IPv6 Maintenance
Internet-Draft
Intended status: Standards Track
Expires: February 12, 2020

L. Colitti
J. Linkova
Google

August 11, 2019

# Discovering PREF64 in Router Advertisements
## draft-ietf-6man-ra-pref64-04

Abstract

This document specifies a Router Advertisement option to communicate NAT64 prefixes to clients.

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 https://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 February 12, 2020.

Table of Contents



## 1. Introduction

NAT64 [RFC6146] with DNS64 [RFC6147] is a widely-deployed mechanism to provide IPv4 access on IPv6-only networks. In various scenarios, the host must be aware of the NAT64 prefix in use by the network. This document specifies a Router Advertisement [RFC4861] option to communicate the NAT64 prefix to hosts.

## 1.1. 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].

## 1.2. Terminology

Pref64 (or NAT64 prefix): an IPv6 prefix used for IPv6 address synthesis [RFC6146];

NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers ([RFC6146]);

RA: Router Advertisement, a message used by IPv6 routers to advertise their presence together with various link and Internet parameters ([RFC4861]);

DNS64: a mechanism for synthesizing AAAA records from A records ([RFC6147]);

## 2. Use cases for communicating the NAT64 prefix to hosts

On networks employing NAT64, it is useful for hosts to know the NAT64 prefix for several reasons, including the following:

- Local DNSSEC validation. As discussed in [RFC6147] section 2, the stub resolver in the host "will try to obtain (real) AAAA RRs, and in case they are not available, the DNS64 function will synthesize AAAA RRs for internal usage." This is required in order to use DNSSEC on a NAT64 network.

- IPv4 address literals on an IPv6-only host. As described in [RFC8305] section 7.1, IPv6-only hosts connecting to IPv4 address literals can resolve the IPv4 literal to an IPv6 address.

- 464XLAT [RFC6877]. 464XLAT is widely deployed and requires that the host be aware of the NAT64 prefix.

- Trusted DNS server. AAAA synthesis is required for the host to be able to use a DNS server not provided by the network (e.g., a DNS-over-TLS server ([RFC7858]) with which the host has an existing trust relationship).

- Networks with no DNS64 server. Hosts that support AAAA synthesis and that are aware of the NAT64 prefix in use do not need the network to perform the DNS64 function at all.

## 3. Why include the NAT64 prefix in Router Advertisements

Fate sharing: NAT64 requires a routing to be configured. IPv6 routing configuration requires receiving an IPv6 Router Advertisement [RFC4861]. Compared to currently-deployed NAT64 prefix discovery methods such as [RFC7050], including the NAT64 prefix in the Router Advertisement minimizes the number of packets required to configure a host. This speeds up the process of connecting to a network that supports NAT64/DNS64, and simplifies host implementation by removing the possibility that the host can have an incomplete layer 3 configuration (e.g., IPv6 addresses and prefixes, but no NAT64 prefix).

Updatability: it is possible to change the NAT64 prefix at any time, because when it changes, it is possible to notify hosts by sending a new Router Advertisement.

Deployability: all IPv6 hosts and networks are required to support [RFC4861]. Other options such as [RFC7225] require implementing other protocols.

## 4. Semantics

To support prefix lengths defined in ([RFC6052]) this option contains the prefix length field. However as /96 prefix is considered to be the most common use case, the prefix length field is optional and only presents for non-/96 prefixes. It allows to keep the option length to a minimum (16 bytes) for the most common case and increase it to 20 bytes for non-/96 prefixes only (see Section 5 below for more details).

This option specifies exactly one NAT64 prefix for all IPv4 destinations. If the network operator desires to route different parts of the IPv4 address space to different NAT64 devices, this can be accomplished by routing more specifics of the NAT64 prefix to those devices. For example, if the operator would like to route 10.0.0.0/8 through NAT64 device A and the rest of the IPv4 space through NAT64 device B, and the operator's NAT64 prefix is 2001:db8:a:b::/96, then the operator can route 2001:db8:a:b::a00:0/104 to NAT64 A and 2001:db8:a:b::/64 to NAT64 B.

This option may appear more than once in a Router Advertisement (e.g. in case of graceful renumbering the network from one NAT64 prefix to another). Host behaviour with regards to synthesizing IPv6 addresses from IPv4 addresses SHOULD follow the recommendations given in Section 3 of [RFC7050], limited to the NAT64 prefixes that have non-zero lifetime.

In a network (or a provisioning domain) that provides both IPv4 and NAT64, it may be desirable for certain IPv4 addresses not to be translated. An example might be private address ranges that are local to the network/provisioning domain and should not be reached through the NAT64. This type of configuration cannot be conveyed to hosts using this option, or through other NAT64 prefix provisioning mechanisms such as [RFC7050] or [RFC7225]. This problem does not apply in IPv6-only networks, because in such networks, the host does not have an IPv4 address and cannot reach any IPv4 destinations without the NAT64. The multihoming and multiple provisioning domains scenarios are discussed in Section 7.

## 5. Option format

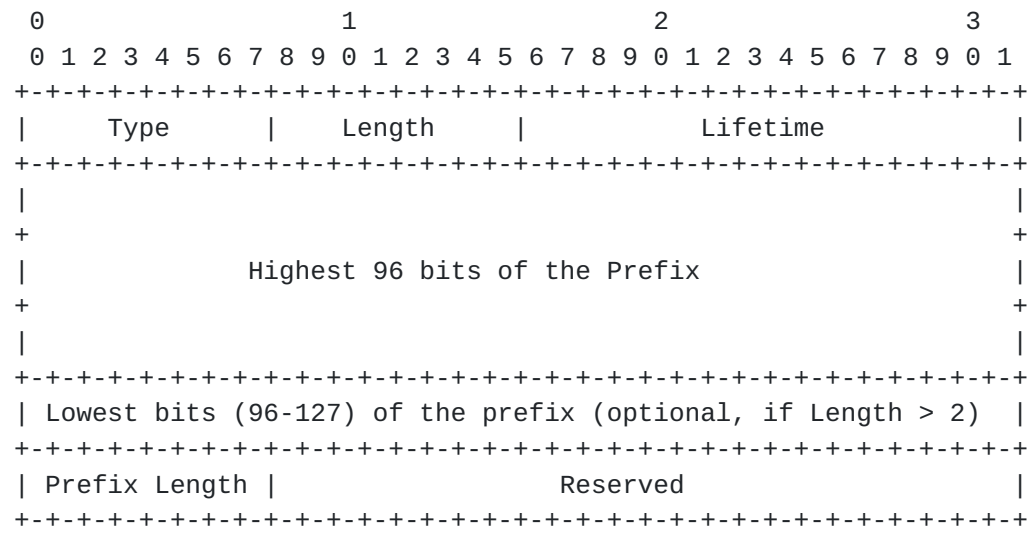


Figure 1: NAT64 Prefix Option Format

Fields:

| | |
|---|---|
| Type | 8-bit identifier of the Pref64 option type as assigned by IANA: TBD |
| Length | 8-bit unsigned integer. The length of the option (including the Type and Length fields) is in units of 8 octets. If the prefix length is 96 bits the sender MUST set the Length to 2 and include the 96 bits of the prefix in the option. If the prefix length is not 96 bits then the sender MUST set the length to 3 and include all 128 bits of the prefix in the Prefix field and set the Prefix Length field to the prefix length. The receiver MUST ignore the Pref64 option if the length field value is 1. If the Length field value exceeds 3, the receiver MUST utilize the first 21 octets and ignore the rest of the option. |
| Lifetime | 16-bit unsigned integer. The maximum time in seconds over which this NAT64 prefix MAY be used. The value of Lifetime SHOULD by default be set to lesser of 3 x MaxRtrAdvInterval or 65535 seconds. A value of zero means that the prefix MUST no longer be used. |
| Highest 96 bits of the prefix | 96-bit unsigned integer. Contains bits 0 - 95 of the NAT64 prefix. |
| Lowest bits of the prefix | 32-bit unsigned integer. Contains bits 96 - 127 of the NAT64 prefix. This field is optional and presents only if the prefix length is not 96 bits. |
| Prefix Length | 8-bit unsigned integer. Optional field which present only if the prefix length is not 96 bits. The sender MUST set it only to one of the following values: 32, 40, 48, 56, 64 ([RFC6052]. The receiver MUST ignore the Pref64 option if the prefix length value is not set to one of those numbers. |
| Reserved | A 3-byte unused field. If present it MUST be initialized to zero by the sender and MUST be ignored by the receiver. This field is optional and presents only if the prefix length is not 96 bits. |

## 6. Handling Multiple NAT64 Prefixes

In some cases a host may receive multiple NAT64 prefixes from different sources. Possible scenarios include (but are not limited to):

- the host is using multiple mechanisms to discover Pref64 prefixes (e.g. by using PCP ([RFC7225]) and/or by resolving IPv4-only fully qualified domain name ([RFC7050]) in addition to receiving the Pref64 RA option);

- The pref64 option presents in a single RA more than once;

- the host receives multiple RAs with different Pref64 prefixes on one or multiple interfaces.

When multiple Pref64 were discovered via RA Pref64 Option (the Option presents more than once in a single RA or multiple RAs were received), host behaviour with regards to synthesizing IPv6 addresses from IPv4 addresses SHOULD follow the recommendations given in Section 3 of [RFC7050], limited to the NAT64 prefixes that have non-zero lifetime..

When different Pref64 are discovered by using multiple mechanisms, hosts SHOULD select one source of information only. The RECOMMENDED order is:

- PCP-discovered prefixes ([RFC7225]), if supported;

- Pref64 discovered via RA Option;

- Pref64 resolving IPv4-only fully qualified domain name ([RFC7050])

Note that if the network provides Pref64 both via this RA option and [RFC7225], hosts that receive the Pref64 via RA option may choose to use it immediately before waiting for PCP to complete, and therefore some traffic may not reflect any more detailed configuration provided by PCP.

## 7. Multihoming

Like most IPv6 configuration information, the Pref64 option is specific to the network on which it is received. For example, a Pref64 option received on a particular wireless network may not be usable unless the traffic is also sourced on that network. Similarly, a host connected to a cellular network that provides NAT64 generally cannot use that NAT64 for destinations reached through a VPN tunnel that terminates outside that network.

Thus, correct use of this option on a multihomed host generally requires the host to support the concept of multiple Provisioning Domains (PvD, a set of configuration information associated with a network, [RFC7556]) and to be able to use these PvDs.

This issue is not specific to the Pref64 RA option and, for example, is quite typical for DNS resolving on multihomed hosts (e.g. a host might resolve a destination name by using the corporate DNS server via the VPN tunnel but then send the traffic via its Internet-facing interface).

## 8. Pref64 Consistency

Section 6.2.7 of [RFC4861] recommends that routers inspect RAs sent by other routers to ensure that all routers onlink advertise the consistent information. Routers SHOULD inspect valid Pref64 options received on a given link and verify the consistency. Detected inconsistencies indicate that one or more routers might be misconfigured. Routers SHOULD log such cases to system or network management. Routers SHOULD check and compare the following information:

- set of Pref64 with non-zero lifetime;
- set of Pref64 with zero lifetime.

PvD-aware routers MUST only compare information scoped to the same implicit or explicit PvD.

## 9. IANA Considerations

The IANA is requested to assign a new IPv6 Neighbor Discovery Option type for the PREF64 option defined in this document.

| Option Name | Type |
|---|---|
| PREF64 option | (TBD) |

Table 1

The IANA registry for these options is:

https://www.iana.org/assignments/icmpv6-parameters [1]

## 10. Security Considerations

Because Router Advertisements are required in all IPv6 configuration scenarios, on IPv6-only networks, Router Advertisements must already be secured, e.g., by deploying RA guard [RFC6105]. Providing all configuration in Router Advertisements increases security by ensuring

that no other protocols can be abused by malicious attackers to provide hosts with invalid configuration.

The security measures that must already be in place to ensure that Router Advertisements are only received from legitimate sources eliminate the problem of NAT64 prefix validation described in section 3.1 of [RFC7050].

## 11. Acknowledgements

Thanks to the following people (in alphabetical order) for their review and feedback: Mikael Abrahamsson, Mark Andrews, Brian E Carpenter, David Farmer, Nick Heatley, Martin Hunek, Tatuya Jinmei, Erik Kline, David Lamparter, Jordi Palet Martinez, Tommy Pauly, Michael Richardson, David Schinazi.

## 12. References

### 12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://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, <https://www.rfc-editor.org/info/rfc4861>.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <https://www.rfc-editor.org/info/rfc6052>.

### 12.2. Informative References

[I-D.ietf-intarea-provisioning-domains] Pfister, P., Vyncke, E., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", draft-ietf-intarea-provisioning-domains-05 (work in progress), June 2019.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <https://www.rfc-editor.org/info/rfc4033>.

[RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, DOI 10.17487/RFC6105, February 2011, <https://www.rfc-editor.org/info/rfc6105>.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <https://www.rfc-editor.org/info/rfc6146>.

[RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, DOI 10.17487/RFC6147, April 2011, <https://www.rfc-editor.org/info/rfc6147>.

[RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <https://www.rfc-editor.org/info/rfc6877>.

[RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <https://www.rfc-editor.org/info/rfc7050>.

[RFC7225] Boucadair, M., "Discovering NAT64 IPv6 Prefixes Using the Port Control Protocol (PCP)", RFC 7225, DOI 10.17487/RFC7225, May 2014, <https://www.rfc-editor.org/info/rfc7225>.

[RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <https://www.rfc-editor.org/info/rfc7556>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <https://www.rfc-editor.org/info/rfc8305>.

### 12.3. URIs

[1] https://www.iana.org/assignments/icmpv6-parameters

## Authors' Addresses

Lorenzo Colitti
Google
Roppongi 6-10-1
Minato, Tokyo 106-6126
JP

Email: lorenzo@google.com

Jen Linkova
Google
1 Darling Island Rd
Pyrmont, NSW 2009
AU

Email: furry@google.com