M. Ohta
Tokyo Institute of Technology

February 2004

Preconfigured DNS Server Addresses

Status of this Memo

Abstract

This memo describes how to make addresses of DNS servers preconfigured in clients. Well known anycast addresses, which are assigned to DNS servers, should be preconfigured to resolvers of clients. It is also explained why anycast does not offer redundancy.

1. Introduction

The requirement addressed by this memo is to eliminate configuration effort of DNS clients both for IPv4 and IPv6 network.

This memo defines several IP addresses to be used as a well known anycast IP addresses of DNS servers.

The addresses are intended to be preconfigured to client resolvers to reduce (or eliminate) initial configuration effort on clients.

The addresses may also be preconfigured to DHCP servers to reduce configuration effort of the DHCP servers.

This memo also describes sever and client operations.

2. Redundancy and Anycast

There is widespread misunderstanding on anycast (and multicast) in, including but not limited to, RFC1546 and RFC2461 that anycast (and multicast) could have provided meaningful redundancy or fault tolerance.

For example, it is wrongly stated in RFC 1546:

> One approach is to ARP for the anycast address. Servers which support the anycast address can reply to the ARP request, and the sending host can transmit to the first server that responds. This approach is reminiscent of the ARP hack (RFC 1027) and like the ARP hack, requires ARP cache timeouts for the anycast addresses be kept small (around 1 minute), so that if an anycast server goes down, hosts will promptly flush the ARP entry and query for other servers supporting the anycast address.

which makes treatment of anycast in RFC2461 unnecessarily complex.

However, anycast does not provide meaningful redundancy.

It is true that anycast and multicast tolerate some route faults.

However, a fault mode where a server process crashes on an anycast server a route to which is still alive, can not be tolerated. Clients keep asking to the original anycast server with no server process in vain. Even if the server replies with ICMP that the designated port is unreachable, the clients have no way to access other anycast servers sharing an anycast address with the original server, as the route for the anycast address is still directed to the original server.

Scalable multicast protocols, such as CBT and PIM with a statically configured multicast server such as core or rendez vous point, are no better, because the multicast server is the single point of failure.

Multicast protocols using broadcast extensively, such as DVMRP, certainly has redundancy, not because of multicast but because of broadcast.

With anycast or multicast, redundancy with no single point of failure can only be provided by using multiple anycast (or multicast) addresses served by different anycast (or multicast) servers.

Thus, it is meaningless that RFC1546 considers, because of

redundancy, a case where there are multiple anycast servers on a single subnet. Like unicast, it is a configuration error if there are two or more anycast servers sharing an anycast address in a subnet.

That is, IPv4 ARP works as is for anycast servers.

3. Server Operation

To serve stub resolvers, preconfigured DNS servers MUST offer recursive service if requested by their clients.

While it should be assumed that the clients have no information to validate their identities to the servers, the servers may still limit their services to valid clients based on some (autoconfigured) property of the clients. For (perhaps the only) example, the servers may reply to requests from clients with certain ranges of IP addresses only.

It is a configuration error if there are two or more anycast servers sharing an anycast address in a single subnet.

4. Client Operation

Clients SHOULD behave as a stub resolver forwarding queries to any of the preconfigured anycast address.

Client resolvers SHOULD NOT have SBELT (root server) information, because root servers and their addresses sometimes changes requiring reconfiguration.

Client resolvers MAY be able to perform recursive queries.

Clients SHOULD be preconfigured with three anycast addresses of DNS servers, as specified in section 5, and if query to an anycast address is not responded, they SHOULD try other addresses.

5. Assigned Addresses

The following three IPv4 addresses are assigned as preconfigured DNS server addresses:

<to be assigned by IANA>

<to be assigned by IANA>

<to be assigned by IANA>

route information of which should be propagated with a netmask of

<to be assigned by IANA>

The following three IPv6 addresses are assigned as preconfigured DNS server addresses:

<to be assigned by IANA>

<to be assigned by IANA>

<to be assigned by IANA>

route information of which should be propagated with a netmask of

<to be assigned by IANA>

A client resolver can receive a UDP reply from an IP address different from one used to send a corresponding query, that UDP reply may use source address of non-anycast IP address of server.

However, such a behavior is not available with TCP and anycast IP addresses MUST be used as source addresses of reply packets of TCP.

5.1. Scope Considerations

The assigned anycast addresses are globally well known and have global scopes.

On the other hand, servers with the anycast addresses have local scopes, boundaries between them is determined by a routing system.

The boundaries many be dynamically determined with routing metric.

In addition, route administrators may configure static boundaries.

While there may be static boundaries at site boundaries, always requiring them eliminates useful cases, for example, of ISPs providing DNS servers for customer sites.

Within a static boundary, there may be a single server for each anycast address, or there may be multiple servers sharing an anycast address, which may be useful for load distribution.

Even if identities of the servers changes dynamically between a query and its reply, a simple exchange of UDP packets is not affected and TCP transactions detect sequence number inconsistencies and should attempt new ones.

However, having multiple servers sharing an anycast address within a static boundary does not improve redundancy, because a DNS process on a server may die while a route to the server is still alive.

Three anycast addresses are provided for the redundancy.

6. Security Considerations

Cryptographic security requires some information securely (at least with authentication, even when public key cryptography is used) shared between servers and clients, which requires manual configuration.

Thus, no cryptographic security, which makes the protocol complex with no security improvement, should be used by preconfigured client resolvers.

The DNS server with the preconfigured addresses are still reasonably reliable, if local environment is reasonably secure, that is, there is no active attackers receiving queries to the anycast addresses of the servers and reply to them.

7. Author's Address

Masataka Ohta
Graduate School of Information Science and Engineering
Tokyo Institute of Technology
2-12-1, O-okayama, Meguro-ku, Tokyo 152-8552, JAPAN

Phone: +81-3-5734-3299
Fax: +81-3-5734-3299
EMail: mohta@necom830.hpcl.titech.ac.jp