

MIF Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 17, 2016

M. Stenberg
S. Barth
Independent
October 15, 2015

**Multiple Provisioning Domains using Domain Name System
draft-stenberg-mif-mpvd-dns-00**

Abstract

This document describes a mechanism to transmit and secure provisioning domain information for IPv6 and IPv4 addresses by using reverse DNS resolution. In addition it specifies backwards-compatible extensions to IPv6 host configuration to support special-purpose global IPv6 prefixes which can only be used to access certain isolated services.

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 17, 2016.

Copyright Notice

Copyright (c) 2015 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.

Table of Contents

- [1. Introduction](#) [2](#)
- [2. Terminology](#) [3](#)
 - [2.1. Requirements Language](#) [3](#)
- [3. PVD information discovery using DNS](#) [3](#)
 - [3.1. PVD TXT Record Fomat](#) [4](#)
 - [3.2. PVD TXT Record Keys](#) [4](#)
 - [3.2.1. Reachable Services](#) [4](#)
 - [3.2.2. DNS Configuration](#) [5](#)
 - [3.2.3. Connectivity Characteristics](#) [5](#)
 - [3.2.4. Private Extensions](#) [6](#)
- [4. Special-purpose IPv6 prefixes](#) [6](#)
 - [4.1. Extensions to Stateless Address Autoconfiguration](#) [6](#)
 - [4.2. Extensions to DHCPv6](#) [7](#)
- [5. Security Considerations](#) [9](#)
- [6. IANA Considerations](#) [9](#)
- [7. References](#) [10](#)
 - [7.1. Normative references](#) [10](#)
 - [7.2. Informative references](#) [11](#)
- [Appendix A. This solution compared to doing this in DHCPv6/NDP](#)
[RFC Editor: please remove] [11](#)
- [Appendix B. Discussion Points](#) [RFC Editor: please remove] [12](#)
- [Appendix C. Changelog](#) [RFC Editor: please remove] [13](#)
- [Appendix D. Draft Source](#) [RFC Editor: please remove] [13](#)
- [Appendix E. Acknowledgements](#) [13](#)
- Authors' Addresses [13](#)

1. Introduction

Given multiple address prefixes or multiple interfaces, hosts require more information to make educated choices about the interfaces and addresses to use. [RFC7556] describes the provision domains (PVDs) that provide the additional information the hosts need.

This document describes where and how the provision domain information is encoded in the Domain Name System (DNS). For optional authentication DNSSEC is used.

A backwards compatible way of adding IPv6 prefixes without generic internet connectivity is also provided so that the hosts that are not aware of the provisioning domain prefixes do not inadvertently use those for general network access.

2. Terminology

PVD	a provisioning domain, usually with a set of provisioning domain information; for more information, see [RFC7556].
special-purpose IPv6 prefix	a global IPv6 source prefix that cannot be used to reach the public IPv6 internet but instead only allows access to certain special services (e.g., VoIP, IPTV).

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

3. PVD information discovery using DNS

Each PVD is stored within the DNS, encoded as a single TXT record [[RFC1035](#)]. [Section 3.1](#) describes the syntax and [Section 3.2](#) the semantics of the enclosed data.

To find the per-PVD TXT records that apply to a source address, the host queries the DNS for PTR records of the domain `_pvd.<domain>`. `<domain>` is a .arpa domain for reverse lookups derived from the respective prefix or subnet the source address is assigned from and generated as specified in [[RFC3596](#)] for IPv6 addresses and [[RFC1034](#)] for IPv4 addresses respectively. If the query returned any PTR records the host then subsequently queries the DNS for TXT records located in each domain indicated in the PTR records and handles their contents as individual PVDs.

As prefixes can be sub-delegated arbitrarily, PTR records SHOULD be provided for any subprefixes contained within a particular prefix. For example, given a prefix `2001:db8:8765:4321::/64`, a host with an address of `2001:db8:8765:4321:1234:5678:abcd:beef` queries for PTR records in `_pvd.0.1.2.3.4.5.6.7.8.8.b.d.0.1.0.0.2.ip6.arpa`. However, if the address is assigned from `2001:db8:8765:4321:1234::/80`, the request would be made for `_pvd.0.4.3.2.1.1.2.3.4.5.6.7.8.8.b.d.0.1.0.0.2.ip6.arpa` instead.

If for example, the retrieved PTR record is assumed to point at the FQDN `_pvd.example.com`. The next query is then sent for TXT records in this domain, and if successful, the node has retrieved up-to-date information about PVDs applicable to that particular address.

A host MUST support handling multiple PTR records for the initial .arpa domain as well as multiple TXT records for all domains pointed to by the PTR records. This facilitates handling of multiple PVDs with minimal amount of state in the network. A host MUST honor both the time-to-live of the received records, and negative replies that conform to [RFC2308]. A host MUST NOT use addresses from a prefix as the source for new packet flows once the TTL has passed until it did successfully retrieve updated PVD information.

3.1. PVD TXT Record Fomat

PVD information within DNS is encoded using TXT records, similar to those of DNS-SD [RFC6763] and defined as follows. TXT records consist of key/value pairs, each encoded as a string of up to 255 octets preceded by a length byte storing the number of octets. The strings are in the form "key=value" or simply "key" (without quotation marks) where everything up to the first '=' character (if any, otherwise the whole string) is the key and everything after it (if any, including subsequent '=' characters) is the value. Due to the use of a length byte, quotation marks or similar are not required around the value. Keys are case-sensitive. Hosts MUST ignore TXT records which do not conform to these rules.

3.2. PVD TXT Record Keys

The following keys are defined to be used inside PVD TXT records. Unknown keys inside PVD Information MUST be ignored.

3.2.1. Reachable Services

The following set of keys can be used to specify the set of services for which the respective PVD should be used. If present they MUST be honored by the client, i.e., if the PVD is marked as not usable for internet access it MUST NOT be used for internet access, if the usability is limited to a certain set of domain or address prefixes, then a different PVD MUST be used for other destinations.

Key	Description	Value	Example
n	User-visible service name	human-readable UTF-8 string	n=FooBar Service
s	Internet inaccessible	(none)	s
z	DNS zones accessible	comma-separated list of DNS zone	z=foo.com,sub.bar.com
6	IPv6-prefixes accessible	comma-separated list of IPv6 prefixes	6=2001:db8:a::/48,2001:db8:b:c::/64
4	IPv4-prefixes accessible	comma-separated list of IPv4 prefixes in CIDR	4=1.2.3.0/24,2.3.0.0/16

3.2.2. DNS Configuration

The following set of keys can be used to specify the DNS configuration for the respective PVD. If present, they MUST be honored and used by the client whenever it wishes to access a resource of the PVD.

Key	Description	Value	Example
r	Recursive DNS server	comma-separated list of IPv6 and IPv4 addresses	r=2001:db8::1,192.0.2.2
d	DNS search domains	comma-separated list of search domains	d=foo.com,sub.bar.com

3.2.3. Connectivity Characteristics

The following set of keys can be used to signal certain characteristics of the connection towards the PVD.

Key	Description	Value	Example
bw	Maximum achievable bandwidth	1 symmetrical or 2 comma-separated ingress, egress values in kilobits per second	bw=5000 or bw=1000,100
lt	Minimum achievable latency	1 symmetrical or 2 comma-separated ingress, egress values in milliseconds	lt=20 ot lt=20,100
rl	Maximum achievable reliability	1 symmetrical or 2 comma-separated ingress, egress values in 1/1000	rl=1000 or rl=900,800
tm	Traffic metered (cut-off / limited over threshold)	(none) or traffic threshold in kilobytes	tm or tm=1000000
cp	Captive portal	(none)	cp
nat	IPv4 NAT in place	(none)	nat

3.2.4. Private Extensions

keys starting with "x-" are reserved for private use and can be utilized to provide vendor-, user- or enterprise-specific information. It is RECOMMENDED to use one of the patterns "x-FQDN-KEY[=VALUE]" or "x-PEN-KEY[=VALUE]" where FQDN is a fully qualified domain name or PEN is a private enterprise number [PEN] under control of the author of the extension to avoid collisions.

4. Special-purpose IPv6 prefixes

A service provider might wish to assign certain global unicast address prefixes which can be used to reach a limited set of services only. In the presence of legacy hosts it must be ensured however that these prefixes are not mistakenly used as source addresses for other destinations. This section therefore defines optional extensions to NDP [RFC4861], DHCPv6 [RFC3315] and DHCPv6-PD [RFC3633] to indicate this state.

4.1. Extensions to Stateless Address Autoconfiguration

NDP [RFC4861] defines the Prefix Information option to announce prefixes for stateless address configuration. The "A-bit" is set, whenever hosts may autonomously derive addresses from a given prefix. For special-purpose prefixes this document defines the first bit of the Reserved1-field as the "S-Bit".

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type										Length										Prefix Length										L A S Reserved1									
...																																							

The following additional requirements apply to hosts intending to support global special-purpose IPv6 prefixes:

Upon reception of a Prefix Information Option with the S-Bit set, it should behave as if the A-Bit was set, however (unless the A-Bit was actually set by the sending router) it MUST delay using any addresses derived from the prefix until it has queried, retrieved and honored (see [Section 3](#)) at least all mandatory provisioning domain information related to the given prefix.

A host SHOULD NOT interpret the S-Bit being clear as an indicator that no provisioning domain information is available for a given prefix.

The following additional requirements apply to routers:

A router MUST NOT set the A-Bit for global unicast address prefixes which cannot be used to reach the public IPv6 internet.


A router SHOULD use the S-Bit to indicate that PVD-aware hosts can statelessly assign themselves addresses from a given prefix. It MAY use the S-Bit in addition to the A-Bit to indicate that a prefix usable to reach the public IPv6 internet has additional (optional) provisioning domain information.

A router announcing one or more Prefix Information options with the S-Bit set MUST also announce one or more recursive DNS servers using a Recursive DNS Server Option [[RFC6106](#)]. If none of the Prefix Information Options it announces have the A-Bit set then at least one of these recursive DNS servers MUST be reachable using a link-local address.

4.2. Extensions to DHCPv6

Using DHCPv6 [[RFC3315](#)] and DHCPv6-PD [[RFC3633](#)] servers can actively decide which addresses or prefixes are assigned to clients and requesting routers, however a mechanism is needed to distinguish PVD-aware devices and in the same sense PVD-aware devices need to be able to detect which prefixes and addresses are special-purpose. Therefore, a zero-length DHCPv6 option OPTION_SPECIAL_PURPOSE is

defined to be added as a suboption to OPTION_IAADDR and OPTION_IAPREFIX options.



The following additional requirements apply to clients and requesting routers intending to support global special-purpose IPv6 prefixes via DHCPv6:

A client or requesting router MUST include the option code OPTION_SPECIAL_PURPOSE in an option OPTION_ORO in its SOLICIT and REQUEST messages, whenever it wishes to accept special-purpose prefixes in its identity associations.

Upon reception of an OPTION_IAADDR or OPTION_IAPREFIX option having an embedded OPTION_SPECIAL_PURPOSE option it MUST delay using any addresses derived from the prefix as source address for its own communication until it has queried, retrieved and honored (see Section 3) at least all mandatory provisioning domain information related to the given prefix or address. If it is a requesting router, it MAY however subdelegate prefixes or assign addresses from special-purpose prefixes to clients without doing so as long as the requirements in the following paragraph are honored.

The following additional requirements apply to routers assigning addresses from or delegating (parts of) special-purpose prefixes using DHCPv6:

A router MUST include a zero-length suboption of type OPTION_SPECIAL_PURPOSE in every OPTION_IAADDR and OPTION_IAPREFIX option it assigns or delegates containing a global unicast address or prefix which cannot be used to reach the public IPv6 internet. It MUST NOT assign or delegate such an address or prefix to a client or requesting router not including the option code of OPTION_SPECIAL_PURPOSE inside an OPTION_ORO option.

A router announcing one or more addresses or prefixes with an embedded OPTION_SPECIAL_PURPOSE option MUST also announce one or more recursive DNS servers using a OPTION_DNS_SERVERS option [RFC3646]. If all of the addresses in a DHCPv6 reply carry the embedded OPTION_SPECIAL_PURPOSE option then at least one of the announced recursive DNS servers MUST be reachable using a link-local address.

5. Security Considerations

The security implications of using PVDs in general depend on two factors: what they are allowed to do, and whether or not the authentication and authorization of the PVD information received is sufficient for the particular usecase. As the defined scheme uses DNS for retrieval of the connection parameters, the retrieval of both the PTR and the TXT records should be secured, if they are to be trusted. The PVD information allows for the following types of attacks:

- o Traffic redirection, both by providing custom DNS server, as well as actual potentially different next-hop and/or source address selection.
- o Faking of connection capabilities to e.g. prefer some connection fraudulently over others.

If a host requires DNSSEC authentication and the retrieved information is not sufficiently secured, they MUST be ignored as the defined way of using them in [Section 3.2](#) requires honoring the supplied information.

Security properties of NDP and DHCPv6 are inherited for the respective extensions, therefore relevant sections of [[RFC4861](#)] and [[RFC3315](#)] should be consulted. In any case, signaling addresses and prefixes to be special-purpose does not have a significant impact on the underlying assignment and delegation mechanisms.

6. IANA Considerations

IANA is requested to setup a PVD DNS TXT Record Key registry with the initial types: s, z, 6, 4 ([Section 3.2.1](#)); r, d ([Section 3.2.2](#)); bw, lt, rl, tm, cp, nat ([Section 3.2.3](#)) and a prefix x- ([Section 3.2.4](#)) for Private Use [[RFC5226](#)]. The policy for further additions to the registry is requested to be RFC Required [[RFC5226](#)].

This document defines a new bit for the NDP Prefix Information Option (the S-bit). There is currently no registration procedure for such bits, so IANA should not take any action on this matter.

IANA should assign a DHCPv6 option code OPTION_SPECIAL_PURPOSE to the DHCPv6 option code space defined in [[RFC3315](#)].

7. References

7.1. Normative references

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), DOI 10.17487/RFC1034, November 1987, <<http://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), DOI 10.17487/RFC1035, November 1987, <<http://www.rfc-editor.org/info/rfc1035>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/[RFC2119](#), March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", [RFC 2308](#), DOI 10.17487/RFC2308, March 1998, <<http://www.rfc-editor.org/info/rfc2308>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3646](#), DOI 10.17487/RFC3646, December 2003, <<http://www.rfc-editor.org/info/rfc3646>>.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", [RFC 6106](#), DOI 10.17487/RFC6106, November 2010, <<http://www.rfc-editor.org/info/rfc6106>>.

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", [RFC 7556](#), DOI 10.17487/RFC7556, June 2015, <<http://www.rfc-editor.org/info/rfc7556>>.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", [RFC 3596](#), DOI 10.17487/RFC3596, October 2003, <<http://www.rfc-editor.org/info/rfc3596>>.

7.2. Informative references

- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", [RFC 6763](#), DOI 10.17487/RFC6763, February 2013, <<http://www.rfc-editor.org/info/rfc6763>>.
- [PEN] IANA, "Private Enterprise Numbers", <<https://www.iana.org/assignments/enterprise-numbers>>.

Appendix A. This solution compared to doing this in DHCPv6/NDP [RFC Editor: please remove]

The angle of attack of the MIF work to date (autumn 2015) has been to add container options and their transfer mechanisms to DHCPv6 and NDP. This document details a different approach, and therefore it is sensible to compare it to the existing solutions in terms of (highly subjective) pros and cons. The authors consider pros of this proposal to be:

- o No overhead for hosts that do not care (possibly most; no spurious RA options, ...)
- o No problems with relaying data; if the first-hop device does not care, DNS requests propagate onward.
- o Little/no changes to DHCP, DHCPv6, DHCPv6-PD or RA.
- o Much more scalable; no worries about multicast packet size limits.
- o No duplication of specifications / TLVs for DHCP, DHCPv6 and RA.
- o Solves m:n prefix <-> PVD elegantly: no need to either duplicate applying prefix for each PVD or duplicate each PVD for each applying prefix.

- o Easily extensible (TXT records, no TLV definitions, parsing and generation necessary)
- o Probably not affected by IPR on [draft-ietf-mif-mpvd-dhcp-support](#)
- o Reuses the existing reverse DNS infrastructure

The authors consider cons of this proposal to be:

- o This scheme requires DNS servers 'close' on the path to the user, if changed information is to be sent; otherwise centralized solution would work (with some synthesized records).
- o Security using either DNSSEC or in-band hashes is rather painful (but possibly not more than the scheme in the current DHCP/RA drafts), so the default would most likely be insecure. That is not much different from DHCP*/RA, which are also 99.999...% of the time not secured.

Appendix B. Discussion Points [RFC Editor: please remove]

- o Besides special purpose prefixes, it might be desirable to have special purpose routers which only provide access to certain services but not the entire internet. These services could be announced by only using more-specific routes, however if the destination addresses are possibly changing, extension of the RIO mechanism might be needed. One possibility would be to add a new RIO S-flag with semantics like: "When the host receives a Route Information Option with the S-Bit set, it MUST ignore the value in the Prf-field (even if it is (10)) and instead assume the preference to have a value greater than (11). However, it MUST only use the route for packets having a source prefix announced by the same router.". This would allow selective routes (Prf=(10)) only applying to MIF-hosts.
- o DNSSEC delegation of reverse zones might be difficult in some cases. It is debatable, whether we want a complementary in-band signing mechanism as well, e.g., pre-shared public keys associated the domain name of the TXT records and "sig-X=..." keys (where X identifies the specific key) and ... is an EdDSA or ECDSA signature over all records not starting with "sig-". Care would need to be taken wrt. TTL and negative caching though.
- o Should PVD-aware hosts be recommended or even required to always prefer routers that announced the respective source address in a PIO over those that didn't when making routing decisions? This takes on the points made in [draft-baker-6man-multi-homed-host](#).

Appendix C. Changelog [RFC Editor: please remove]

[draft-stenberg-mif-mpvd-dns-00](#):

- o Initial version.

Appendix D. Draft Source [RFC Editor: please remove]

As usual, this draft is available at <https://github.com/fingon/ietf-drafts/> in source format (with nice Makefile too). Feel free to send comments and/or pull requests if and when you have changes to it!

Appendix E. Acknowledgements

Thanks to Eric Vyncke for the original idea of using DNS for transmitting PVD information.

Authors' Addresses

Markus Stenberg
Independent
Helsinki 00930
Finland

Email: markus.stenberg@iki.fi

Steven Barth
Independent
Halle 06114
Germany

Email: cyrus@openwrt.org

