Peeresolver: Distributed Name Resolution using Peer-to-Peer Overlay

Vivek Sawant, Jasleen Kaur
Department of Computer Science, University of North Carolina at Chapel Hill

Extended Abstract

A naming scheme is a key component of any distributed system. Names or references are *mapped* to objects or their locations and vice versa using a *name resolution service*. The Domain Name Service (DNS) serves this purpose for the Internet. Primarily, DNS provides a mapping between domain names and IP addresses of Internet hosts. The hierarchical DNS domain names are resolved using a hierarchical service structure of DNS *name servers*. In current DNS design and deployment, these two hierarchies are closely tied to each other. This has at least two undesirable consequences. First, the hierarchical service structure creates dependencies between name servers serving different components of a domain name. These dependencies are configured and managed manually, and hence are prone to human errors [1]. Second, due to mostly one-to-one or many-to-one association between domain names and their DNS name server, the service often does not scale well for domain names that see high rate of resolution requests. Moreover, this association opens up DNS to denial of service (DoS) attacks as demonstrated by the October 21, 2002 DNS DoS attack [1].

We propose Peeresolver, a generic reference resolution infrastructure that uses two key mechanisms to address these limitations (i) separation of resolution service structure from the reference naming scheme, and (ii) use of complementary replication policies. The scalable storage and resolution of references in Peeresolver is based on a Peer-to-Peer (P2P) Distributed Object Location and Routing (DOLR) service, which makes the service structure independent of the naming structure of Peeresolver applications. To support scalability for load-intensive applications such as DNS, Peeresolver design supports application-customizable replication. This is made feasible by the use of DOLR service that enables insertion (*publishing*) of a reference at multiple Peeresolver nodes, and efficient routing of the lookup request to the closest copy of the reference, in the P2P overlay. In addition, Peeresolver offers the following *replication policies*.

- **Load-balancing Replica:** To provide scalable resolution of references that are in high demand (e.g. root/gTLD.)
- **Popularity Replica:** To respond to sudden rise in the request rate for a reference by localizing its resolution.
- **Home Replica:** To provide continued resolution of local references in the face of a network partition.
- **Root-backup Replica:** To provide quite recovery from a Peeresolver failure.
- **Random Replica:** To improve resolution efficiency for less popular references.

It is worth noting that a recent evaluation of a P2P overlay-based design for DNS name resolution, called DDNS [3], concluded that the relatively high number of RPCs per lookup in P2P overlays make them an unattractive alternative for the current DNS infrastructure. However, DDNS does not support or evaluate flexible replication of references. Our current focus is on using Peeresolver to evaluate the effectiveness of reference replication for efficient reference resolution for an application like DNS. Peeresolver is being implemented over Tapestry. In evaluation experiments to begin shortly, we will analyze the effectiveness of the above replication policies. The system will be subjected to realistic distribution of DNS lookups obtained from traces of real DNS traffic taken at UNC-Chapel Hill egress link.

[1] Paul Albitz and Cricket Liu. *DNS and BIND*. O'Reilly & Associates, 1998.

[2] Paul Mockapetris, Keeping ahead of DNS attacks, *ZDNet News Commentary*. http://zdnet.com.com/2100-1107-979650.html. January 8, 2003.

[3] Cox, R., Muthitacharoen, A., Morris, R. T. Serving DNS using a Peer-to-Peer Lookup Service. *In the proceedings of the First International Workshop on Peer-to-Peer Systems (IPTPS '02)*, March, 2002; Cambridge, MA.