The Effectiveness of DNSSEC Under a Partial Deployment

Sruthi Rangavajhula, Thilina Buddhika, Zhao Zhao, Jun Li
Network and Security Lab, University of Oregon

BACKGROUND

The Domain Name System is an integral part of today's Internet which is responsible for conversion of domain names to IP addresses. An attack on any component of the DNS infrastructure can severely impact the proper functioning of the Internet. Among various attacks on DNS infrastructure, cache poisoning attacks are very prominent. In a DNS cache poisoning attack, attackers exploit certain characteristics of DNS protocol, for instance its connectionless nature to inject a malicious DNS response into a cache of a DNS resolver. As a result, subsequent DNS requests served by the infected DNS server will return malicious DNS responses. In a more sophisticated attack, attackers are able to retain the malicious responses within the DNS resolvers for an extended period of time.

DNSSEC addresses this issue by introducing origin authentication and data integrity to the DNS responses. Based on public key infrastructure, DNSSEC introduces digital signatures and relevant key information fields to a DNS message. By using this information, DNS resolvers can validate DNS responses before sending them back to clients and caching them for future reference.

MOTIVATION

Due to various complexities and challenges involved in deploying DNSSEC, the adoption rate of DNSSEC is low. To be effective, it requires a complete deployment from the root to individual zones in the DNS hierarchy. With a partial deployment, DNSEC cannot deliver the expected results. Deploying DNSSEC at a particular DNS zone does not necessarily assure it's not vulnerable to cache poisoning attacks. In our research, we are trying to explore how effective DNSSEC is against cache poisoning attacks under a partial deployment.

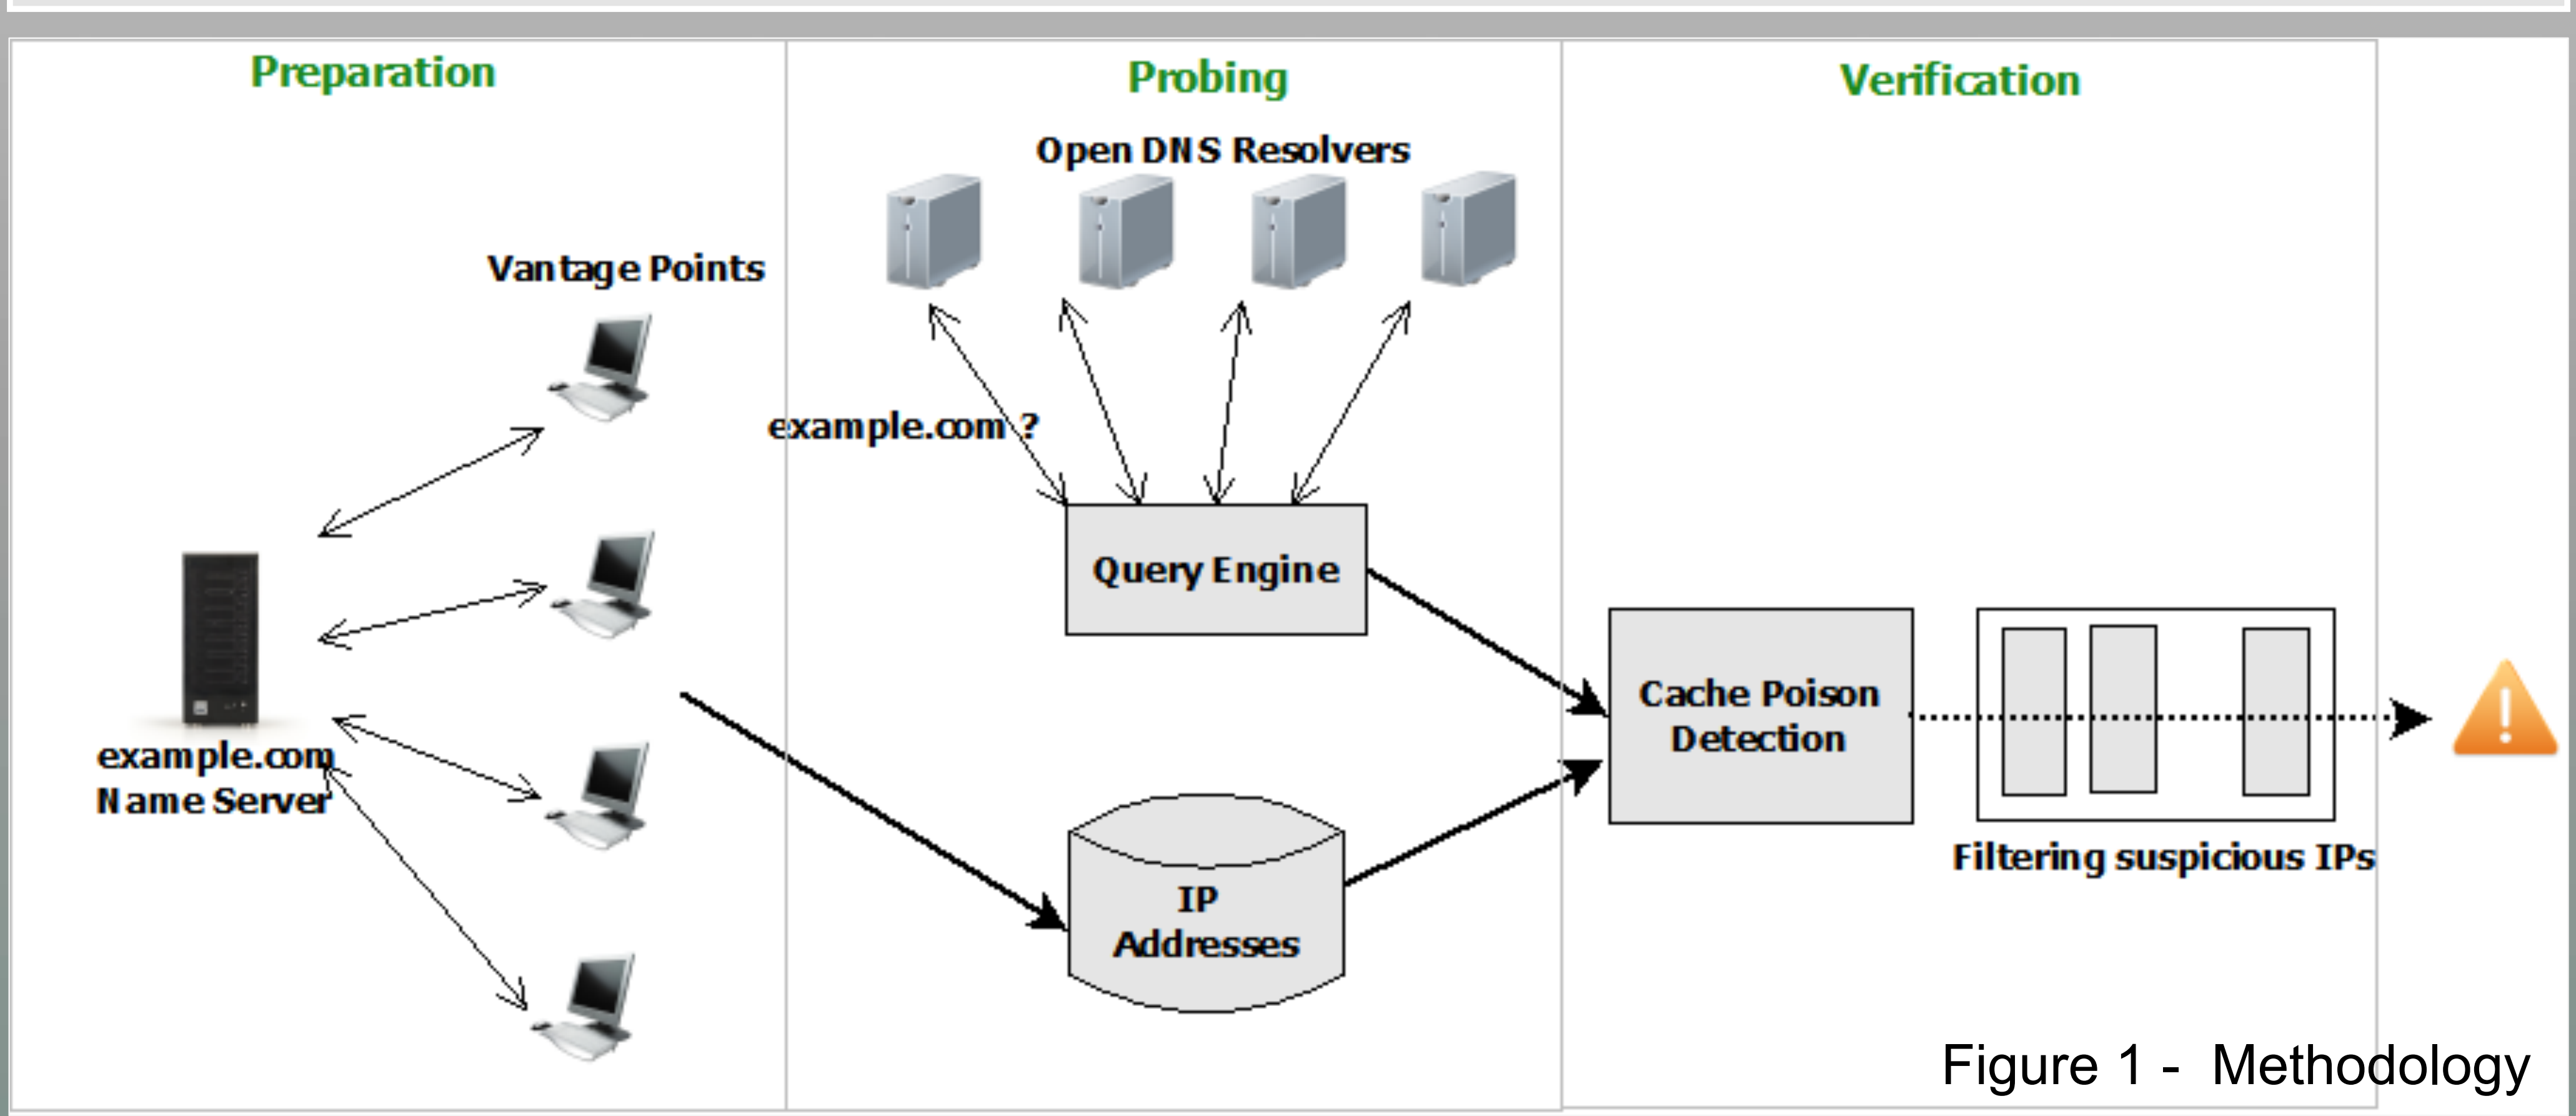

Figure 1 - Methodology

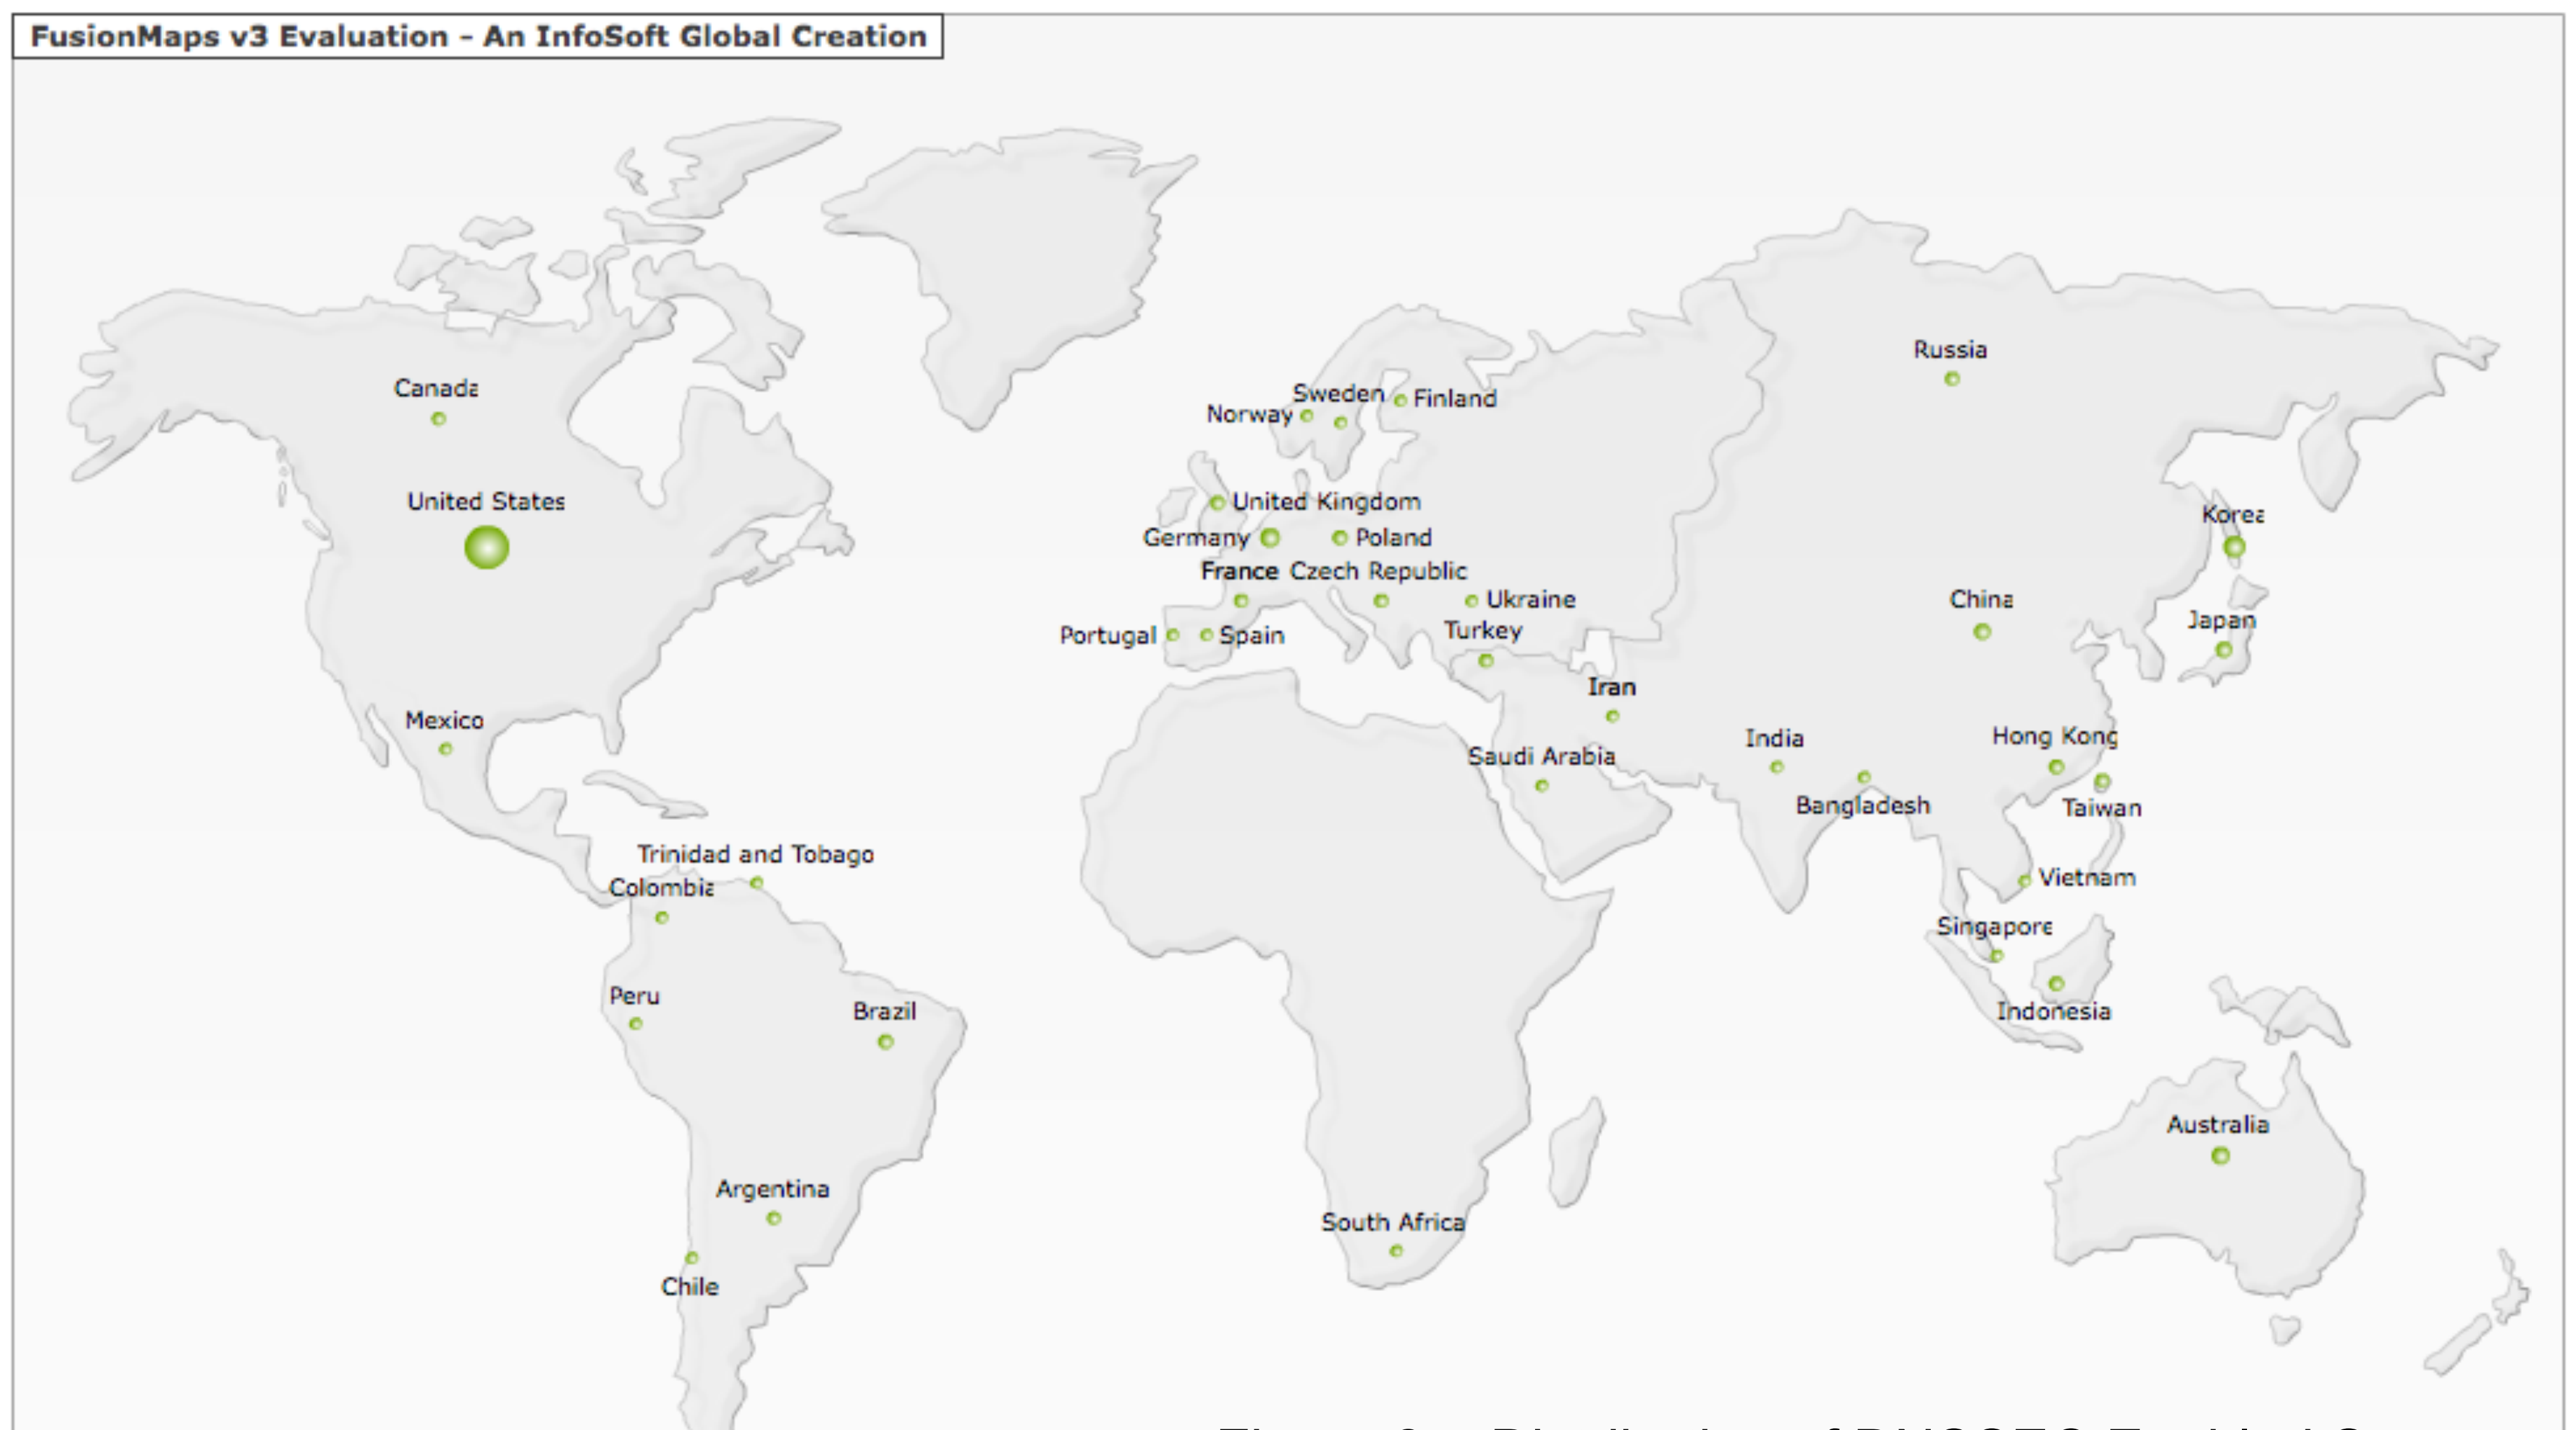

Figure 2 - Distribution of DNSSEC Enabled Servers

The size of the circle represents the number of DNSSEC enabled servers.

METHODOLOGY

First we identified a set of domains which can be potential victims of a cache poisoning attack. Using several vantage points located in multiple geographical locations, we have gathered valid IP addresses of the nodes used by each of those domains. We are querying over one million open DNS resolvers collected from our previous work to analyze the cached entries corresponding to the domains we have chosen. If this process could detect a suspicious IP, then it will be processed by a filter chain to eliminate any false positives. At the end of this experiment, we should be able to observe the degree of resilience shown by DNSSEC enabled name servers against the cache poisoning attacks.

CHALLENGES

- Reliably identifying the complete set of nodes used by a domain including nodes it uses from CDN providers and cloud vendors.
- Single probe provides only a snapshot of the cache. But it may contain poisoned entries in the long run.
- DNS responses can be altered due to DNS censorship, e.g. Great Firewall of China.
- Optimizing the filtering to minimize human intervention.

CONTRIBUTION

Based on the experiments performed on real world data gathered from a large collection of DNS resolvers, we will be able to provide an accurate approximation of the degree of cache poisoning that exists in today's Internet and the level of protection provided by the DNSSEC. The conclusion drawn on the effectiveness of DNSSEC under a partial deployment will provide incentives to accelerate its adoption or encourage community to look for an alternative.