

Distributed Mobility Management (DMM)
Internet-Draft
Updates: [4862](#) (if approved)
Intended status: Standards Track
Expires: August 28, 2016

J. Korhonen
Broadcom
S. Gundavelli
Cisco
P. Seite
Orange
D. Liu
Alibaba
February 25, 2016

IPv6 Prefix Properties
draft-korhonen-dmm-prefix-properties-05.txt

Abstract

This specification defines an extension to the IPv6 stateless address autoconfiguration procedure. New options with meta data are defined that describe the properties and other prefix class meta data associated with the prefix. The stateless address autoconfiguration procedure and end hosts can make use of the additional properties and class information when selecting source address prefixes for a particular uses and use cases. This specification updates [RFC4862](#).

Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

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 August 28, 2016.

Copyright Notice

Copyright (c) 2016 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.](#) Background and Motivation [3](#)
- [3.](#) Option Formats [4](#)
 - [3.1.](#) Prefix Meta Data [5](#)
 - [3.2.](#) Meta Data Suboptions [6](#)
- [4.](#) Host Considerations [7](#)
 - [4.1.](#) Stateless Address Autoconfiguration Enhancements [7](#)
 - [4.2.](#) Internal Data Structures [8](#)
 - [4.3.](#) Default Address Selection [8](#)
- [5.](#) Router Considerations [8](#)
- [6.](#) Multiple Provisioning Domain Considerations [9](#)
- [7.](#) Security Considerations [9](#)
- [8.](#) IANA Considerations [10](#)
- [9.](#) Acknowledgements [10](#)
- [10.](#) References [10](#)
 - [10.1.](#) Normative References [10](#)
 - [10.2.](#) Informational References [11](#)
- Authors' Addresses [12](#)

1. Introduction

This specification defines a new neighbor discovery protocol message option, the Prefix Information Option with Meta Data (PIOMD), that indicate, for example, the mobility management properties associated to the prefix, and a class value that conveys metadata associated to the prefix with a local administrative domain wide importance. The solution may use of Multiple Provisioning Domains (MPVD) framework [[RFC7556](#)] [[I-D.ietf-mif-mpvd-ndp-support](#)]. Furthermore, the specification discusses corresponding source address selection hint issues with the IPv6 Socket API and applications in general [[I-D.ietf-dmm-ondemand-mobility](#)].

For example, the IPv6 Socket API for Source Address Selection [[RFC5014](#)] already covers Mobile IPv6 [[RFC6275](#)] and allows selecting between a home address (HoA) and a care-of address (CoA). A mobile node (MN) with a client based mobility IP stack is supposed to know which prefixes are CoA(s) and/or HoA(s). However, this is not the case with network based mobility management where the MN is expected to be agnostic of the mobility support.

The extensions are minimal in a sense that they do not define new functionality, for example, to any existing mobility protocol but instead add an explicit indication of network based mobility knowledge into the IPv6 stateless address autoconfiguration (SLAAC) [[RFC4862](#)]. The heavy lifting is mostly on the applications side and on the IP stack providing interface for applications, since they need to make use of the new functionality. The new functionality is achieved by defining a new, backward compatible, IPv6 neighbor discovery protocol options that convey the required prefix related meta data information the SLAAC procedure may take use of.

This would allow for network based mobility solutions, such as Proxy Mobile IPv6 [[RFC5213](#)] or GTP [[TS.29274](#)] to explicitly indicate that their prefixes have mobility, and therefore, the MN IP stack or specifically applications can make an educated selection between prefixes that have mobility and those that do not. There is also a potential need to extend both [[RFC3493](#)] and [[RFC5014](#)] in order to provide required hooks into socket APIs.

The underlying assumption is that a MN has multiple prefixes to choose from. Typically this means either the MN has multiple interfaces or an interface has been configured with multiple prefixes. This specification does not make a distinction between these alternatives and does not either make any assumptions how the possible transfer of a prefix is done between interfaces in the case a network based mobility solution is used.

2. Background and Motivation

This section discusses the motivations behind adding metadata and other address selection decision making affecting information into IPv6 prefixes. The additional information is conveyed from the network to a end host during the IPv6 address configuration phase. The motivation example taken from and discussed below is from the mobile networks.

IP mobility and its centralized topological anchoring of IP addresses has known issues. For instance, non-optimal routing is a classical example. Another concerns include excessive tunneling, increased signaling due the maintenance of mobility related bindings,

aggregation of traffic to centralized mobility anchor gateways and unnecessary IP mobility related state management for IP traffic that does not as such benefit from mobility. In general, it is observed that most applications do not need IP level mobility, and work just fine with "temporary" IP addresses that come and go. However, IP mobility still has its virtues making the applications unaware of mobility, and certain wireless mobile networking architecture make extensive use of network based IP mobility.

In order to overcome some of the above issues, use of local resources and topologically local addressing could be enhanced. In many cases this would lead to use of multiple addresses of which some provide mobility and some do not. However, an end host has to have means to distinguish between addresses that provide mobility, and those that are short lived and usable only within a limited topological area.

[RFC7333] discussed the requirements for distributed mobility management and [RFC7429] describes the gaps from current best practices and the desired approaches for de-centralized mobility management. One approach is using the dynamic anchoring for distributed de-centralized mobility management. The idea is to use the local allocated prefix for any newly initiated 'IP session' and use the previously allocated prefix for the ongoing sessions. This specification can be used to implement the prefix selection for dynamic anchoring. For example, both the locally allocated and the remotely allocated/anchored prefixes can be identified by the prefix property option as described in [Section 3.2](#).

The solution described in this document also shares similar motivations for classifying the prefix as described in [[I-D.bhandari-dhc-class-based-prefix](#)]. Some service providers may wish to allocate specific prefixes for some services or type of traffic. In this situation, the end host must be able to classify prefixes according to type of service.

This specification provides tools for extending the IPv6 address management and source address selection so that end hosts (and their applications) can select a proper address for their needs. This specification complements [[I-D.bhandari-dhc-class-based-prefix](#)] by providing the SLAAC version of the additional prefix related meta data information delivery compared to the DHCPv6 stateful approach.

3. Option Formats

3.1. Prefix Meta Data

This specification defines a new neighbor discovery protocol message option, the Prefix Information Option with Meta Data (PIOMD), to be used in router advertisement messages. The PIOMD is treated as the same as [RFC4861] Prefix Information Option (PIO) except with an addition of new meta data suboptions.

The PIOMD can coexist with RFC4861 PIO. The prefixes advertised in both PIOMD and PIO can even be the same. It is up to the receiving end host to select the appropriate prefix(es) for configuring its IPv6 addresses. In a case the PIO and the PIOMD share the same prefix, then all the other parameter (like flags and lifetimes) MUST be the same.

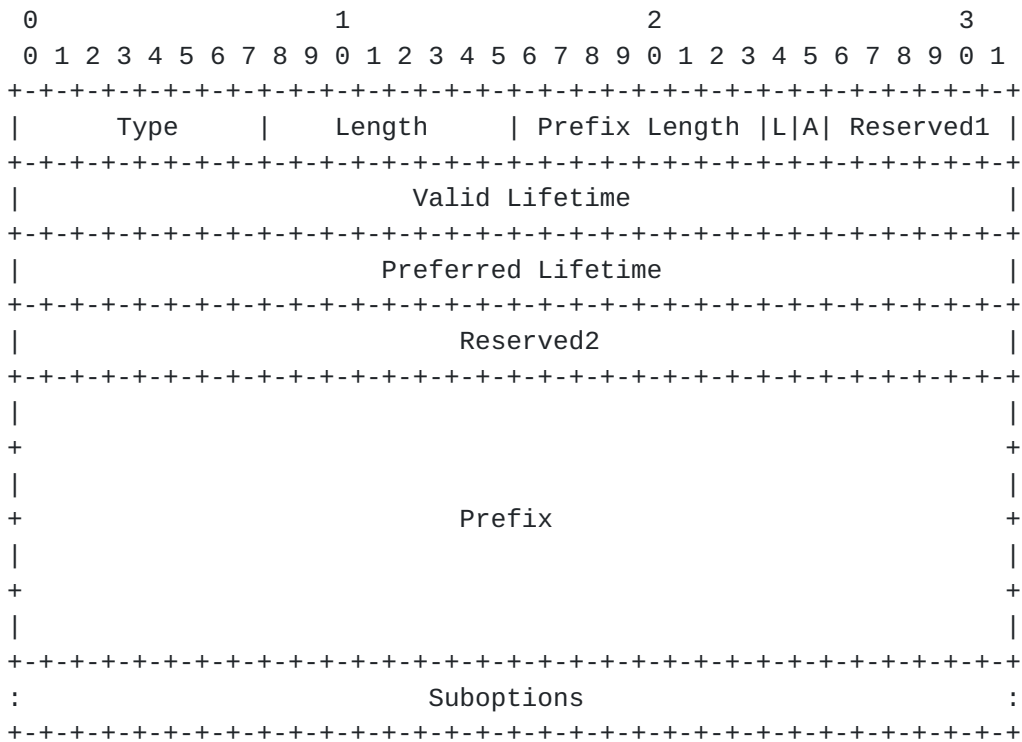


Figure 1: Prefix Information Option with additional meta data

Type

Set to TBD1.

Length

4 if no suboptions are present. Greater than 4 if one or more suboptions are present.

Suboptions

Zero or more suboptions that describe properties and other meta data attached to the advertised prefix. See Section 3.2 for description of the meta data suboption format and suboptions already defined in this specification. The existence of suboptions can be determined from the length field. If the length is greater than 4, then at least one suboption MUST be present.

Rest of the fields are handled exactly as described in Section 4.6.2. of RFC4861 [RFC4861].

3.2. Meta Data Suboptions

The generic suboption format for the PIO with meta data (PIOMD) is shown in Figure 2. The suboption follows the alignment and length rules familiar from [RFC4861]. On a particular note, the flag 'C' describes whether the suboption is mandatory to understand by the receiver or not. If 'C' is set to zero (0), the receiver can silently discard an unknown suboption and skip to the next suboption. If 'C' is set to one (1), then an unknown suboption causes the receiver to silently discard the entire PIOMT and no further suboptions need to be parsed. There can be multiple instances of the same suboption type in one PIOMD option.




Figure 2: Generic meta data suboption format

Figure 3 shows the Prefix Properties suboption. The prefix properties values are defined in Section 6.1. of [I-D.bhandari-dhc-class-based-prefix]. When an end host receives a router advertisement message with a PIOMD and the prefix properties suboption, it can use the suboption information as an additional hint for selecting the prefix for a desired purpose and use case. The prefix properties have global meaning i.e., they have the same treatment and handling cross administrative domains. The value for the 'C' flag SHOULD be one (1). This also implies that if the prefix properties bit vector has a flag bit set, which the receiving end host does not understand and the 'C' flag is also set, then the whole PIOMD option MUST be discarded.




Figure 3: Prefix Properties suboption

Figure 4 shows the Prefix Class suboption. The prefix class values and usage follow what has been defined in Section 2.3. of [I-D.bhandari-dhc-class-based-prefix]. When an end host receives a router advertisement message with a PIOMD and the prefix class suboption, it can use the suboption information as an additional hint for selecting the prefix for a desired purpose and use case. The prefix class has only local administrative meaning i.e., they are local to the access network and may overlap both semantically and registry wise across different administrative domains. How the boundaries of an administrative domain are determined is outside the scope of this specification. The value for the 'C' flag SHOULD be zero (0).




Figure 4: Prefix Class suboption

Future specifications MAY define new suboptions. One potential example could be a suboption to identify the provisioning domain where the configuration information originates.

4. Host Considerations

4.1. Stateless Address Autoconfiguration Enhancements

This specification extends to the [RFC4862] Stateless Address Autoconfiguration (SLAAC). As described in Section 3.1, a new [RFC4861] PIO like option PIOMD can be used to either complement or entirely replace the PIO in a router advertisement. An end host that understands the PIOMD option MUST always prefer a prefix found in the PIOMD over the same prefix found in the PIO option.

[4.2.](#) Internal Data Structures

The host internal data structures need to be extended with the 'prefix property' and the 'prefix class' information associated to the learned prefix and configured addresses. How this is accomplished is host implementation specific. It is also a host implementation issue how an application can learn or query both properties or class of an address or a prefix. One possibility is to provide such information through the socket API extensions (see discussion in [[I-D.ietf-dmm-ondemand-mobility](#)]). Other possibilities include the use of e.g., `ioctl()` or NetLink [[RFC3549](#)] extensions.

[4.3.](#) Default Address Selection

The 'prefix property' is only used as a hint. It does not affect the existing [[RFC6724](#)] automatically. A specific rule to host's policy table has to be inserted by an application or some daemon process. Alternatively, an application can express its address mobility property preferences through the socket API extensions (see discussion in [[I-D.ietf-dmm-ondemand-mobility](#)]), which means the socket library or middleware has to modify [[RFC6724](#)] policy table or algorithm.

The 'prefix properties' flags MAY define the prefix preference for an IP stack that understands the extensions defined in this specification. The IP stack SHOULD use the properties preferences to supersede [[RFC6724](#)] Source Address Selection Rule 8 when selecting a default source address among multiple choices and an application has not explicitly indicate what kind of source address it prefers.

The 'prefix class' defines an application 'class' the advertised prefix is intended to be used for. The class has only local administrative domain significance. The 'prefix class' can be used, for example, to identify prefixes that are meant to be used reach a voice over IP (VoIP) service or a video streaming application within the local administrative network. A specific application in the end host MAY use this additional class information when enumerating through multiple available addresses and then select a specific address to be used for its purposes.

[5.](#) Router Considerations

A network administrator MAY configure routers complying to this specification also send router advertisements with the PIOMD option into every router advertisement that also contains the [[RFC4861](#)] PIO option. Since the PIOMD sending router has no prior knowledge whether the end hosts on the link support the PIOMD option, it is strongly RECOMMENDED that both [[RFC4861](#)] PIO and the PIOMD are always

included in the router advertisement, even if the advertised prefixes were the same. Alternatively (or in addition) multiple provisioning domains [[I-D.ietf-mif-mpvd-ndp-support](#)] can be used to separate prefixes advertised using PIOMD options. See [Section 6](#) for further details.

A router can also make use of the 'C' flag handling in the PIOMD suboptions when introducing new functionality into the network. Since it is possible to include multiple suboptions of the same type into the PIOMD option, the router can easily make a difference between e.g., prefix properties that must be understood by the receiver and those that can safely be ignored.

6. Multiple Provisioning Domain Considerations

Multiple Provisioning Domains (MPVD) framework [[RFC7556](#)] allows grouping network configuration information under an explicitly named provisioning domain [[I-D.ietf-mif-mpvd-id](#)]. This would allow network operators to place mobility related configuration information (including prefixes) under a specific explicit provisioning domain and non-mobile configuration information into other explicit domain or implicit provisioning domain.

MPVDs are the RECOMMENDED way to deliver PIOMD options. This allows mobile network operators selectively advertise mobility related network configurations. MPVDs also provide adequate security features for mobile hosts to verify the authenticity of the configuration information.

7. Security Considerations

Existing Prefix Information Option related security considerations apply as described in [[RFC4861](#)] and [[RFC4191](#)]. A malicious node on the shared link could include stale metadata in a PIOMD causing the host to learn wrong information regarding the prefix and thus make misguided selection of prefixes on the link. Similarly a malicious middleman on the link could modify or remove metadata in the PIOMD causing misguided selection of prefixes. In order to avoid on-link attacks, SeND [[RFC3971](#)] can be used to reject Router Advertisements from potentially malicious nodes and guarantee integrity protection of the Router Advertisements.

If MPVD support for NDP [[I-D.ietf-mif-mpvd-ndp-support](#)] is used, then the mobile host can use its security features to verify the authenticity and correctness of the received PIOMD information.

8. IANA Considerations

[Section 3.1](#) defines a new IPv6 Neighbor Discovery protocol option type TBD1 for the Prefix Information Option with Meta Data. The type value is defined in the existing 'IPv6 Neighbor Discovery Option Formats' IANA registry.

[Section 3.2](#) defines a new IANA registry for the Prefix Information Option with Meta Data suboptions. The registry allocation policy is Standards Action [[RFC5226](#)]. The initial allocations for the prefix properties and prefix class suboptions are listed in [Section 3.2](#).

9. Acknowledgements

The authors thank Ole Troan for his feedback and suggestions on this document (the Classed PIO).

10. References

10.1. Normative References

- [I-D.ietf-mif-mpvd-id]
Krishnan, S., Korhonen, J., Bhandari, S., and S. Gundavelli, "Identification of provisioning domains", [draft-ietf-mif-mpvd-id-02](#) (work in progress), October 2015.
- [I-D.ietf-mif-mpvd-ndp-support]
Korhonen, J., Krishnan, S., and S. Gundavelli, "Support for multiple provisioning domains in IPv6 Neighbor Discovery Protocol", [draft-ietf-mif-mpvd-ndp-support-02](#) (work in progress), October 2015.
- [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>>.
- [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>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.

- [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>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", [RFC 6724](#), DOI 10.17487/RFC6724, September 2012, <<http://www.rfc-editor.org/info/rfc6724>>.

10.2. Informational References

- [I-D.bhandari-dhc-class-based-prefix]
Systems, C., Halwasia, G., Gundavelli, S., Deng, H., Thiebaut, L., Korhonen, J., and I. Farrer, "DHCPv6 class based prefix", [draft-bhandari-dhc-class-based-prefix-05](#) (work in progress), July 2013.
- [I-D.ietf-dmm-ondemand-mobility]
Yegin, A., Kweon, K., Lee, J., Park, J., and D. Moses, "On Demand Mobility Management", [draft-ietf-dmm-ondemand-mobility-02](#) (work in progress), February 2016.
- [RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", [RFC 3493](#), DOI 10.17487/RFC3493, February 2003, <<http://www.rfc-editor.org/info/rfc3493>>.
- [RFC3549] Salim, J., Khosravi, H., Kleen, A., and A. Kuznetsov, "Linux Netlink as an IP Services Protocol", [RFC 3549](#), DOI 10.17487/RFC3549, July 2003, <<http://www.rfc-editor.org/info/rfc3549>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), DOI 10.17487/RFC3971, March 2005, <<http://www.rfc-editor.org/info/rfc3971>>.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", [RFC 4191](#), DOI 10.17487/RFC4191, November 2005, <<http://www.rfc-editor.org/info/rfc4191>>.
- [RFC5014] Nordmark, E., Chakrabarti, S., and J. Laganier, "IPv6 Socket API for Source Address Selection", [RFC 5014](#), DOI 10.17487/RFC5014, September 2007, <<http://www.rfc-editor.org/info/rfc5014>>.

- [RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", [RFC 5213](#), DOI 10.17487/RFC5213, August 2008, <<http://www.rfc-editor.org/info/rfc5213>>.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", [RFC 6275](#), DOI 10.17487/RFC6275, July 2011, <<http://www.rfc-editor.org/info/rfc6275>>.
- [RFC7333] Chan, H., Ed., Liu, D., Seite, P., Yokota, H., and J. Korhonen, "Requirements for Distributed Mobility Management", [RFC 7333](#), DOI 10.17487/RFC7333, August 2014, <<http://www.rfc-editor.org/info/rfc7333>>.
- [RFC7429] Liu, D., Ed., Zuniga, JC., Ed., Seite, P., Chan, H., and CJ. Bernardos, "Distributed Mobility Management: Current Practices and Gap Analysis", [RFC 7429](#), DOI 10.17487/RFC7429, January 2015, <<http://www.rfc-editor.org/info/rfc7429>>.
- [RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", [RFC 7556](#), DOI 10.17487/RFC7556, June 2015, <<http://www.rfc-editor.org/info/rfc7556>>.
- [TS.29274] 3GPP, "3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C)", 3GPP TS 29.060 8.11.0, December 2010.

Authors' Addresses

Jouni Korhonen
Broadcom
3151 Zanker Rd.
CA San Jose
USA

Email: jouni.nospam@gmail.com

Sri Gundavelli
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

Email: sgundave@cisco.com

Pierrick Seite
Orange
4, rue du Clos Courtel, BP 91226
Cesson-Sevigne 35512
France

Email: pierrick.seite@orange.com

Dapeng Liu
Alibaba

Email: max.ldap@alibaba-inc.com