Protecting DNS from Routing Attacks
-Two Alternative Anycast Implementations

Boran Qian StudentID 317715

Abstract

The *Domain Names System* (DNS) is an important role of internet infrastructure and supporting many of internet applications. Because it can translate the host names into IP addresses and organized in hierarchy. <<The backbone DNS servers are vulnerable to routing attacks in which adversaries controlling part of the routing system try to hijack the server address space [1].>> In this paper, I present a overview of two alternative anycast application of DNS server, they are operating at network layer and application layer.

1 Introduction

The *Domain Names System* (DNS) is an important role of internet infrastructure and translating the host names, such as ***www.tu-berlin.de*** into IP address for the internet community. If some applications can not receive the reply of the DNS server for their DNS queries, then they are denied services. In some other worse case, if an application received some wrong DNS reply, contains wrong IP address, it may send the date to the server, which selected by attacker.

<<Due to the hierarchical design, failure to reach all 13 root DNS servers would cripple the entire DNS service and make all destinations unreachable by most applications [2].>> And those 13 root DNS servers are used for the generic top level domains (gTLDs) including com, net and org. If we can not reach these gTLDs, we will failure to communicate millions of destinations in com, net, and org name domains.

Sometimes the attacker might prevent the backbone servers from answering legitimate requests or provide a fake response to the DNS requests. Worse still, the attacker control one or more routers and uses the routing system to hijack the address space of the victim servers. It will compromise availability and integrity of DNS server.

Server replication provides scalability by deploying multiple copies of a server and sharing client load across the copies. It offers a relatively straightforward method to potentially improve client performance and reduce network load. So do the DNS

server, replicate the server to improve the ability of top DNS server, large number of DNS requests from clients will be anycast to replicas, shown as figure 2. The send1 and the send2 send the query to serverA, and the Query will anycast to the “nearest” (or “best”) replicate.

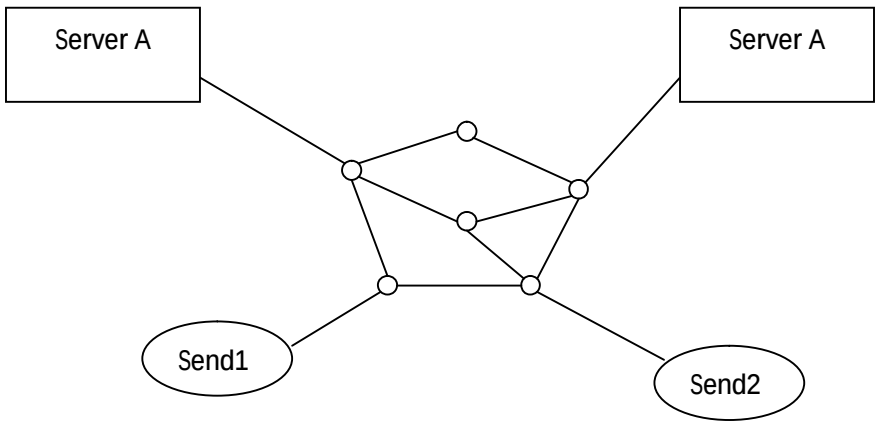

Figure1. Anycast Communication

There are two kinds of anycast implementations, used in practice. First is operating at network layer, second is operating at application layer. Network Layer anycast based on a simple implementation, can be implemented very easy. Application-Layer anycast must be implementing with an oracle able to distinguish legitimate from adversarial routes.

The filtering path technique guarantee the Reachability of the top level DNS server, different from that the Network-layer and Application-layer anycast can be considered as the countermeasure. Furthermore, we compare the secure performance of these two anycast implementation. The conclusion is that, with the help of Network-layer and Application-layer anycast we can solve the problem of protecting the Backbone DNS server from routing attack.

2 Background

2.1 DNS Overview and Terminology

Domain name system actually is a distribute database that map host names to IP addresses. The hostnames assigned by DNS, such as *www.tu-berlin.de*, called domain name. The DNS name space is organized in hierarchy, as shown in Figure 1. To answer each query, it is depending on the corresponding zone.

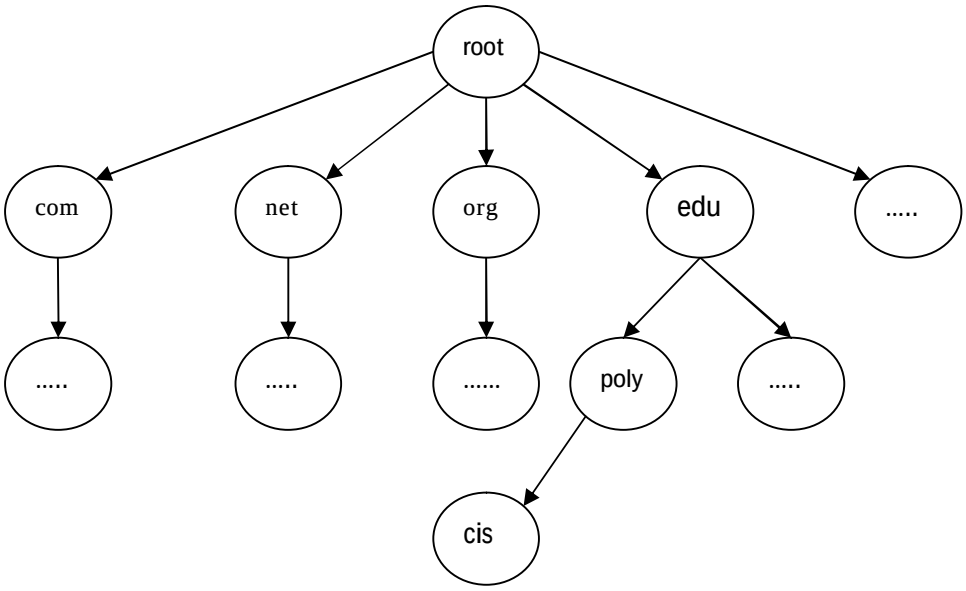

Figure2. The DNS Name Space Tree Structure

<<A DNS resolver requests data by first querying one of the root servers. Using the referral information from the root server the resolver proceeds down the tree until the desired data is obtained. [2]>> For example, a resolver is searching an IP address, www.poly.edu. At the beginning the host sends a query to any of root DNS server, and the root DNS server provides a referral to the DNS servers for the *edu* domain. A query to any of *edu* DNS server returns a referral to the *poly.edu* server. Finally a query to *poly.edu* server returns the IP address, which the host is searching, as shown in figure 3.

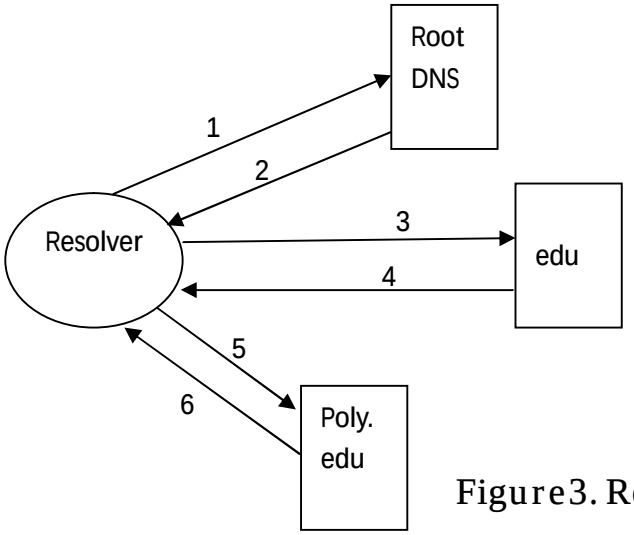

Figure3. Resolve Process

<<Each administrative entity managing a network may also manage its own portion of the DNS database stored in corresponding authoritative nameservers. However, unlike the authoritative nameservers that are responsible for answering queries about a small portion of the database, the root and top-level nameservers (operated by entities such as Verisign) must in principle be contacted for every DNS query made in the Internet.[1]>> Therefore, backbone nameservers are attractive targets for attackers.

2.2 BGP Overview and Terminology

The Internet is divided into thousands of Autonomous Systems (ASes), loosely defined as networks and routers under the same administrative control, and BGP [3] is the de facto inter-AS routing protocol.

2.3 Interdomain Routing

Resolvers and nameservers connect with each other on the network layer, so the route protocol is responsible for establishing the paths along which DNS traffic is forwarded. Certainly DNS traffic might be though thousand ASes. Route within AS is administered by inter-AS routing protocol, route between the ASes is administered by intra-AS routing protocol.(Border Gateway Protocol). The root and top-level nameservers are typically accessible through BGP paths, and normally those paths to the root and top-level nameservers are Stability. ASes are very rare to chane the primary path to root/gTLD servers.

Between ASes they announce block of IP address, called prefixes, and the BGP establishes the AS-path with it. A BGP announcement contains the sequence of ASes that must be traversed en route to the destination.

<<BGP selects routes according to AS-specific routing policies. These policies depend on multiple objectives including performance (as measured by the AS hop count) and the business relationships with neighboring network [1]. >>

3 Network Layer Anycast

Network layer anycast is, assign the common IP address to two or more endhosts. And the routing protocol routes datagrams to the closest server, using the routing distance metric. Standard intradomain unicast routing protocols can complete this job, assuming each server advertises the common IP address. BGP determine which target receives the traffic, and treat the anycast IP prefixes as any other regular IP prefixes. Certainly it will cause two or more alternative BGP route, and BGP route selection

process will handle it.

<<During the cross-continent attack, which lasted for about two-and-a-half hours on February 6 2007, unknown attackers used hijacked computers in the Asia-Pacific region to bombard six of the 13 root servers with data measuring a whopping 1Gb per second but, because the targets were using the Anycast technology, end users were not affected. [5] >> That proved network layer anycast effective in defending against DoS attack by exactly balancing the offending traffic across the anycast targets.

<<Network layer anycast also have limitation, it lack of flexibility in the selection criteria—the routing protocol determines the (single) criteria, typically hop count—and difficulty in extending to wide-area selection. [4]>>

4 Application Layer Anycast

Application layer anycast defines an anycast group with a set of unicast or multicast IP address. First, a set of servers may be grouped together based on equivalence *from a user's perspective*. That is, "exact" replication is not required for membership in the same group. A user might define an anycast group to contain. Second, allowing multicast IP addresses means we can support services that require multiple servers to provide a single instance of the service.

The client interacts with an anycast group via a query–response protocol as shown in figure 4. The query from client contains the *anycastdomainname* (ADN), which is used for identify the group, and the selection criteria to be used in choosing from the group. The anycast response contains the IP address for the selected server. The architecture centers around the use of hierarchy of *anycast resolvers* that perform the ADN to IP address mapping. The resolver receives the anycast query and applies a *filter* to control the selection. A filter operates on a set of anycast group members and returns a (possibly empty) subset. A second filter may be applied at the client. Filters may be content-independent (e.g., select any member at random), or based on performance metrics or policy information.

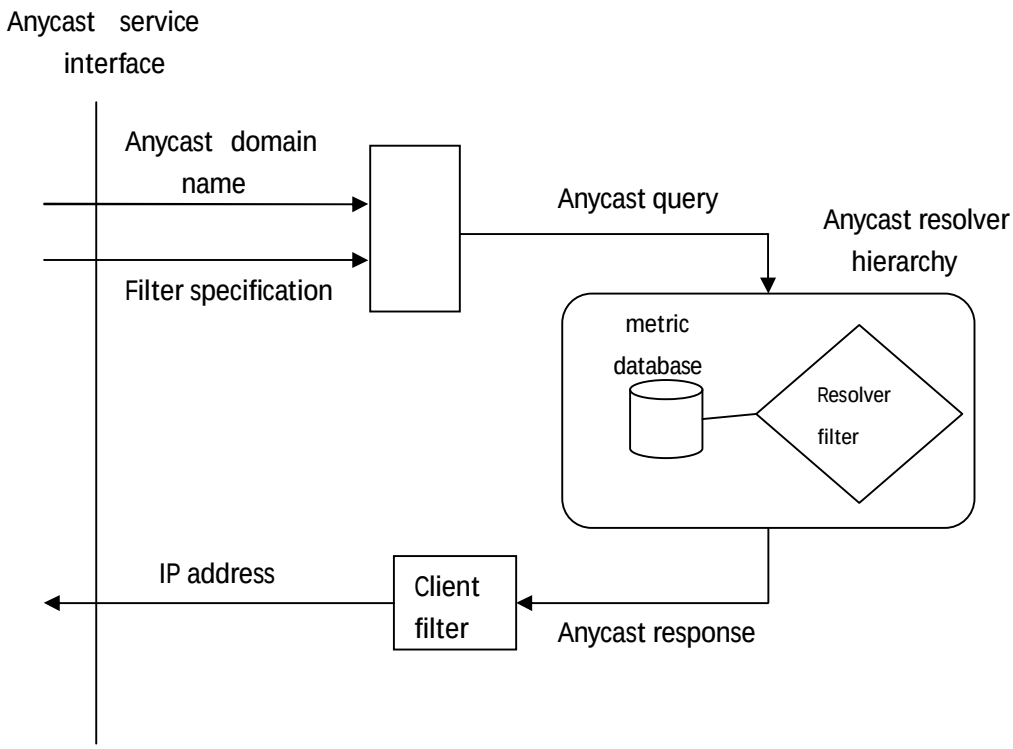

Figure4. Anycast name resolution query/response cycle

If the attacker wants to attack application-layer anycast, he must announce each one of those prefixes, because each participating AS announces a separate prefix. Against network-layer anycast, the adversary simply announces the anycast prefix.

5 Comparing the Security Performance of Network Layer and Application Layer Anycast

With the simulate experiments [1] we can evaluate the security performance of network layer and application layer anycast, as shown in figure 5 and 6. Those two figures show the percentage of ASes accepting a route toward a legitimate anycast target as a function of the size of the anycast group.

We find that, the security performance of network layer anycast is roughly equal to (the size of anycast group) / ((the size of anycast group) + (the size of the

adversarial group)). In application layer anycast, although each client is able to choose one of several possible routes toward the service by using a corresponding IP address as the destination address, more choice by the clients does not result in better security performance.

In application layer anycast without using an oracle, clients typically send the request the first IP address in the list, send by an authoritative DNS. It has poor security performance, roughly 1 / (1 + the size of the adversarial group). But differently if the client selects the target using an oracle that, given a route, is able to determine whether the route leads to a legitimate target or the adversary. Using this oracle a client selects a route to a legitimate target as long as such a route is available at the source. In this situation the security performance is very close to the security performance of network layer anycast as shown in figure 5, the difference being at most 0.1% at any given experimental configuration. Idea application layer anycast (using an orcale) has not only get better security performance, but also has some very hard problem to implement an oracle.

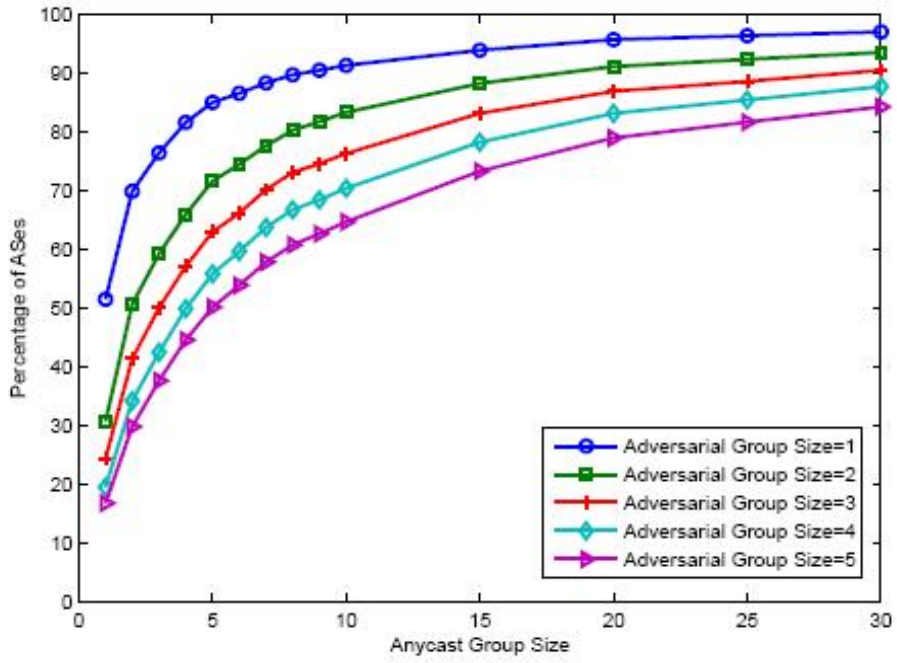

Figure5. Security performance of network-layer anycast

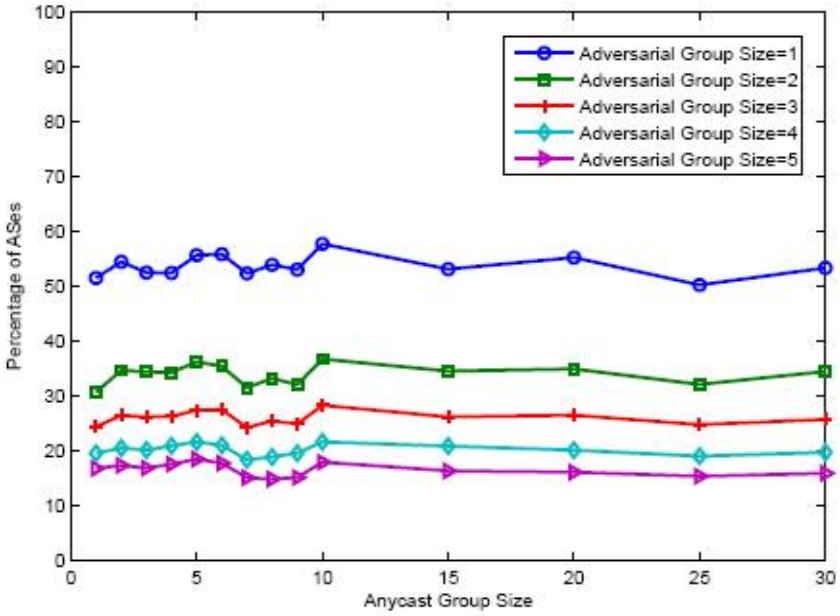

Figure6. Security performance of application-layer anycast

6 Conclusion and Future

In this paper I give a overview of two methods-network layer and application layer anycast for protecting DNS from routing attack, and compare the security performance of them. And the experiments also show that, network layer and application layer anycast both can achieve comparable security, network-layer anycast can be effective using a simple and practical implementation.

In the future I think the simple techniques such as round-robin (application layer anycast without using oracle) or nearest selection (network layer anycast) cannot accommodate the diversity of selection criteria that developing services will demand, because of the poor security performance of round-robin and low flexibility of nearest selection.

References

[1] Ioannis Avramopoulos. Martin Suchara. Protecting DNS from Routing Attacks: A Comparison of Two Alternative Anycast Implementations

[2] Lan Wang, Xiaoliang Zhao, Dan Pei, Randy Bush, Daniel Massey,Allison Mankin, S. Felix Wu, Lixia Zhang, Protecting BGP Routes to Top Level DNS Servers

[3] Y. Rekhter and T. Li. Border Gateway Protocol 4. RFC 1771

[4] EllenW. Zegura, *Member, IEEE*, Mostafa H. Ammar, *Senior Member, IEEE*, Zongming Fei, and Samrat Bhattacharjee Application-Layer Anycasting: A Server Selection Architecture and Use in a Replicated Web Service

[5] ***http://blogs.zdnet.com/security/?p=118*****.**