

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: December 17, 2015

P. Pfister  
B. Paterson  
Cisco Systems  
J. Arkko  
Ericsson  
June 15, 2015

Distributed Prefix Assignment Algorithm  
draft-ietf-homenet-prefix-assignment-07

Abstract

This document specifies a distributed algorithm for automatic prefix assignment. Thus it provides an alternative to manual or centralized prefix and address assignment techniques. Given a set of delegated prefixes, it ensures that at most one prefix is assigned from each delegated prefix to each link. Nodes may assign available prefixes to the links they are directly connected to, or for other private purposes. The algorithm eventually converges and ensures that all assigned prefixes do not overlap.

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 December 17, 2015.

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. Subroutine Specific Terminology . . . . .	5
3. Applicability Statement . . . . .	6
4. Algorithm Specification . . . . .	8
4.1. Prefix Assignment Algorithm Subroutine . . . . .	8
4.2. Overriding and Destroying Existing Assignments . . . . .	10
4.3. Other Events . . . . .	12
5. Prefix Selection Considerations . . . . .	12
6. Implementation Capabilities and Node Behavior . . . . .	14
7. Algorithm Parameters . . . . .	15
8. Security Considerations . . . . .	16
9. IANA Considerations . . . . .	17
10. Acknowledgments . . . . .	17
11. References . . . . .	17
11.1. Normative References . . . . .	17
11.2. Informative References . . . . .	17
Appendix A. Static Configuration Example . . . . .	17
Authors' Addresses . . . . .	18

## 1. Introduction

This document specifies a distributed algorithm for automatic prefix assignment. The algorithm provides a generic alternative to centralized (human or software based) approaches for network prefixes and addresses assignment. Although it does not require to be configured to operate properly, it supports custom configuration by means of variable priority assignments, and can therefore be used in fully autonomic as well as professionally managed networks.

Given a set of delegated prefixes, Nodes may assign available prefixes to links they are directly connected to, or for their private use. The algorithm ensures that the following assertions are satisfied after a finite convergence period:

1. At most one prefix from each delegated prefix is assigned to each link.
2. Assigned prefixes are non-overlapping (i.e., an assigned prefix never includes another assigned prefix).

3. Assigned prefixes do not change in the absence of topology or configuration changes.

In the rest of this document the two first conditions are referred to as the correctness conditions of the algorithm while the third condition is referred to as its convergence condition.

Each assignment has a priority specified by the Node making the assignment, allowing for custom assignment policies. When multiple Nodes assign different prefixes from the same delegated prefix to the same link, or when multiple Nodes assign overlapping prefixes (to the same link or to different links), the assignment with the greatest priority is kept and other assignments are removed.

The prefix assignment algorithm requires that participating Nodes share information through a flooding mechanism. If the flooding mechanism ensures that all messages are propagated to all Nodes within a given time window, the algorithm also ensures that all assigned prefixes used for networking operations (e.g., host configuration) remain unchanged, unless another Node assigns an overlapping prefix with a higher assignment priority, or the topology changes and renumbering cannot be avoided.

## 2. Terminology

In this document, the key words "MAY", "MUST", "MUST NOT", "OPTIONAL", and "SHOULD", are to be interpreted as described in [RFC2119].

This document makes use of the following terminology. The terms defined here are ordered in such a way as to avoid forward references, and therefore are not sorted alphabetically.

**Node:** An entity executing the algorithm specified in this document and able to communicate with other Nodes using the Flooding Mechanism.

**Flooding Mechanism:** A mechanism allowing participating Nodes to reliably share information with all other participating Nodes.

**Link:** An object the distributed algorithm will assign prefixes to. A Node may only assign prefixes to Links it is directly connected to. A Link is either Shared or Private.

**Shared Link:** A Link multiple Nodes may be connected to. Most of the time, a Shared Link is a multi-access link or point-to-point link, virtual or physical, requiring prefixes to be assigned to it.

**Private Link:** A Private Link is an abstract concept defined for the sake of this document. It allows Nodes to make assignments for their private use or delegation. For instance, every DHCPv6-PD [RFC3633] requesting router MAY be considered as a different Private Link.

**Delegated Prefix:** A prefix provided to the algorithm and used as a prefix pool for Assigned Prefixes.

**Node ID:** A value identifying a given participating Node. The set of identifiers MUST be strictly and totally ordered (e.g., using the alphanumeric order).

**Flooding Delay:** A value which MUST be provided by the Flooding Mechanism and SHOULD be a deterministic or likely upper bound on the information propagation delay among participating Nodes.

**Advertised Prefix:** A prefix advertised by another Node and delivered to the local Node by the Flooding Mechanism. It has an Advertised Prefix Priority and, when assigned to a directly connected Shared Link, is associated with that Shared Link.

**Advertised Prefix Priority:** A value that defines the priority of an Advertised Prefix received from the Flooding Mechanism or a published Assigned Prefix. Whenever multiple Advertised Prefixes are conflicting (i.e., overlapping or from the same Delegated Prefix and assigned to the same link), all Advertised Prefixes but the one with the greatest priority will eventually be removed. In case of a tie, the assignment advertised by the Node with the greatest Node ID is kept and others are removed. In order to ensure convergence, the range of priority values MUST have an upper bound.

**Assigned Prefix:** A prefix included in a Delegated Prefix and assigned to a Shared or Private Link. It represents a local decision to assign a given prefix from a given Delegated Prefix to a given Link. The algorithm ensures that there is never more than one Assigned Prefix per Delegated Prefix and Link pair. When destroyed, an Assigned Prefix is set as not applied, ceases to be advertised, and is removed from the set of Assigned Prefixes.

**Applied (Assigned Prefix):** When an Assigned Prefix is applied, it MAY be used (e.g., for host configuration, routing protocol configuration, prefix delegation). When not applied, it MUST NOT be used for any purpose outside of the prefix assignment algorithm. Each Assigned Prefix is associated with a timer (Apply Timer) used to apply the Assigned Prefix. An Assigned Prefix is unapplied when destroyed.

**Published (Assigned Prefix):** The Assigned Prefix is advertised through the Flooding Mechanism as assigned to its associated Link. A published Assigned Prefix **MUST** have an Advertised Prefix Priority. It will appear as an Advertised Prefix to other Nodes, once received through the Flooding Mechanism.

**Prefix Adoption:** When an Advertised Prefix which does not conflict with any other Advertised Prefix or published Assigned Prefix stops being advertised, any other Node connected to the same Link **MAY**, after some random delay, start advertising the same prefix. This procedure is called adoption and provides seamless assignment transfer from a Node to another, e.g., in case of Node failure.

**Backoff Timer:** Every Delegated Prefix and Link pair is associated with a timer counting down to zero. It is used to reduce the probability of colliding assignments made by multiple Nodes by delaying the creation of new Assigned Prefixes or the advertisement of adopted Assigned Prefixes by a random amount of time.

**Renumbering:** Event occurring when an Assigned Prefix which was applied is destroyed. Renumbering is undesirable as it usually implies reconfiguring routers or hosts.

## 2.1. Subroutine Specific Terminology

In addition to the terms defined in Section 2, the subroutine specified in Section 4 makes use of the following terms.

**Current Assignment:** For a given Delegated Prefix and Link, the Current Assignment is the Assigned Prefix (if any) included in the Delegated Prefix and assigned to the given Link by the Node executing the algorithm. At some point in time, the Current Assignment from different Nodes may differ, but the algorithm ensures that eventually, all Nodes directly connected to a Shared Link have the same Current Assignment for any given Delegated Prefix.

**Precedence:** An Advertised Prefix takes precedence over an Assigned Prefix if and only if one of the following conditions is met:

- \* The Assigned Prefix is not published.
- \* The Assigned Prefix is published and the Advertised Prefix Priority from the Advertised Prefix is strictly greater than the Advertised Prefix Priority from the Assigned Prefix.

- \* The Assigned Prefix is published, the priorities are identical, and the Node ID from the Node advertising the Advertised Prefix is strictly greater than the local Node ID.

Best Assignment: For a given Delegated Prefix and Link, the Best Assignment is the unique Advertised Prefix (if any) that:

- \* Includes or is included in the Delegated Prefix (i.e., the Advertised Prefix is a sub-prefix of the Delegated Prefix, or the Delegated Prefix is a sub-prefix of the Advertised Prefix).
- \* Is assigned on the given Link.
- \* Has the greatest Advertised Prefix Priority among Advertised Prefixes fulfilling the two preceding conditions (and, in case of a tie, the prefix advertised by the Node with the greatest Node ID among all prefixes with greatest priority).
- \* Takes precedence over the Current Assignment associated with the same Link and Delegated Prefix (if any).

Valid (Assigned Prefix): An Assigned Prefix is valid if and only if the following two conditions are met:

- \* No Advertised Prefix including or included in the Assigned Prefix takes precedence over the Assigned Prefix.
- \* No Advertised Prefix including or included in the same Delegated Prefix as the Assigned Prefix and assigned to the same Link takes precedence over the Assigned Prefix.

### 3. Applicability Statement

Each Node MUST have a set of non-overlapping Delegated Prefixes (i.e., which do not include each other). This set MAY change over time and be different from one Node to another at some point, but Nodes MUST eventually have the same set of disjoint Delegated Prefixes.

Given this set of disjoint Delegated Prefixes, Nodes may assign available prefixes from each Delegated Prefix to the Links they are directly connected to. The algorithm ensures that at most one prefix from a given Delegated Prefix is assigned to any given Link.

The algorithm can be applied to any address space and can be used to manage multiple address spaces simultaneously. For instance, an implementation can make use of IPv4-mapped IPv6 addresses [RFC4291]

in order to manage both IPv4 and IPv6 prefix assignment using a single prefix space.

The algorithm supports dynamically changing topologies:

- o Nodes may join or leave the set of participating Nodes.
- o Nodes may join or leave Links.
- o Links may be joined or split.

All Nodes MUST run a common Flooding Mechanism in order to share published Assigned Prefixes. The set of participating Nodes is defined as the set of Nodes participating in the Flooding Mechanism.

The Flooding Mechanism MUST:

- o Provide a way to flood Assigned Prefixes assigned to a directly connected Link along with their respective Advertised Prefix Priority and the Node ID of the Node which is advertising them.
- o Specify whether an Advertised Prefix was assigned to a directly connected Shared Link, and if so, on which one.
- o Provide a Flooding Delay value, which SHOULD represent a deterministic or likely upper bound on the information propagation delay among participating Nodes. Whenever the Flooding Mechanism is unable to adhere to the provided Flooding Delay, renumbering may happen. As such a delay often depends on the size of the network, it MAY change over time and MAY be different from one Node to another. Furthermore, the process of selecting this value is subject to a tradeoff between convergence speed and lower renumbering probability (e.g., the value 0 may be used when renumbering is harmless), and is therefore out of scope of this document.

The algorithm ensures that whenever the Flooding Delay is provided and respected, and in the absence of any topology change or Delegated Prefix removal, renumbering only happens when a Node deliberately overrides an existing assignment.

Each Node MUST have a Node ID. Node IDs MAY change over time and be the same on multiple Nodes at some point, but each Node MUST eventually have a Node ID which is unique among the set of participating Nodes.

#### 4. Algorithm Specification

This section specifies the behavior of Nodes implementing the prefix assignment algorithm. The terms 'Current Assignment', 'Precedence', 'Best Assignment' and 'Valid' are used as defined in Section 2.1.

##### 4.1. Prefix Assignment Algorithm Subroutine

This section specifies the prefix assignment algorithm subroutine. It is defined for a given Delegated Prefix and Link pair and takes a BackoffTriggered boolean as parameter (indicating whether the subroutine execution was triggered by the Backoff Timer or by another event).

For a given Delegated Prefix and Link pair, the subroutine MUST be run with the BackoffTriggered boolean set to false whenever:

- o An Advertised Prefix including or included in the considered Delegated Prefix is added or removed.
- o An Assigned Prefix included in the considered Delegated Prefix and associated with a different Link than the considered Link was destroyed, while there is no Current Assignment associated with the given pair. This case MAY be ignored if the creation of a new Assigned Prefix associated with the considered pair is not desired.
- o The considered Delegated Prefix is added.
- o The considered Link is added.
- o The Node ID is modified.

Furthermore, for a given Delegated Prefix and Link pair, the subroutine MUST be run with the BackoffTriggered boolean set to true whenever:

- o The Backoff Timer associated with the considered Delegated Prefix and Link pair fires while there is no Current Assignment associated with the given pair.

When such an event occurs, a Node MAY delay the execution of the subroutine instead of executing it immediately, e.g., while receiving an update from the Flooding Mechanism, or for security reasons (see Section 8). Even if other events occur in the meantime, the subroutine MUST be run only once. It is also assumed that, whenever one of these events is the Backoff Timer firing, the subroutine is executed with the BackoffTriggered boolean set to true.



In order to execute the subroutine for a given Delegated Prefix and Link pair, first look for the Best Assignment and Current Assignment associated with the Delegated Prefix and Link pair, then execute the corresponding case:

1. If there is no Best Assignment and no Current Assignment: Decide whether the creation of a new assignment for the given Delegated Prefix and Link pair is desired (As any result would be valid, the process of making this decision is out of the scope of this document) and do the following:

- \* If it is not desired, stop the execution of the subroutine.
- \* Else if the Backoff Timer is running, stop the execution of the subroutine.
- \* Else if the BackoffTriggered boolean is set to false, set the Backoff Timer to some random delay between ADOPT\_MAX\_DELAY and BACKOFF\_MAX\_DELAY (see Section 7) and stop the execution of the subroutine.
- \* Else, continue the execution of the subroutine.

Select a prefix for the new assignment (see Section 5 for guidance regarding prefix selection). This prefix MUST be included in or be equal to the considered Delegated Prefix and MUST NOT include or be included in any Advertised Prefix. If a suitable prefix is found, use it to create a new Assigned Prefix:

- \* Assigned to the considered Link.
- \* Set as not applied.
- \* The Apply Timer set to '2 \* Flooding Delay'.
- \* Published with some selected Advertised Prefix Priority.

2. If there is a Best Assignment but no Current Assignment: Cancel the Backoff Timer and use the prefix from the Best Assignment to create a new Assigned Prefix:

- \* Assigned to the considered Link.
- \* Set as not applied.
- \* The Apply Timer set to '2 \* Flooding Delay'.
- \* Set as not published.

3. If there is a Current Assignment but no Best Assignment:
  - \* If the Current Assignment is not valid, destroy it, and execute the subroutine again with the BackoffTriggered boolean set to false.
  - \* If the Current Assignment is valid and published, stop the execution of the subroutine.
  - \* If the Current Assignment is valid and not published, the Node MUST either:
    - + Adopt the prefix by canceling the Apply Timer and set the Backoff Timer to some random delay between 0 and ADOPT\_MAX\_DELAY (see Section 7). This procedure is used to avoid renumbering when the Node advertising the prefix left the Shared Link.
    - + Destroy it and go to case 1.
4. If there is a Current Assignment and a Best Assignment:
  - \* Cancel the Backoff Timer.
  - \* If the two prefixes are identical, set the Current Assignment as not published. If the Current Assignment is not applied and the Apply Timer is not set, set the Apply Timer to '2 \* Flooding Delay'.
  - \* If the two prefixes are not identical, destroy the Current Assignment and go to case 2.

When the prefix assignment algorithm subroutine requires an assignment to be created or adopted, any Advertised Prefix Priority value can be used. Other documents MAY provide restrictions over this value depending on the context the algorithm is operating in, or leave it as implementation-specific.

#### 4.2. Overriding and Destroying Existing Assignments

In addition to the behaviors specified in Section 4.1, the following procedures MAY be used in order to provide additional behavior options (Section 6):

Overriding Existing Assignments: For any given Link and Delegated Prefix, a Node MAY create a new Assigned Prefix using a chosen prefix and Advertised Prefix Priority such that:

- \* The chosen prefix is included in or is equal to the considered Delegated Prefix.
- \* The Current Assignment, if any, as well as all existing Assigned Prefixes which include or are included inside the chosen prefix, are destroyed.
- \* It is not applied.
- \* The Apply Timer is set to ' $2 * \text{Flooding Delay}$ '.
- \* It is published.
- \* The Advertised Prefix Priority is greater than the Advertised Prefix Priority from all Advertised Prefixes which include or are included in the chosen prefix.
- \* The Advertised Prefix Priority is greater than the Advertised Prefix Priority from all Advertised Prefixes which include or are included in the considered Delegated Prefix and are assigned to the considered Link.

In order to ensure algorithm convergence:

- \* Such overriding assignments MUST NOT be created unless there was a change in the Node configuration, a Link was added, or an Advertised Prefix was added or removed.
- \* The chosen Advertised Prefix Priority for the new Assigned Prefix SHOULD be greater than all priorities from the destroyed Assigned Prefixes. If not, simple topologies with only two Nodes may not converge. Nodes which do not adhere to this rule MUST implement a mechanism which detects whether the distributed algorithm does not converge and, whenever this would happen, stop creating overriding Assigned Prefixes which do not adhere to this rule. The specifications for such safety procedures are out of the scope of this document.

Removing an Assigned Prefix: A Node MAY destroy any Assigned Prefix which is published. Such an event reflects the desire of a Node to not assign a prefix from a given Delegated Prefix to a given Link anymore. In order to ensure algorithm convergence, such a procedure MUST NOT be executed unless there was a change in the Node configuration. Furthermore, whenever an Assigned Prefix is destroyed in this way, the prefix assignment algorithm subroutine MUST be run for the Delegated Prefix and Link pair associated with the destroyed Assigned Prefix.

The two procedures specified in this section are OPTIONAL. They could be used for various purposes, e.g., for providing custom prefix assignment configuration or reacting to prefix space exhaustion (by overriding short Assigned Prefixes and assigning longer ones).

#### 4.3. Other Events

When the Apply Timer fires, the associated Assigned Prefix MUST be applied.

When the Backoff Timer associated with a given Delegated Prefix and Link pair fires while there is a Current Assignment associated with the same pair, the Current Assignment MUST be published with some associated Advertised Prefix Priority and, if the prefix is not applied, the Apply Timer MUST be set to '2 \* Flooding Delay'.

When a Delegated Prefix is removed from the set of Delegated Prefixes (e.g., when the Delegated Prefix expires), all Assigned Prefixes included in the removed Delegated Prefix MUST be destroyed.

When one Delegated Prefix is replaced by another one that includes or is included in the deleted Delegated Prefix, all Assigned Prefixes which were included in the deleted Delegated Prefix but are not included in the added Delegated Prefix MUST be destroyed. Others MAY be kept.

When a Link is removed, all Assigned Prefixes assigned to that Link MUST be destroyed.

#### 5. Prefix Selection Considerations

When the prefix assignment algorithm subroutine specified in Section 4.1 requires a new prefix to be selected, the prefix MUST be selected either:

- o Among prefixes included in the considered Delegated Prefix which were previously assigned and applied on the considered Link. For that purpose, Applied Prefixes may be stored in stable storage along with their associated Link.
- o Randomly, picked in a set of at least RANDOM\_SET\_SIZE (see Section 7) prefixes included in the considered Delegated Prefix and not including or included in any Assigned or Advertised Prefix. If less than RANDOM\_SET\_SIZE candidates are found, the prefix MUST be picked among all candidates.
- o Based on some custom selection process specified in the configuration.

A simple implementation MAY randomly pick the prefix among all available prefixes, but this strategy is inefficient in terms of address space use as a few long prefixes may exhaust the pool of available short prefixes.

The rest of this section describes a more efficient approach which MAY be applied any time a Node needs to pick a prefix for a new assignment. The two following definitions are used:

**Available prefix:** The prefix of the form Prefix/PrefixLength is available if and only if it satisfies the three following conditions:

- \* It is included in the considered Delegated Prefix.
- \* It does not include and is not included in any Assigned or Advertised Prefix.
- \* It is equal to the considered Delegated Prefix or Prefix/(PrefixLength-1) includes an Assigned or Advertised Prefix.

**Candidate prefix:** A prefix of desired length which is included in or is equal to an available prefix.

The procedure described in this section takes the three following criteria into account:

**Prefix Stability:** In some cases, it is desirable that the selected prefix should remain the same across executions and reboots. For this purpose, prefixes previously applied on the Link or pseudo-random prefixes generated based on Node- and Link-specific values may be considered.

**Randomness:** When no stored or pseudo-random prefix is chosen, a prefix may be randomly picked among RANDOM\_SET\_SIZE candidates of desired length. If less than RANDOM\_SET\_SIZE candidates can be found, the prefix is picked among all candidates.

**Addressing-space usage efficiency:** In the process of assigning prefixes, a small set of badly chosen long prefixes may prevent any shorter prefix from being assigned. For this reason, the set of RANDOM\_SET\_SIZE candidates is created from available prefixes with longest prefix lengths and, in case of a tie, preferring numerically small prefix values.

When executing the procedure, do as follows:

1. For each prefix stored in stable storage, check if the prefix is included in or equal to an available prefix. If so, pick that prefix and stop.
2. For each prefix length, count the number of available prefixes of the given length.
3. If the desired prefix length was not specified, select one. The available prefixes count computed previously may be used to help pick a prefix length such that:

- \* There is at least one candidate prefix.
- \* The prefix length is chosen large enough to not exhaust the address space.

Let N be the chosen prefix length.

4. Iterate over available prefixes starting with prefixes of length N down to length 0 and create a set of RANDOM\_SET\_SIZE candidate prefixes of length exactly N included in or equal to available prefixes. The end goal here is to create a set of RANDOM\_SET\_SIZE candidate prefixes of length N included in a set of available prefixes of maximized prefix length. In case of a tie, smaller prefix values (as defined by the bit-wise lexicographical order) are preferred.
5. Generate a set of prefixes of desired length, which are pseudo-randomly chosen based on Node- and Link-specific values. For each pseudo-random prefix, check if the prefix is equal to a candidate prefix. If so, pick that prefix and stop.
6. Choose a random prefix from the set of selected candidates.

The complexity of this procedure is equivalent to the complexity of iterating over available prefixes. Such operation may be accomplished in linear time, e.g., by storing Advertised and Assigned Prefixes in a binary trie.

## 6. Implementation Capabilities and Node Behavior

Implementations of the prefix assignment algorithm may vary from very basic to highly customizable, enabling different types of fully interoperable behaviors. The three following behaviors are given as examples:

**Listener:** The Node only acts upon assignments made by other Nodes, i.e., it never creates new assignments nor adopts existing ones.

Such behavior does not require the implementation of the considerations specified in Section 5 or Section 4.2. The Node never checks the validity of existing assignments, which makes this behavior particularly suited to lightweight devices which can rely on more capable neighbors to make assignments on directly connected Shared Links.

**Basic:** The Node is capable of assigning new prefixes or adopting prefixes which do not conflict with any other existing assignment. Such behavior does not require the implementation of the considerations specified in Section 4.2. It is suited to situations where there is no preference over which prefix should be assigned to which Link, and there is no priority between different Links.

**Advanced:** The Node is capable of assigning new prefixes, adopting existing ones, making overriding assignments and destroying existing ones. Such behavior requires the implementation of the considerations specified in Section 5 and Section 4.2. It is suited when the administrator desires some particular prefix to be assigned on a given Link, or some Link to be assigned prefixes with a greater priority when there are not enough prefixes available for all Links.

Note that if all Nodes directly connected to some Link are listener Nodes or none of these Nodes are willing to make an assignment from a given Delegated Prefix to the given Link, no prefix from the given Delegated Prefix will ever be assigned to the Link (and such existing prefixes will be removed). This situation may be detected by watching whether no prefix from a given Delegated Prefix has been assigned to the Link for longer than BACKOFF\_MAX\_DELAY plus the Flooding Delay.

## 7. Algorithm Parameters

This document does not provide values for ADOPT\_MAX\_DELAY, BACKOFF\_MAX\_DELAY and RANDOM\_SET\_SIZE. The algorithm ensures convergence and correctness for any chosen values, even when these are different from Node to Node. They MAY be adjusted depending on the context, providing a tradeoff between convergence time, efficient addressing, reduced control traffic (generated by the Flooding Mechanism), and low collision probability.

ADOPT\_MAX\_DELAY (respectively BACKOFF\_MAX\_DELAY) represents the maximum backoff time a Node may wait before adopting an assignment (respectively making a new assignment). BACKOFF\_MAX\_DELAY MUST be greater than or equal to ADOPT\_MAX\_DELAY. The greater ADOPT\_MAX\_DELAY and (BACKOFF\_MAX\_DELAY - ADOPT\_MAX\_DELAY), the lower

the collision probability and the lesser the amount of control traffic, but the greater the convergence time.

RANDOM\_SET\_SIZE represents the desired size of the set a random prefix will be picked from. The greater RANDOM\_SET\_SIZE, the better the convergence time and the lower the collision probability, but the worse the addressing-space usage efficiency.

## 8. Security Considerations

The prefix assignment algorithm functions on top of two distinct mechanisms, the Flooding Mechanism and the Node ID assignment mechanism.

An attacker able to publish Advertised Prefixes through the Flooding Mechanism may perform the following attacks:

- \* Publish a single overriding assignment for a whole Delegated Prefix or for the whole address space, thus preventing any Node from assigning prefixes to Links.
- \* Quickly publish and remove Advertised Prefixes, generating traffic at the Flooding Mechanism layer and causing multiple executions of the prefix assignment algorithm in all participating Nodes.
- \* Publish and remove Advertised Prefixes in order to prevent the convergence of the algorithm.

An attacker able to prevent other Nodes from accessing a portion or the whole set of Advertised Prefixes may compromise the correctness of the algorithm.

An attacker able to cause repetitive Node ID changes may cause traffic to be generated in the Flooding Mechanism and multiple executions of the prefix assignment algorithm in all participating Nodes.

An attacker able to publish Advertised Prefixes using a Node ID used by another Node may prevent the correctness and convergence of the algorithm or cause the result to violate the correctness conditions.

Whenever the security of the Flooding Mechanism and Node ID assignment mechanism cannot be ensured, the convergence of the algorithm may be prevented. In environments where such attacks may be performed, the execution of the prefix assignment algorithm subroutine SHOULD be rate limited, as specified in Section 4.1.



## 9. IANA Considerations

This document has no actions for IANA.

## 10. Acknowledgments

The authors would like to thank those who participated in the previous document's version development as well as the present one. In particular, the authors would like to thank Tim Chown, Fred Baker, Mark Townsley, Lorenzo Colitti, Ole Troan, Ray Bellis, Markus Stenberg, Wassim Haddad, Joel Halpern, Samita Chakrabarti, Michael Richardson, Anders Brandt, Erik Nordmark, Laurent Toutain, Ralph Droms, Acee Lindem and Steven Barth for interesting discussions and document review.

## 11. References

### 11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

### 11.2. Informative References

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

## Appendix A. Static Configuration Example

This section describes an example of how custom configuration of the prefix assignment algorithm may be implemented.

The Node configuration is specified as a finite set of rules. A rule is defined as:

- o A prefix to be used.
- o A Link on which the prefix may be assigned.
- o An Assigned Prefix Priority (smallest possible Assigned Prefix Priority if the rule may not override other Assigned Prefixes).
- o A rule priority (0 if the rule may not override existing Advertised Prefixes).

In order to ensure the convergence of the algorithm, the Assigned Prefix Priority MUST be an increasing function (not necessarily strictly) of the configuration rule priority (i.e., the greater is the configuration rule priority, the greater the Assigned Prefix Priority must be).

Each Assigned Prefix is associated with a rule priority. Assigned Prefixes which are created as specified in Section 4.1 are given a rule priority of 0.

Whenever the configuration is changed or the prefix assignment algorithm subroutine is run: For each Link/Delegated Prefix pair, look for the configuration rule with the greatest configuration rule priority such that:

- o The prefix specified in the configuration rule is included in the considered Delegated Prefix.
- o The Link specified in the configuration rule is the considered Link.
- o All the Assigned Prefixes which would need to be destroyed in case a new Assigned Prefix is created from that configuration rule (as specified in Section 4.2) have an associated rule priority which is strictly lower than the one of the considered configuration rule.
- o The assignment would be valid when published with an Advertised Prefix Priority equal to the one specified in the configuration rule.

If a rule is found, a new Assigned Prefix is created based on that rule as specified in Section 4.2. The new Assigned Prefix is associated with the Advertised Prefix Priority and the rule priority specified in the considered configuration rule.

Note that the use of rule priorities ensures the convergence of the algorithm.

#### Authors' Addresses

Pierre Pfister  
Cisco Systems  
Paris  
France

Email: pierre.pfister@darou.fr

Benjamin Paterson  
Cisco Systems  
Paris  
France

Email: paterson.b@gmail.com

Jari Arkko  
Ericsson  
Jorvas 02420  
Finland

Email: jari.arkko@piuha.net