

Network Working Group	T. Bastian
Internet-Draft	Ecole Normale Supérieure, Paris
Updates: 6126bis (if approved)	J. Chroboczek
Intended status: Experimental	IRIF, University of Paris-Diderot
Expires: November 15, 2020	May 14, 2020

Announcing IPv4 routes with an IPv6 next-hop in the Babel routing
protocol
draft-bastian-babel-v4ov6-01

Abstract

This document defines an extension to the Babel routing protocol that allows announcing routes to an IPv4 prefix with an IPv6 next-hop, which makes it possible for IPv4 traffic to flow through interfaces that have not been assigned an IPv4 address.

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 November 15, 2020.

Copyright Notice

Copyright (c) 2020 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 (<https://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
1.1. Specification of Requirements	3
2. Protocol operation	3
2.1. Announcing v4-over-v6 routes	3
2.2. Receiving v4-over-v6 routes	3
2.3. Prefix and seqno requests	4
2.4. Other TLVs	4
3. Backwards compatibility	4
4. Protocol encoding	5
4.1. Prefix encoding	5
4.2. Changes for existing TLVs	5
5. IANA Considerations	6
6. Security Considerations	6
7. References	7
7.1. Normative References	7
7.2. Informative References	7
Authors' Addresses	7

1. Introduction

Traditionally, a routing table maps a network prefix of a given address family to a next-hop address in the same address family. The sole purpose of this next-hop address is to serve as an input to a protocol that will map it to a link-layer address, Neighbour Discovery (ND) [RFC4861] in the case of IPv6, Address Resolution (ARP) [RFC0826] in the case of IPv4. Therefore, there is no reason why the address family of the next hop address should match that of the prefix being announced: an IPv6 next-hop yields a link-layer address that is suitable for forwarding both IPv6 or IPv4 traffic.

We call a route towards an IPv4 prefix that uses an IPv6 next hop a "v4-over-v6" route. Since an IPv6 next-hop can use a link-local address that is autonomously configured, the use of v4-over-v6 routes enables a mode of operation where the network core has no statically assigned IP addresses of either family, thus significantly reducing the amount of manual configuration.

This document describes an extension that allows the Babel routing protocol [RFC6126bis] to announce routes towards IPv6 prefixes with IPv4 next hops. The extension is inspired by a previously defined extension to the BGP protocol [RFC5549].

1.1. Specification of Requirements

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Protocol operation

The Babel protocol fully supports double-stack operation: all data that represent a neighbour address or a network prefix are tagged by an Address Encoding (AE), a small integer that identifies the address family (IPv4 or IPv6) of the address of prefix, and describes how it is encoded. This extension defines a new AE, called v4-over-v6, which has the same format as the existing AE for IPv4 addresses. This new AE is only allowed in TLVs that carry network prefixes: TLVs that carry a neighbour address use the normal encodings for IPv6 addresses.

2.1. Announcing v4-over-v6 routes

A Babel node that needs to announce an IPv4 route over an interface that has no assigned IPv4 address MAY make a v4-over-v6 announcement. In order to do so, it first establishes an IPv6 next-hop address in the usual manner (either by sending the Babel packet over IPv6, or by including a Next Hop TLV containing an IPv6 address); it then sends an Update with AE equal to TBD containing the IPv4 prefix being announced.

If the outgoing interface has been assigned an IPv4 address, then, in the interest of maximising compatibility with existing routers, the sender SHOULD prefer an ordinary IPv4 announcement; even in that case, however, it MAY use a v4-over-v6 announcement. A node SHOULD NOT send both ordinary IPv4 and v4-over-v6 announcements for the same prefix over a single interface (if the update is sent to a multicast address) or to a single neighbour (if sent to a unicast address), since doing that doubles the amount of routing traffic while providing no benefit.

2.2. Receiving v4-over-v6 routes

Upon reception of an Update TLV with a v4-over-v6 AE, a Babel node computes the IPv6 next-hop, as described in Section 4.6.9 of [RFC6126bis]. If no IPv6 next-hop exists, then the Update MUST be silently ignored. If an IPv6 next-hop exists, then the node MAY acquire the route being announced, as described in Section 3.5.3 of [RFC6126bis]; the parameters of the route are as follows:

- o the prefix, plen, router-id, seqno, metric MUST be computed as for an IPv4 route, as described in Section 4.6.9 of [RFC6126bis];
- o the next-hop MUST be computed as for an IPv6 route, as described in Section 4.6.9 of [RFC6126bis]: it is taken from the last preceding Next-Hop TLV with an AE field equal to 2 or 3; if no such entry exists, and if the Update TLV has been sent in a Babel packet carried over IPv6, then the next-hop is the network-layer source address of the packet.

As usual, a node MAY ignore the update, e.g., due to filtering (Appendix C of [RFC6126bis]). If a node cannot install v4-over-v6 routes, eg., due to hardware or software limitations, then routes to an IPv4 prefix with an IPv6 next-hop MUST NOT be selected, as described in Section 3.5.3 of [RFC6126bis].

2.3. Prefix and seqno requests

Prefix and seqno requests are used to request an update for a given prefix. Since they are not related to a specific Next-Hop, there is no semantic difference between ordinary IPv4 and v4-over-v6 requests.

A node SHOULD NOT send requests of either kind with the AE field being set to TBD (v4-over-v6); instead, it SHOULD request IPv4 updates using requests with the AE field being set to 1 (IPv4).

When receiving requests, AEs 1 (IPv4) and TBD (v4-over-v6) MUST be treated in the same manner: the receiver processes the request as described in Section 3.8 of [RFC6126bis]. If an Update is sent, then it MAY be sent with AE 1 or TBD, as described in Section 2.1 above, irrespective of which AE was used in the request.

When receiving a request with AE 0 (wildcard), the receiver SHOULD send a full route dump, as described in Section 3.8.1.1 of [RFC6126bis]. Any IPv4 routes contained in the route dump MAY use either AE 1 or AE TBD, as described in Section 2.1 above.

2.4. Other TLVs

The only other TLV defined by [RFC6126bis] that carries an AE field is the IHU TLV. IHU TLVs MUST NOT carry the AE TBD (v4-over-v6).

3. Backwards compatibility

This protocol extension adds no new TLVs or sub-TLVs.

This protocol extension uses a new AE. As discussed in Appendix D of [RFC6126bis] and specified in the same document, implementations that

do not understand the present extension will silently ignore the various TLVs that use this new AE. As a result, incompatible versions will ignore v4-over-v6 routes. They will also ignore requests with AE TBD, which, as stated in Section 2.3, are NOT RECOMMENDED.

Using a new AE introduces a new compression state, used to parse the network prefixes. As this compression state is separate from other AEs' states, it will not interfere with the compression state of unextended nodes.

This extension reuses the next-hop state from AEs 2 and 3 (IPv6), but makes no changes to the way it is updated, and therefore causes no compatibility issues.

As mentioned in Section 2.1, ordinary IPv4 announcements are preferred to v4-over-v6 announcements when the outgoing interface has an assigned IPv4 address; doing otherwise would prevent routers that do not implement this extension from learning the route being announced.

4. Protocol encoding

This extension defines the v4-over-v6 AE, whose value is TBD. This AE is solely used to tag network prefixes, and MUST NOT be used to tag peers' addresses, eg. in Next-Hop or IHU TLVs.

This extension defines no new TLVs or sub-TLVs.

4.1. Prefix encoding

Network prefixes tagged with AE TBD MUST be encoded and decoded as prefixes tagged with AE 1 (IPv4), as described in Section 4.3.1 of [RFC6126bis].

A new compression state for AE TBD (v4-over-v6) distinct from that of AE 1 (IPv4) is introduced, and MUST be used for address compression of prefixes tagged with AE TBD, as described in Section 4.6.9 of [RFC6126bis]

4.2. Changes for existing TLVs

The following TLVs MAY be tagged with AE TBD:

- o Update (Type = 8)
- o Route Request (Type = 9)

- o Seqno Request (Type = 10)

As AE TBD is suitable only to tag network prefixes, IHU (Type = 5) and Next-Hop (Type = 7) TLVs MUST NOT be tagged with AE TBD. Such TLVs MUST be silently ignored.

4.2.1. Update

An Update (Type = 8) TLV with AE = TBD is constructed as described in Section 4.6.9 of [RFC6126bis] for AE 1 (IPv4), with the following specificities:

- o Prefix. The Prefix field is constructed according to the Section 4.1 above.
- o Next hop. The next hop is determined as described in Section 2.2 above.

4.2.2. Other valid TLVs tagged with AE = TBD

Any other valid TLV tagged with AE = TBD MUST be constructed and decoded as described in Section 4.6 of [RFC6126bis]. Network prefixes within MUST be constructed and decoded as described in Section 4.1 above.

5. IANA Considerations

IANA is requested to allocate a value (4 suggested) in the "Babel Address Encodings" registry as follows:

AE	Name	Reference
TBD	v4-over-v6	(this document)

6. Security Considerations

This extension does not fundamentally change the security properties of the Babel protocol: as described in Section 6 of [RFC6126bis], Babel must be protected by a suitable cryptographic mechanism in order to be made secure.

However, enabling this extension will allow IPv4 traffic to flow through sections of a network that have not been assigned IPv4 addresses, which, in turn, might allow IPv4 traffic to reach areas of the network that were previously inaccessible to such traffic. If this is undesirable, the flow of IPv4 traffic must be restricted by

the use of suitable filtering rules (Appendix C of [RFC6126bis]) together with matching access control rules in the data plane.

7. References

7.1. 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.
- [RFC6126bis] Chroboczek, J. and D. Schinazi, "The Babel Routing Protocol", draft-ietf-babel-rfc6126bis-17 (work in progress), February 2020.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017.

7.2. Informative References

- [RFC0826] Plummer, D., "An Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, DOI 10.17487/RFC0826, November 1982.
- [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.
- [RFC5549] Le Faucheur, F. and E. Rosen, "Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop", RFC 5549, DOI 10.17487/RFC5549, May 2009.

Authors' Addresses

Theophile Bastian
Ecole Normale Supérieure, Paris
France

Email: theophile.bastian@ens.fr

Juliusz Chroboczek
IRIF, University of Paris-Diderot
Case 7014
75205 Paris Cedex 13
France

Email: jch@irif.fr