Jaeyeon Jung, Emil Sit, Hari Balakrishnan, and Robert Morris
MIT Laboratory for Computer Science
Cambridge, MA 02139
{jyjung, sit, hari, rtm}@lcs.mit.edu

The Domain Name System (DNS) is an essential part of the Internet infrastructure and most Internet applications depend on the correct operation of DNS prior to establishing a connection to a remote host. We seek to understand the client-perceived performance and behavior of DNS and investigate the effectiveness of its caching mechanisms. Our goal is to identify the factors that affect DNS response latency, the errors and failure modes of DNS, and its scalability. An analysis of the effectiveness of DNS caching is especially important in light of several recent changes in the way DNS is used. Content distribution networks (CDNs) and popular Web sites with multiple servers are increasingly using DNS as a level of indirection to help balance load across servers, provide fault tolerance, or to direct each client request to a topologically nearby server.

Methodology

Our analysis is based on an extensive collection of packet traces. The novel idea in our approach is to jointly collect both DNS packets and associated TCP connection traffic: since TCP applications drive most DNS traffic, a joint trace collection where all TCP SYN/FIN/RST packets are collected together with DNS packets has the potential to allow us to infer various ways in which DNS is used. Without accompaning TCP traces, we might be able to infer things like DNS response latencies and failure modes, but not infer much about its caching effectiveness.

We have been collecting data at the border router connecting MIT's LCS and AI Lab to the rest of the Internet since Fall 1999. We have analyzed two weeks worth of data collected in January 2000 and December 2000. We also collected data from KAIST in Korea in Spring 2001. We have analyzed a week-long data from May 2001.

Our analysis has two parts: first, we study the packet traces to characterize DNS performance as seen by our clients, and draw more fundamental conclusions about its failure modes and its retransmission protocol. Then, we conduct simulations driven by these traces to explore the effectivness of DNS caching as a function of time-to-live (TTL) fields and degree of cache sharing.

We are currently investigating a mathematical framework to capture asymptotic properties observed by simulations.

Key Results

A paper detailing our methodology and explaining our results is available [1].

Our most surprising, non-obvious findings and conclusions are:

1. About a quarter of all DNS lookups never get an answer. More than 50% of the DNS-related packets in the wide-area correspond to such lookups.
2. The DNS retransmission protocol appears to be overly persistent: while most successful answers are received in at most 2-3 retransmissions, the lack of an answer or response causes a much larger number of retransmissions and a corresponding number of DNS packets traverse the wide-area.
3. Setting all A-record TTL's to a value as small as 10 minutes is not likely to degrade the scalability of DNS in any noticeable way. A-record caching does not seem to add much more to the per-host or per-application caching done by end clients today. Figure 1 show the simulation results. It shows the hit rate as a function of TTL. Since the results depend on the number of clients that share a cache, it includes separate curves for per-client caches, groups of 25 clients per cache, and one cache shared among all clients.
4. The scalability of DNS has little to do with its hierarchical organization or to good A-record caching; in fact, most of the DNS name space is a two-level structure. This suffices to partition the name space so that authority is handled by different name servers. The cacheability of these NS-records allows scalability by reducing load on the root and top-level name servers.

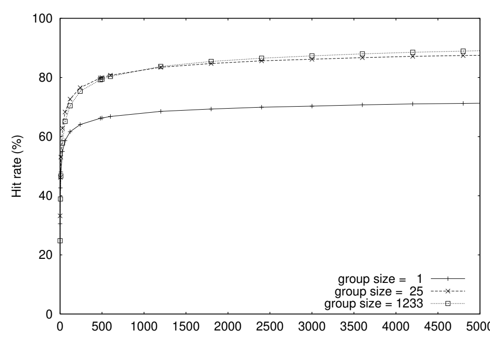

Figure 1: Impact of TTL on hit rate for the most recent MIT trace.

References

[1] Jung, J., Sit, E., Balakrishnan, H., and Morris, R. DNS Performance and the Effectiveness of Caching. In *Proc. 1st ACM SIGCOMM Internet Measurement Workshop* (San Francisco, CA, Nov. 2001). Project information available from `http://nms.lcs.mit.edu/dns/`.