

Internet Engineering Task Force (IETF)
Request for Comments: 9080
Category: Standards Track
ISSN: 2070-1721

J. Chroboczek
IRIF, University of Paris-Diderot
August 2021

Homenet Profile of the Babel Routing Protocol

Abstract

This document defines the exact subset of the Babel routing protocol and its extensions that is required by an implementation of the Homenet protocol suite, as well as the interactions between the Home Networking Control Protocol (HNCP) and Babel.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9080>.

Copyright Notice

Copyright (c) 2021 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
 - 1.1. Requirements Language
 - 1.2. Background
2. The Homenet Profile of Babel
 - 2.1. Requirements
 - 2.2. Optional Features
3. Interactions between HNCP and Babel
 - 3.1. Requirements
 - 3.2. Optional Features
4. Security Considerations
5. IANA Considerations
6. References
 - 6.1. Normative References
 - 6.2. Informative References

Acknowledgments

Author's Address

1. Introduction

The core of the Homenet protocol suite consists of the Home Networking Control Protocol (HNCP) [RFC7788], a protocol used for flooding configuration information and assigning prefixes to links, combined with the Babel routing protocol [RFC8966]. Babel is an extensible, flexible, and modular protocol: minimal implementations of Babel have been demonstrated that consist of a few hundred lines of code, while the "large" implementation includes support for a number of extensions and consists of over ten thousand lines of C code.

This document consists of two parts. The first specifies the exact subset of the Babel protocol and its extensions that is required by an implementation of the Homenet protocol suite. The second specifies how HNCP interacts with Babel.

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

1.2. Background

The Babel routing protocol and its extensions are defined in a number of documents:

- * RFC 8966 [RFC8966] defines the Babel routing protocol. It allows Babel's control data to be carried either over link-local IPv6 or over IPv4 and in either case allows announcing both IPv4 and IPv6 routes. It leaves link cost estimation, metric computation, and route selection to the implementation. Distinct implementations of Babel [RFC8966] will interoperate, in the sense that they will maintain a set of loop-free forwarding paths. However, if they implement conflicting options, they might not be able to exchange a full set of routes. In the worst case, an implementation that only implements the IPv6 subset of the protocol and an implementation that only implements the IPv4 subset of the protocol will not exchange any routes. In addition, if implementations use conflicting route selection policies, persistent oscillations might occur.
- * The informative Appendix A of [RFC8966] suggests a simple and easy-to-implement algorithm for cost and metric computation that has been found to work satisfactorily in a wide range of topologies.
- * While RFC 8966 does not provide an algorithm for route selection, its Section 3.6 suggests selecting the route with the smallest metric with some hysteresis applied. An algorithm that has been found to work well in practice is described in Section III.E of [DELAY-BASED].
- * Four documents define optional extensions to Babel: authentication based on Hashed Message Authentication Code (HMAC) [RFC8967], source-specific routing [RFC9079], delay-based routing [BABEL-RTT], and ToS-specific (Type of Service) routing [ToS-SPECIFIC]. All of these extensions interoperate with the core protocol as well as with each other.

2. The Homenet Profile of Babel

2.1. Requirements

- REQ1: A Homenet implementation of Babel MUST encapsulate Babel control traffic in IPv6 packets sent to the IANA-assigned port 6696 and either the IANA-assigned multicast group ff02::1:6 or to a link-local unicast address.

Rationale: Since Babel is able to carry both IPv4 and IPv6 routes over either IPv4 or IPv6, choosing the protocol used for carrying control traffic is a matter of preference. Since IPv6 has some features that make implementations somewhat simpler and more reliable (notably properly scoped and reasonably stable link-local addresses), we require carrying control data over IPv6.

- REQ2: A Homenet implementation of Babel MUST implement the IPv6 subset of the protocol defined in the body of RFC 8966.

Rationale: Support for IPv6 routing is an essential component of the Homenet architecture.

- REQ3: A Homenet implementation of Babel SHOULD implement the IPv4 subset of the protocol defined in the body of RFC 8966. Use of other techniques for acquiring IPv4 connectivity (such as multiple layers of NAT) is strongly discouraged.

Rationale: Support for IPv4 will likely remain necessary for years to come, and even in pure IPv6 deployments, including code for supporting IPv4 has very little cost. Since HNCP makes it easy to assign distinct IPv4 prefixes to the links in a network, it is not necessary to resort to multiple layers of NAT, with all of its problems.

- REQ4: A Homenet implementation of Babel MUST implement source-specific routing for IPv6, as defined in RFC 9079 [RFC9079].

Rationale: Source-specific routing is an essential component of the Homenet architecture. Source-specific routing for IPv4 is not required, since HNCP arranges things so that a single nonspecific IPv4 default route is announced (Section 6.5 of [RFC7788]).

- REQ5: A Homenet implementation of Babel must use metrics that are of a similar magnitude to the values suggested in Appendix A of [RFC8966]. In particular, it SHOULD assign costs that are no less than 256 to wireless links and SHOULD assign costs between 32 and 196 to lossless wired links.

Rationale: If two implementations of Babel choose very different values for link costs, combining routers from different vendors will cause suboptimal routing.

- REQ6: A Homenet implementation of Babel SHOULD distinguish between wired and wireless links; if it is unable to determine whether a link is wired or wireless, it SHOULD make the worst-case hypothesis that the link is wireless. It SHOULD dynamically probe the quality of wireless links and derive a suitable metric from its quality estimation. Appendix A of [RFC8966] gives an example of a suitable algorithm.

Rationale: Support for wireless transit links is a distinguishing feature of Homenet, and one that is requested by our users. In the absence of dynamically computed metrics, the routing protocol attempts to minimise the number of links crossed by a route and therefore prefers long, lossy links to shorter, lossless

ones. In wireless networks, "hop-count routing is worst-path routing".

While it would be desirable to perform link-quality probing on some wired link technologies, notably power-line networks, these kinds of links tend to be difficult or impossible to detect automatically, and we are not aware of any published link-quality algorithms for them. Hence, we do not require link-quality estimation for wired links of any kind.

2.2. Optional Features

OPT1: A Homenet implementation of Babel MAY perform route selection by applying hysteresis to route metrics, as suggested in Section 3.6 of [RFC8966] and described in detail in Section III.E of [DELAY-BASED]. However, hysteresis is not required, and the implementation may simply pick the route with the smallest metric.

Rationale: Hysteresis is only useful in congested and highly dynamic networks. In a typical home network, which is stable and uncongested, the feedback loop that hysteresis compensates for does not occur.

OPT2: A Homenet implementation of Babel may include support for other extensions to the protocol, as long as they are known to interoperate with both the core protocol and source-specific routing.

Rationale: A number of extensions to the Babel routing protocol have been defined over the years; however, they are useful in fairly specific situations, such as routing over global-scale overlay networks [BABEL-RTT] or multi-hop wireless networks with multiple radio frequencies [BABEL-Z]. Hence, with the exception of source-specific routing, no extensions are required for Homenet.

3. Interactions between HNCP and Babel

The Homenet architecture cleanly separates configuration, which is done by HNCP, from routing, which is done by Babel. While the coupling between the two protocols is deliberately kept to a minimum, some interactions are unavoidable.

All the interactions between HNCP and Babel consist of HNCP causing Babel to perform an announcement on its behalf (under no circumstances does Babel cause HNCP to perform an action). How this is realised is an implementation detail that is outside the scope of this document; while it could conceivably be done using a private communication channel between HNCP and Babel, in existing implementations, HNCP installs a route in the operating system's kernel that is later picked up by Babel using the existing redistribution mechanisms.

3.1. Requirements

REQ7: If an HNCP node receives a DHCPv6 prefix delegation for prefix P and publishes an External-Connection TLV containing a Delegated-Prefix TLV with prefix P and no Prefix-Policy TLV, then it MUST announce a source-specific default route with source prefix P over Babel.

Rationale: Source-specific routes are the main tool that Homenet uses to enable optimal routing in the presence of multiple IPv6 prefixes. External connections with

nontrivial prefix policies are explicitly excluded from this requirement, since their exact behaviour is application specific.

REQ8: If an HNCP node receives a DHCPv4 lease with an IPv4 address and wins the election for NAT gateway, then it MUST act as a NAT gateway and MUST announce a (nonspecific) IPv4 default route over Babel.

Rationale: The Homenet stack does not use source-specific routing for IPv4; instead, HNCP elects a single NAT gateway and publishes a single default route towards that gateway ([RFC7788], Section 6.5).

REQ9: If an HNCP node assigns a prefix P to an attached link and announces P in an Assigned-Prefix TLV, then it MUST announce a route towards P over Babel.

Rationale: Prefixes assigned to links must be routable within the Homenet.

3.2. Optional Features

OPT3: An HNCP node that receives a DHCPv6 prefix delegation MAY announce a nonspecific IPv6 default route over Babel in addition to the source-specific default route mandated by requirement REQ7.

Rationale: Since the source-specific default route is more specific than the nonspecific default route, the former will override the latter if all nodes implement source-specific routing. Announcing an additional nonspecific route is allowed, since doing that causes no harm and might simplify operations in some circumstances, e.g., when interoperating with a routing protocol that does not support source-specific routing.

OPT4: An HNCP node that receives a DHCPv4 lease with an IPv4 address and wins the election for NAT gateway SHOULD NOT announce a source-specific IPv4 default route.

Rationale: Homenet does not require support for IPv4 source-specific routing. Announcing IPv4 source-specific routes will not cause routing pathologies (blackholes or routing loops), but it might cause packets sourced in different parts of the Homenet to follow different paths, with all the confusion that this entails.

4. Security Considerations

Both HNCP and Babel carry their control data in IPv6 packets with a link-local source address, and implementations are required to drop packets sent from a global address. Hence, they are only susceptible to attacks from a directly connected link on which the HNCP and Babel implementations are listening.

The security of a Homenet network relies on having a set of "Internal", "Ad Hoc", and "Hybrid" interfaces (Section 5.1 of [RFC7788]) that are assumed to be connected to links that are secured at a lower layer. HNCP and Babel packets are only accepted when they originate on these trusted links. "External" and "Guest" interfaces are connected to links that are not trusted, and any HNCP or Babel packets that are received on such interfaces are ignored. ("Leaf" interfaces are a special case since they are connected to trusted links, but HNCP and Babel traffic received on such interfaces is ignored.) This implies that the security of a Homenet network

depends on the reliability of the border discovery procedure described in Section 5.3 of [RFC7788].

If untrusted links are used for transit, which is NOT RECOMMENDED, then any HNCP and Babel traffic that is carried over such links MUST be secured using an upper-layer security protocol. While both HNCP and Babel support cryptographic authentication, at the time of writing, no protocol for autonomous configuration of HNCP and Babel security has been defined.

5. IANA Considerations

This document has no IANA actions.

6. References

6.1. 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>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<https://www.rfc-editor.org/info/rfc7788>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8966] Chroboczek, J. and D. Schinazi, "The Babel Routing Protocol", RFC 8966, DOI 10.17487/RFC8966, January 2021, <<https://www.rfc-editor.org/info/rfc8966>>.
- [RFC9079] Boutier, M. and J. Chroboczek, "Source-Specific Routing in the Babel Routing Protocol", RFC 9079, DOI 10.17487/RFC9079, August 2021, <<https://www.rfc-editor.org/rfc/rfc9079>>.

6.2. Informative References

- [BABEL-RTT] Jonglez, B. and J. Chroboczek, "Delay-based Metric Extension for the Babel Routing Protocol", Work in Progress, Internet-Draft, draft-ietf-babel-rtt-extension-00, 26 April 2019, <<https://datatracker.ietf.org/doc/html/draft-ietf-babel-rtt-extension-00>>.
- [BABEL-Z] Chroboczek, J., "Diversity Routing for the Babel Routing Protocol", Work in Progress, Internet-Draft, draft-chroboczek-babel-diversity-routing-01, 15 February 2016, <<https://datatracker.ietf.org/doc/html/draft-chroboczek-babel-diversity-routing-01>>.
- [DELAY-BASED] Jonglez, B., Boutier, M., and J. Chroboczek, "A delay-based routing metric", March 2014, <<http://arxiv.org/abs/1403.3488>>.
- [RFC8967] D, C., Kolodziejak, W., and J. Chroboczek, "MAC Authentication for the Babel Routing Protocol", RFC 8967, DOI 10.17487/RFC8967, January 2021, <<https://www.rfc-editor.org/info/rfc8967>>.
- [ToS-SPECIFIC]

Chouasne, G. and J. Chroboczek, "TOS-Specific Routing in Babel", Work in Progress, Internet-Draft, draft-chouasne-babel-tos-specific-00, 3 July 2017, <<https://datatracker.ietf.org/doc/html/draft-chouasne-babel-tos-specific-00>>.

Acknowledgments

A number of people have helped with defining the requirements listed in this document. I am especially indebted to Barbara Stark and Markus Stenberg.

Author's Address

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

Email: jch@irif.fr