Strong Cache Consistency Support for Domain Name System

Xin Chen, Haining Wang, Shansi Ren, and Xiaodong Zhang
College of William and Mary
{xinchen,hnw,sren,zhang}@cs.wm.edu

***Abstract*— Caching is critical to the performance and scalability of Domain Name System (DNS). However, DNS only supports weak cache consistency by using Time-To-Live (TTL) mechanism. Without strong cache consistency among DNS servers, it is not only cumbersome and inefficient to invalidate out-of-date cache entries, but also highly likely to lose connection services due to cache inconsistency. The loss of service availability is a much more serious problem than the service degradation. We propose an active DNS cache update protocol, called *DNScup*, to maintain strong cache consistency. Our trace-driven simulation and prototype implementation demonstrate that *DNScup* achieves the strong cache consistency of DNS, and hence, significantly improves its availability, performance and scalability.**

I. Motivation

DNS performance has been well studied. However, none of previous work focuses on the DNS cache consistency. The DNS cache inconsistency may induce the loss of service availability, which is much more serious than performance degradation. To investigate the dynamics of domain-name-to-IP-address (DN2IP) mapping changes, we have conducted a wide range of DNS measurements. The purpose of our DNS dynamics measurement is to answer the question about how often a DN2IP mapping changes. In general, a mapping change may cause two different effects. If the original DN2IP mapping is one to one, the change may lead to the loss of connection services. We call this kind of changes as physical changes. However, if the original DN2IP mapping is one to many, the change may be anticipated to balance the workload of a web site as CDN does. We classify these changes as logical changes.

II. Measurement

Since the DNS is predominately used by Web sessions to resolve the IP addresses of Web sites, our measurements are focused on the dynamics of the mappings between Web domain names and their corresponding IP addresses. We collected the Web domain names from the recent IRCache proxy traces between November 6, 2003 to November 12, 2003. Based on its TTL value, each domain name in our collection is periodically resolved to check if the mapping has been changed. What we have found are summarized as follows:

- While physical mapping changes per Web domain name rarely happen, the probability of a physical change per minute within a large number of Web domains is certain to one;
- Compared with the frequencies of logical mapping changes, the corresponding TTLs are set to much smaller values, resulting in redundant DNS traffic;
- The TTL value of a Web domain name is independent of its popularity, but its logical mapping change frequency does depend on the popularity of the Web domain.

Based on our measurements, we conclude that maintaining strong cache consistency is essential to prevent the potential losses of service availability, especially for critical or popular Internet services. Furthermore, with the strong cache consistency support, CDNs may provide fine-grained load-balance with reduced DNS traffic.

III. DNScup

To reduce the storage overhead and communication overhead, we introduce the dynamic lease technique for maintaining DNS cache consistency. A DNS cache keeps track of the local query rate of a cached DNS record. The authoritative DNS name server grants the lease of a DNS record to the DNS cache on-the-fly based on its query rate. The lease duration is determined by the record's DN2IP change frequency. Overall, the activities of *DNScup* include: (1) the query rate estimation at the cache-side; (2) the dynamic lease granting at the server-side; and (3) the communication between the DNS cache and the authoritative DNS server. We use trace-driven simulations to evaluate the effectiveness of dynamic lease of DNScup.

As an extension to DNS Dynamic Update protocol that maintains a consistency among a master DNS sever and its slaves within a single zone, we build a *DNScup* prototype with minor modifications on top of BIND 9.2.3. The components of *DNScup* implementation include the detecting module, the listening module, the notification module and the track file. In order to deal with the wide area DNS cache update propagation, we define a new type of message called `CACHE-UPDATE`. The interactions among all components are illustrated in Figure 1.

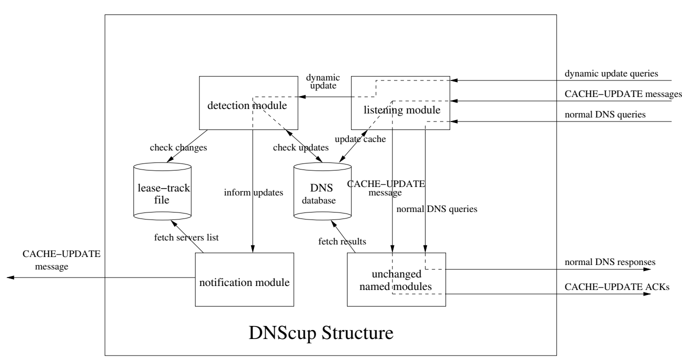

Fig. 1. The Structure of *DNScup* Prototype

Our trace-driven simulation and prototype implementation demonstrate that *DNScup* achieves the strong cache consistency in DNS and significantly improves its availability, performance and scalability.
URL: *http://www.cs.wm.edu/∼xinchen/DNScup.html*