

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

T. Lemon
Nominum
October 21, 2013

Populating the DNS Reverse Tree for DHCP Delegated Prefixes
draft-ietf-dhc-dns-pd-01.txt

Abstract

This document describes three alternatives for populating the DNS reverse tree for prefixes delegated using DHCP, and provides mechanisms for implementing each alternative.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 24, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Internet-Draft

DNS for Delegated Prefixes

October 2013

Table of Contents

1.	Introduction	2
2.	Terminology	2
3.	Methods for populating the reverse tree	3
3.1.	Site-managed reverse tree	3
3.2.	Provider-managed reverse tree	3
3.3.	Provider-managed spoofed reverse tree	3
3.4.	Other solutions not documented here	3
4.	Negotiating the reverse tree population method	4
5.	Configuring a site-managed reverse tree	6
5.1.	Requesting Router Behavior	6
5.2.	Delegating Router Behavior	8
6.	Configuring a provider-managed reverse tree	9
6.1.	Requesting Router Behavior	9
6.2.	Delegating Router Behavior	9
7.	Configuring a spoofed reverse tree	10
8.	Configuring no reverse tree	10
9.	Encoding of options	10
9.1.	Prefix Delegation Method Types	10
9.2.	Prefix Delegation Zone Preference Option	11
9.3.	Prefix Delegation Zone Method Option	11
9.4.	Prefix Delegation Zone Server Option	11
10.	Security Considerations	12
11.	IANA Considerations	12
12.	References	12
12.1.	Normative References	12
12.2.	Informative References	13
	Author's Address	13

[1.](#) Introduction

When a site is numbered using DHCP prefix delegation [[RFC3633](#)], there are three ways of populating the Domain Name System [[RFC1035](#)] reverse tree. Which mechanism is chosen depends on the capabilities of the site's DNS infrastructure, if any, on the capabilities and policies of the service provider, and on the preferences of the site administration.

This document does not take a position on which mechanism, if any, is best for populating the reverse tree, but simply documents each of the possible mechanisms for doing so, and provides a means whereby site administrators and service providers can negotiate the mechanism

whereby the reverse tree for a particular site will be populated.

[2.](#) Terminology

Lemon

Expires April 24, 2014

[Page 2]

Internet-Draft

DNS for Delegated Prefixes

October 2013

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

[3.](#) Methods for populating the reverse tree

There are three common methods of populating the reverse tree for a delegated prefix: delegation, dynamic dns, and zone spoofing. In addition, of course, it is possible to leave the reverse tree unpopulated.

[3.1.](#) Site-managed reverse tree

To populate the reverse tree by delegation, the site administrator must provide a DNS authoritative name server for the delegated zone. The site administrator must communicate the IP address of the authoritative name server to the service provider. The service provider must then add a delegation for that zone using the IP address or addresses of the DNS authoritative servers provided by the site administrator.

[3.2.](#) Provider-managed reverse tree

To populate the reverse tree using DNS updates, the service provider must provide an authoritative name server for the zone. The site administrator must provide a key to the service provider that can be used to authenticate DNS updates. The site administrator must then provide a mechanism whereby DNS updates will automatically be generated, using the provided key, whenever IP addresses are allocated within the delegated prefix.

[3.3.](#) Provider-managed spoofed reverse tree

In some cases the site administrator may not be willing or able to populate a reverse tree. However, the service provider may wish to provide meaningful answers to reverse zone queries for the delegated

zone. It's not possible to populate the delegated zone: a fully populated zone for a /64 would require 1.8×10^{19} names. However, the names in such a zone would never change; consequently it is possible for a name server to spoof the zone contents, constructing answers for queries against any name within the zone on the fly. Because the contents of the zone never change, the zone can have a consistent authority record.

[3.4.](#) Other solutions not documented here

It's worth noting that there are several other ways that the zone for a delegated prefix could be populated, but we are not covering these

mechanisms because they seem more difficult to implement and deploy. For instance, nodes configured with addresses within a delegated prefix could issue their own DNS updates to an authoritative server operated by the service provider. The problem of key management in this case becomes intractable, however.

It would also be possible for the site to have its own key management infrastructure, and for some agent on the requesting router to act as an intermediary in updating a zone maintained by the service provider. However, this is substantially more complicated than either of the proposed solutions.

Another option is to simply not populate the reverse tree. This is an attractive option in the IETF in particular because the reverse tree is frequently used for purposes to which it is not suited, and some IETF participants believe that in order to discourage these applications, it's better simply to not populate the reverse tree. This document takes no position on this question, but does offer a means whereby the site administrator can indicate that the reverse tree should not be populated.

[4.](#) Negotiating the reverse tree population method

The prefix delegation process is initiated by a requesting router. If a delegating router chooses to delegate a prefix to the requesting router, it replies with a prefix. The requesting router may receive responses from more than one delegating router, and may choose one or more such delegated prefixes. For delegating routers whose offer is accepted, the requesting router sends a request for the offered

address; at this point the delegating router commits the delegation to stable storage and sends a confirmation to the requesting router.

The messages used to complete this transaction are the DHCP Discover, DHCP Advertise, DHCP Request and DHCP Reply messages, respectively. The negotiation as to how the reverse tree will be populated piggybacks on this four-message process.

In the DHCP Discover message, the requesting router indicates the site administrator's preference for how the reverse tree for the delegated prefix will be populated. It does this by including, in each IA_PD option it sends, a Prefix Delegation Zone Preference option (PDZP) containing one or more preference codes. These codes are listed in order of preference with the most preferred mechanism first. A requesting router that includes a PDZP option MUST send an Option Request option (ORO) that requests the Prefix Delegation Zone Method (PDZM) option.

If the delegating router chooses not to delegate a prefix to the requesting router, no special action need be taken in response to the PDZP option. The remainder of this section describes what happens if the delegating router chooses to delegate a prefix to the requesting router.

Delegating routers that implement this specification can be configured with a list of supported reverse tree population methods. When a requesting router receives an IA_PD option that includes a PDZP option, if it has been configured with a reverse tree population method list, it iterates across the list of methods in the PDZP option. For each entry in the PDZP option, the requesting router tests to see if that method has been configured by the site administrator as being supported. If the method is on the list, the iteration stops at this point.

Upon completion of this iteration, if a method was found in the PDZP that is supported by the delegating router, that is the method that will be used to populate the reverse tree for the delegated zone. The delegating router constructs a PDZM option indicating that this method will be used and includes this in the DHCP Advertise message.

If no supported method was found, this means that the service provider will not cooperate with the site administrator in populating the reverse tree. The delegating router indicates that this is the case by not including a PDZM option in the DHCP Advertise message..

The requesting router may receive one or more DHCP Advertise messages containing delegated prefixes. The requesting router MUST silently discard any DHCP Advertise message containing a PDZM option that indicates a method that was not listed in the PDZP option sent in the DHCP Discover message.

The requesting router may then choose to respond to one or more of the remaining DHCP Advertise messages, if any. The lack of a PDZM option indicates either that the delegating router does not implement DNS for delegated prefixes, or that it is not configured to support DNS for delegated prefixes. The requesting router MAY prefer DHCP Advertise messages containing PDZM options over DHCP Advertise messages that do not contain PDZM options.

When responding to any DHCP Advertise messages containing PDZM options, the requesting router MUST include a PDZM option containing the same method indicated in the received PDZM option.

Each delegating router that receives a DHCP Request message containing a PDZM option MUST check the method indicated in the PDZM option is supported; if not, the delegating router MUST silently discard the DHCP Request option.

The requesting and delegating routers should follow the same procedure specified for the DHCP Request/DHCP Reply sequence whenever a DHCP Renew or DHCP Rebind is sent and a DHCP reply sent in response, if that response renews the delegated prefix. In the case that the response does not renew the prefix, the delegating router MUST NOT send a PDZM in the IA_PD option.

[5.](#) Configuring a site-managed reverse tree

If the PDZM option returned by the delegating router in the DHCP

Advertise message specifies the Site Managed method, the requesting router must arrange to set up one or more authoritative name servers that will provide service for the zone or zones that correspond to the delegated prefix. It must also communicate to the delegating router the IP address or addresses of these servers.

5.1. Requesting Router Behavior

The requesting router MUST include a Prefix Delegation Zone Server (PDZS) option in each IA_PD in the DHCP Request message, which includes zero or more IP addresses of authoritative name servers for the delegated zone. IPv4 addresses MUST be represented as IPv4-Embedded IPv6 addresses using the Well-Known prefix [[RFC6052](#)].

Authoritative name service for these zones may be provided by any or all of the following three types of authoritative name servers:

- o An authoritative name server running on a node that has an IP address known to the requesting router that is not obtained from the prefix being delegated.
- o An authoritative name server running on the requesting router.
- o An authoritative name server running on a node that will obtain its only IP address from the prefix being delegated.

In the first case, it is possible that the reverse zone for the delegated prefix is already configured on the authoritative name server. In this case, the requesting router SHOULD include the IP address of the authoritative name servers for the delegated zone in the PDZS option.

However, if the prefix is being delegated for the first time, the delegating router will not have had an opportunity to configure it prior to sending the DHCP Request message. In this case, the delegating router SHOULD NOT include the IP Address of this name server in the PDZS option that's send in the DHCP Request message; instead, it should send a DHCP Renew once the authoritative server has been configured, and list the server's IP address in the PDZS option in the DHCP Renew message.

In the second case, the requesting router may already have an IP address, and may be able to configure the authoritative server for the delegated zone before sending the DHCP Request. In this case, the requesting router SHOULD include its own IP address in the PDZS option in the DHCP Request message.

If the requesting router does not have an IP address at this time, it SHOULD send a DHCP Renew message containing a PDZS option that lists all the authoritative servers for the reverse zone or zones for the delegated prefix after it has an IP address and has configured the authoritative servers.

If the authoritative name server is running on a node that will configure its IP address from the delegated prefix, this name server cannot even be configured until it has an IP address. The process of configuring this name server is beyond the scope of the document; however, once the name server has been configured, the requesting router SHOULD send a DHCP Renew message for the delegated prefix with an IA_PD containing a PDZS option that lists the IP address of this name server.

In general, if there are any globally-reachable name servers that are authoritative for the zone or zones that provide the reverse tree for the delegated prefix at the time that the DHCP Request message is sent, the requesting router should list the IP addresses of these name servers in the PDZS option in the associated IA_PD option in the DHCP Request message.

If new globally-reachable name servers that are authoritative for the reverse zone or zones become available after the DHCP Request has been sent and the DHCP Reply received, the requesting router SHOULD send a DHCP Renew message containing an IA_PD for the delegated prefix and a PDZS option listing the name servers for that prefix that have come online. The requesting router SHOULD be aware of all outstanding name server configuration processes and minimize the number of DHCP Renew message sent.

When a requesting router sends a DHCP Renew or DHCP Rebind message to renew a delegated prefix, if a site-managed reverse tree was

successfully configured, the requesting router MUST send a PDZM

option containing the same method sent in the original DHCP Request message. The requesting router MUST also send a PDZS option that contains one or more IP addresses for authoritative servers for the reverse tree for the delegated prefix.

5.2. Delegating Router Behavior

When a delegating router receives a valid DHCP Request message containing an IA_PD that contains both a PDZM option indicating the Site Managed method and a PDZS option containing at least one IP address, it compares the IP addresses in the PDZM option to any previous record it may have for that delegation. If the contents of the PDZM option differ from the previous record, or if there is no previous record, the delegating router MUST issue a DNS Update to add a delegation to the parent zone of the reverse tree zone for the delegated prefix.

In the event that the PDZS option contains zero IP addresses, the delegating router does not update the zone.

If the delegated prefix must be represented as more than one zone, the delegating router adds delegations to the parent zone for each such zone.

When a delegating router receives a DHCP Renew or DHCP Rebind message for a prefix it delegated and elects to renew the prefix, it MUST check its record for that prefix to see if a delegation exists. If the contents of the PDZS differ from the recorded list of authoritative name servers for that prefix, the delegating router MUST update the parent zone with the new delegations.

When a delegating router receives a DHCP Renew or DHCP Rebind message for a prefix it delegated, and elects not to renew the delegation, the delegating router MUST check to see if it has a site-managed reverse tree configuration for that pprefix. If it does, it must update the parent zone to remove any delegations that were added, and update its record for the delegated prefix to indicate that no site-managed reverse tree configuration for that prefix is present.

When a delegated prefix expires without being renewed by the requesting router, the same procedure should be followed to update the parent zone.

In all cases where the delegating router updates the delegation for the zone, it must first query the name server or servers listed in the PDZS option for an SOA record for each delegated zone. If the name server does not respond within the standard timeout period, or

does not provide an authoritative answer, the delegating router MUST NOT add a delegation for that name server.

[6.](#) Configuring a provider-managed reverse tree

If the PDZM returned by a delegating router in the DHCP Advertise message specifies the Provider Managed method, the delegating router must arrange to set up a reverse zone for the delegated prefix. The requesting router must communicate a key to the delegating router that can be used to secure updates to the reverse zone.

[6.1.](#) Requesting Router Behavior

In order to update the provider-managed reverse zone, the requesting router must provide a key to the delegating router. Because DHCP does not provide confidentiality, this key must be the public half of a private key.

Typically sites that wish to populate their reverse tree with meaningful information maintain a site-specific or company-wide DNS zone. In order to update the reverse zone, the site administrator must publish a SIG(0) key in this zone. The requesting router MUST include a Prefix Delegation SIG(0) Key FQDN (PDSKF) option in the DHCP Request message and any subsequent DHCP Renew messages. It must use the private half of the SIG(0) key in any DNS updates to the reverse zone.

[6.2.](#) Delegating Router Behavior

There are two cases that the delegating router needs to handle: the case where the prefix being delegated was previously delegated to the same requesting router, and the case where it was not.

In the case where the prefix was previously delegated to the same requesting router, the delegating router need take no action to populate the zone, because it should already be populated.

In the case where the prefix was previously delegated to a different requesting router, the delegating router MUST remove the old zone information from the master authoritative name server for the zone.

In this case, and in the case where no previous delegation had been done, the delegating router must then configure a new reverse zone on the master server.

In any case, the delegating router must configure the reverse zone so

that it can be updated using the SIG(0) key stored on the name provided by the requesting router in the PDSKF option.

[7.](#) Configuring a spoofed reverse tree

A spoofed reverse tree can be configured either unilaterally by the service provider or upon request of the site administrator. The site administrator would list this as an option to indicate a preference for a spoofed reverse tree over no reverse tree; the choice doesn't make any sense otherwise.

Generally speaking, the service provider has the option of either setting up spoofed zones on demand, or setting them up when requested. If the service provider only offers spoofed zones, it makes some sense to set them up in advance; otherwise they should be set up whenever a prefix is delegated to a particular requesting router for the first time.

In some cases the site administration may request a spoofed zone because they do not wish to populate the reverse tree, but wish for it to appear populated. A service provider may support this option in addition to the site-managed option, the provider-managed option, and the no zone option. In this case, when a prefix is delegated to a new router for the first time, there may be an old zone configured differently. In this case, the delegated router **MUST** remove the old zone configuration before setting up the spoofed zone.

[8.](#) Configuring no reverse tree

A service provider may choose to simply not populate reverse trees for delegated prefixes. This is a desirable option in the sense that it minimizes the work required to support the reverse DNS tree, and avoids creating spoofed nonsense records. The service provider may also simply offer it as an option for sites that prefer not to have a populated reverse tree.

In this case, if the non-populated reverse tree is an option, and the prefix had previously been delegated to a different router, the delegating router must remove any previously-existing zone for the delegated prefix.

[9.](#) Encoding of options

easier ways to get a DDoS amplification.

11. IANA Considerations

We request that IANA assign three new option codes from the DHCP Option Codes table of the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) parameters registry maintained in <http://www.iana.org/assignments/dhcpv6-parameters/dhcpv6-parameters.xml>. These option codes will be assigned to the Prefix Delegation Zone Preference (OPT_PDZP), Prefix Delegation Zone Method (OPT_PDZM) and Prefix Delegation Zone Servers (OPT_PDZS) options.

We also request that the IANA add a new table, the Prefix Delegation Zone Method Types table, to the same registry. The first three entries in the table will contain the values specified in the section above titled "Prefix Delegation Zone Method Types." New entries to the table may be added according to the "Specification Required" IANA policy [[RFC5226](#)].

12. References

12.1. Normative References

Lemon Expires April 24, 2014 [Page 12]

Internet-Draft DNS for Delegated Prefixes October 2013

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), December 2003.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), October 2010.

12.2. Informative References

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#),

May 2008.

Author's Address

Ted Lemon
Nominum
2000 Seaport Blvd
Redwood City, CA 94063
USA

Phone: +1 650 381 6000
Email: mellon@nominum.com