

Network Working Group
Request for Comments: 5158
Category: Informational

G. Huston
APNIC
March 2008

6to4 Reverse DNS Delegation Specification

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This memo describes the service mechanism for entering a delegation of DNS servers that provide reverse lookup of 6to4 IPv6 addresses into the 6to4 reverse zone file. The mechanism is based on a conventional DNS delegation service interface, allowing the service client to enter the details of a number of DNS servers for the delegated domain. In the context of a 6to4 reverse delegation, the client is primarily authenticated by its source address used in the delegation request, and is authorized to use the function if its IPv6 address prefix corresponds to an address from within the requested 6to4 delegation address block.

1. Introduction

6to4 [RFC3056] 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 an IPv6 network without necessarily obtaining IPv6 services from its adjacent upstream network provider.

As noted in [6to4-dns], 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 [DNS] 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." [6to4-dns]

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 behaviors 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 [6to4-dns], 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 an end site would use 6to4 is commonly that the upstream Internet Service Provider does not support an IPv6 transit service and the end site is using 6to4 to tunnel through to IPv6 connectivity. A conventional end site environment of this form would have the end site using provider-based IPv4 addresses, where the end site's IPv4 address is a more specific address block drawn from the upstream provider's address block, and the end site would have an entry in the upstream provider's reverse DNS zone file, or it would have authoritative local name servers that are delegated from the upstream provider's DNS zone. In the case of the 6to4 mapped IPv6 space, the upstream may not be providing any IPv6-based services at all, and therefore would not be expected to have a 6to4 reverse DNS delegation for its IPv4 address block. The general observation is that 6to4 IPv6 reverse DNS delegations cannot necessarily always precisely match the corresponding IPv4 reverse DNS delegations.

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 behavior 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 behaviors 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 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 considered here 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.

This service is implemented by web servers that are operated on a dual-stack IPv4 / IPv6 server, accessible via SSL. The web 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 service is a conventional DNS delegation interface, allowing the client to enter the details of a number of DNS servers for the corresponding reverse domain. The targets of the DNS delegation are checked by the delegation manager using IPv4 and IPv6, according to the addresses of the targets, 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 zone at the primary master. In order to avoid the server being used as a denial of service platform, the server should throttle the number of DNS delegation 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 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:

- o 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.
- o IPv4 DHCP-based 6to4 sites, or any 6to4 site that uses dynamically-assigned external IPv4 interface addresses, could inherit nonsense reverse entries created by previous users of the dynamically assigned address. In this case, the client site could request delegation of the reverse zone as required. More generally, given the potential for inheritance of 'stale' reverse DNS information in this context, in those cases where the issues of potential inheritance of 'stale' reverse DNS information is a concern, it is recommended that a 6to4 site either use a static IPv4 address in preference to a dynamically-assigned

address, or ensure that the reverse delegation information is updated using the service mechanism described here upon each dynamic address assignment event.

- o 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 an isolated network where other forms of a 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.
- o 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 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.

There may be circumstances where an IPv4 address controller wishes to "block" the ability for users of these addresses to use this 6to4 scheme. Scenarios that would motivate this concern would include situations when the IPv4 provider is also offering an IPv6 service, and native IPv6 should be deployed instead of 6to4. In such circumstances the 6to4 reverse zone operator should allow for such a 6to4 reverse delegation blocking function upon application to the delegation zone operator.

For a delegation to be undertaken the following conditions should hold:

- o The 6to4 site must have configured a minimum of one primary and one secondary server for the 6to4 IPv6 reverse address zone.
- o At the time of the delegation request, the primary and secondary servers must be online, reachable, correctly configured, and in a mutually consistent state with respect to the 6to4 reverse zone. In this context, "mutually consistent" implies the same SOA RR and identical NS RRsets.

- o The 6to4 reverse delegation service will only accept delegation requests associated with the 6to4 source address of the requesting client.

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

For maintenance of the reverse delegation zones the service maintains an email contact point for each active delegation, derived from the zone's SOA contact address (SOA RNAME), or explicitly entered in the delegation interface. This contact point would be informed upon any subsequent change of delegation details.

The 6to4 reverse DNS management system also undertakes 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 is scheduled for 14 days later. If this second check fails, the delegation is automatically removed, and a further notice is issued to the contact point.

5. Security Considerations

This system offers a rudimentary level of assurance in attempting to ensure that delegation requests from a 6to4 site can only direct the delegation of the corresponding 6to4 reverse DNS domain and no other.

Address-based authentication is not a very robust method from a security perspective, as addresses can be readily spoofed. Accordingly, reverse delegation information does not provide reliable information in either validating a domain name or in validating an IP address, and no conclusions should be drawn from the presence or otherwise of a reverse DNS 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 delegation test system to make a DNS query to any nominated system. The server throttles the number of requests made to any single IP address to mitigate the attendant risk of a high volume of bogus DNS queries being generated by the server. For similar reasons, the server also throttles the number of redelegation requests for any single 6to4 zone.

For a general threat analysis of 6to4, especially the additional risk of address spoofing in 2002::/16, see [RFC3964].

Section 4 notes that the local site administrator could take appropriate access control measures to prevent clients inside a 6to4 site performing unauthorized changes to the delegation details. This may be in the form of a firewall configuration, regarding control of access to the service from the interior of a 6to4 site, or a similar mechanism that enforces local access policies.

6. IANA Considerations

The IANA has delegated the 2.0.0.2.ip6.arpa domain according to delegation instructions provided by the Internet Architecture Board.

7. 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 Peter Koch, Pekka Savola, Jun-ichiro Itojun Hagino, and Catherine Meadows for their helpful review comments.

8. References

8.1. Normative References

- [RFC2136] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, October 2003.
- [RFC4346] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.1", RFC 4346, April 2006.

8.2. Informative References

- [6to4-dns] Moore, K., "6to4 and DNS", Work in Progress, April 2003.
- [RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", RFC 3964, December 2004.

Author's Address

Geoff Huston
APNIC

EMail: gih@apnic.net
URI: <http://www.apnic.net>

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

