NRO Document	G. Huston
	APNIC
	August 18, 2004

6to4 Reverse DNS Delegation

Abstract

This memo describes a potential mechanism for entering a description of DNS servers which provide "reverse lookup" of 6to4 addresses into the 6to4 reverse zone file. The proposed mechanism is a conventional DNS delegation interface, allowing the client to enter the details of a number of DNS servers for the delegated domain. The client is authenticated by its source address and is authorised to use the function if its IPv6 /48 address prefix corresponds to the requested delegation point.

---

1. Introduction

6to4 **[1]** defines a mechanism for allowing isolated IPv6 sites to communicate using IPv6 over the public IPv4 Internet. This is achieved through the use of a dedicated IPv6 global unicast address prefix. A 6to4 'router' can use its IPv4 address value in conjunction with this global prefix to create a local IPv6 site prefix. Local IPv6 hosts use this site prefix to form their local IPv6 address.

This address structure allows any site that is connected to the IPv4 Internet the ability to use IPv6 via automatically created IPv6 over IPv4 tunnels. The advantage of this approach is that it allows the piecemeal deployment of IPv6 using tunnels to traverse IPv4 network segments. A local site can connect to a IPv6 network without necessarily obtaining IPv6 services from its adjacent upstream network provider.

As noted in **[3]**, the advantage of this approach is that: "it decouples deployment of IPv6 by the core of the network (e.g. Internet Service Providers or ISPs) from deployment of IPv6 at the edges (e.g. customer sites), allowing each site or ISP to deploy IPv6 support in its own time frame according to its own priorities. With 6to4, the edges may communicate with one another using IPv6 even if one or more of their ISPs do not yet provide native IPv6 service."

The particular question here is the task of setting up a set of delegations that allows "reverse lookups" for this address space.

> "[This] requires that there be a delegation path for the IP address being queried, from the DNS root to the servers for the DNA zone which provides the PTR records for that IP address. For ordinary IPv6 addresses, the necessary DNS servers and records for IPv6 reverse lookups would be maintained by the each organization to which an address block is delegated; the delegation path of DNS records reflects the delegation of address blocks themselves. However, for IPv6 addresses beginning with the 6to4 address

> prefix, the DNS records would need to reflect IPv4 address delegation. Since the entire motivation of 6to4 is to decouple site deployment of IPv6 from infrastructure deployment of IPv6, such records cannot be expected to be present for a site using 6to4 - especially for a site whose ISP did not yet support IPv6 in any form." **[3]**

The desired characteristics of a reverse lookup delegation mechanism are that it:

- is deployable with minimal overhead or tool development
- has no impact on existing DNS software and existing DNS operations
- performs name lookup efficiently
- does not compromise any DNS security functions

---

2. Potential Approaches

There are a number of approaches to this problem, ranging from a conventional explicit delegation structure to various forms of modified server behaviours that attempt to guess the location of non- delegated servers for fragments of this address space. These approaches have been explored in some detail in terms of their advantages and drawbacks in **[3]**, so only a summary of these approaches will be provided here.

2.1 Conventional Address Delegation

The problem with this form of delegation is the anticipated piecemeal deployment of 6to4 sites. The reason why a site would use 6to4 is commonly that the upstream provider does not support a IPv6 transit service and the end site is using 6to4 to tunnel through to IPv6 connectivity. A conventional environment would have the 6to4 site using provider-based IPv4 addresses. In the IPv4 "in-addr.arpa" domain the local site would have an entry in the upstream's reverse DNS zone file, or would have authoritative local name servers that are delegated from the upstream's DNS zone. In the case of the mapped IPv6 space the upstream is not using IPv6 and therefore would not be expected to have a 6to4 delegation for its IPv4 address block.

Sub-delegations of IPv4 provider address space are not consistently recorded, and any 6to4 reverse zone operator would be required to undertake reverse zone delegations in the absence of reliable current address assignment information, undertaking a "hop over" of the upstream provider's address block. Similarly, a delegated entity may need to support the same "hop over" when undertaking further delegations in their reverse zone.

2.2 Guessing a Non-Delegated 6to4 Reverse Server

One way to avoid such unreliable delegations is to alter server behaviour for reverse servers in this zone. Where no explicit delegation information exists in the zone file the server could look up the in-addr.arpa domain for the servers for the equivalent IPv4 address root used in the 6to4 address. These servers could then be queried for the IPv6 PTR query.

The issues with fielding altered server behaviours for this domain are not to be taken lightly, and the delegation chain for IPv4 will not be the same for 6to4 in any case. An isolated 6to4 site uses a single IPv4 /32 address, and it is improbable that a single address would have explicit in-addr.arpa delegation. In other words it is not likely that the server delegation for IPv4 would parallel that of 6to4.

2.3 Locating Local Servers at Reserved Addresses

Another approach uses an altered server to resolve non-delegated 6to4 reverse queries. The 6to4 query is decoded to recover the original 6to4 IP address. The site-specific part of the address is rewritten to a constant value, and this value is used as the target of a lookup query. This requires that a 6to4 site should reserve local addresses, and configure reverse servers on these addresses. Again this is a weak approach in that getting the DNS to query non-delegated addresses is a case of generation of spurious traffic.

2.4 Synthesized Responses

The final approach is to synthesize an answer when no explicit delegation exists. This approach would construct a pseudo host name using the IPv6 query address as the seed. Given that the host name has no valid forward DNS mapping, then this becomes a case of transforming one invalid DNS object into another.

2.5 Selecting a Reasonable Approach

It would appear that the most reasonable approach is to support a model of conventional standard delegation. The consequent task is to reduce the administrative overheads in managing the zone, supporting delegation of reverse zone files on a basis of providing a delegation capability directly to each 6to4 site.

---

3. 6to4 Networks Address Use

A 6to4 client network is an isolated IPv6 network composed as a set of IPv6 hosts and a dual stack (IPv4 and IPv6) local router connected to the local IPv6 network and the external IPv4 network.

An example of a 6to4 network is as follows:

```
                         +-------------+
IPv6-in-IPv4 packets (A)|             |     IPv6 packets
------------------------| 6to4router  |--------------------------
                        |             |    |  |   |     |   |
                        +-------------+   local IPv6 clients

   IPv4 network                                IPv6 network
```

Figure 1

The IPv4 address used as part of the generation of 6to4 addresses for the local IPv6 network is the external IPv4 network (labelled '(A)' in the above diagram). For example, if the interface (A) has the IPv4 address 192.0.2.1, then the local IPv6 clients will use a common IPv6 address prefix of the form 2002:{192.0.2.1}::/48 (or (2002:C000:201::/48 in hex notation). All the local

IPv6 clients share this common /48 address prefix, irrespective of any local IPv4 address that such host may use if they are operating in a dual stack mode.

An example of a 6to4 network with addressing:

```
                   +-------------+
     IPv4 network  |             | IPv6 network
-------------------| 6to4router  |-------------
          192.0.2.1|             |    |  |    | interface identifier
                   +-------------+   1A  |    | local IPv6 address
                                     2002:C000:201::1A
                                         |    |
                                         1B   |
                                         2002:C000:201::1B
                                              |
                                              1C
                                              2002:C000:201::1C
```

Figure 2

---

4. Delegation Administration

This document proposes to use a a single delegation level in the 2.0.0.2.ip6.arpa zone, delegating zones only at the 48th bit position. The corresponds with individual delegations corresponding to a /32 IPv4 address, or the equivalent of a single 6to4 local site.

The zone files containing the end site delegations are proposed to be operated with a TTL (configured to be a time value in the scale of hours rather than days or weeks), and updates from delegation requests are to be made using incremental DNS updates **[2]**.

The delegation system is proposed to be self-driven by clients residing within 6to4 networks. The server's delegation function is proposed to be accessible only by clients using 6to4 IPv6 source addresses, and the only delegation that can be managed is that corresponding to the /48 prefix of the IPv6 source address of the client.

It is proposed to operate the delegation management service using secure web-based servers. This will ensure that the source address- driven delegation selection function cannot be disrupted through proxy caching of the server's responses.

It is proposed that the secure web servers be operated on a dual-stack IPv4 / IPv6 server. The service is to be available on a) an IPv4 address (instructions only), b) a native IPv6 address (instructions plus delegation service) and c) a 6to4 address (instructions plus delegation service).

The server's actions will be determined by the source address of the client. If the client uses a 6to4 source address the server will present a delegation interface for the corresponding 6to4 reverse zone. Otherwise the server will provide a description of the delegation process.

When accessed by a 6to4 source address, the interface presented by the delegation server is a standard DNS delegation interface, allowing the client to enter the details of a number of DNS servers for the corresponding reverse domain. The delegation servers are checked by the delegation manager to ensure that they are responding, that they are configured consistently and are authoritative for the delegated domain. If these conditions are met the delegation details are entered into the primary zone. In order to avoid the server being used as a denial of service platform the server should throttle the number of DNS requests made to any single IP address, and also throttle the number of redelegation requests for any single 6to4 zone.

In other cases the system provides diagnostic information to the client.

The benefits of this proposed structure include a fully automated mode of operation. The service delivery is on demand and the system only permits self-operation of the delegation function.

The potential issues with this structure include:

- Clients inside a 6to4 site could alter the delegation details without the knowledge of the site administrator. It is noted that this is intended for small-scale sites. Where there are potential issues of unauthorized access to this delegation function the local site administrator could take appropriate access control measures.
- IPv4 DHCP-based 6to4 sites could inherit nonsense reverse entries created by previous users of the DHCP address. In this case the client site could request delegation of the reverse zone as required.
- The approach does not scale efficiently, as there is the potential that the flat zone file may grow considerably. However it is noted that 6to4 is intended to be a transition mechanism useful for a limited period of time in a limited context of isolated network where other forms of tunnelled connection is not feasible. It is envisaged that at some point the density of IPv6 adoption in stub network would provide adequate drivers for widespread adoption of native IPv6 services, obviating the need for continued scaling of 6to4 support services. An estimate of the upper bound of the size of the 6to4 reverse delegation zone would be of the order of millions of entries. It is also noted that the value of a reverse delegation is a questionable proposition and many deployment environments have no form of reverse delegation.
- It is also conceivable that an enterprise network could decide to use 6to4 internally in some form of private context, with the hosts only visible in internal DNS servers. In this proposed mechanism the reverse delegation, if desired, would need to be implemented in an internal private (non-delegated) corresponding zone of the 6to4 reverse domain space.

It is envisaged that there may be circumstances with an IPv4 address controller wishes to "block" the ability for "children" to use this 6to4 scheme. It is envisaged that scenarios that would motivate this concern would include when the IPv4 provider is also offering an IPv6 service, and native IPv6 should be deployed instead of 6to4. In such circumstances the 2002 zone operator should allow for such a delegation blocking function upon application to the delegation zone operator.

For a delegation to be undertaken the following must hold:

- The 6to4 site must have connectivity to the global IPv6 network

- The 6to4 site must have configured a minimum of one primary and one secondary server for the 6to4 IPv6 reverse address zone.
- At the time of the delegation request, the primary and secondary servers should be online, reachable, correctly configured, and in a mutually consistent state with respect to the 6to4 reverse zone.
- The delegation server will only accept delegation requests associated with the 6to4 source address of the requesting client.

The approach suggested here, of a fully automated system driven by the site administrators of the 6to4 client networks, appears to represent an appropriate match the requirements of reverse DNS domains.

For maintenance of the reverse delegation zones it is proposed to maintain an email contact point for each active delegation, derived from the zone's SOA contact address, or explicitly entered in the delegation interface. This contact point would be informed upon any subsequent change of delegation details.

The management system will also undertake a periodic sweep of all active delegations, so that each delegation is checked every 30 days. If the delegation fails this integrity check the email contact point is informed of the problem, and a further check scheduled in a further 14 days. If this second check fails, the delegation is automatically removed, and a further notice is issued to the contact point.

---

5. Security Considerations

The system proposed here offers a moderate level of assurance in attempting to ensure that a 6to4 site can only direct the delegation of the corresponding reverse domain and no other.

Address-based authentication is not useful in a security sense. Accordingly, reverse delegation information does not provide useful information in either validating a domain name or in validating an IP address, and that no conclusions should be drawn from the presence or otherwise of a reverse mapping for any IP address.

The service management interface allows a 6to4 client to insert any server name as a DNS server, potentially directing the server to make a DNS query to any nominated system. The server should throttle the number of requests made to any single IP address to mitigate this risk of a high volume of bogus DNS queries being generated by the server. For similar reasons, the server should also throttle the number of redelegation requests for any single 6to4 zone.

---

6. Acknowledgements

The author acknowledges the prior work of Keith Moore in preparing a document that enumerated a number of possible approaches to undertake the delegation and discovery of reverse zones. The author acknowledges the assistance of

George Michaelson and Andrei Robachevsky in preparing this document, and Pekka Savola and Jun-ichiro itojun Hagino for their review comments.

7. References

[1] Carpenter, B. and K. Moore, "**Connection of IPv6 Domains via IPv4 Clouds**", RFC 3056, February 2001.

[2] Vixie, P., Thomson, S., Rekhter, Y. and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997 (HTML, XML).

[3] Moore, K., "Work in progress: 6to4 and DNS", April 2003.

Author's Address

Geoff Huston
APNIC